

IMS
Version 11

Communications and Connections



IMS
Version 11

Communications and Connections



Note

Before using this information and the product that it supports, be sure to read the general information under “Notices” on page 963.

This edition applies to IMS Version 11 (program number 5635-A02) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1974, 2015.**

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

Contents

About this information xv

Prerequisite knowledge	xv
IMS function names used in this information	xv
How new and changed information is identified	xvi
How to read syntax diagrams	xvi
Accessibility features for IMS Version 11	xviii
How to send your comments	xviii

Changes to the IMS library xix

Part 1. Configuring external Java environment connections 1

Chapter 1. IMS Universal drivers: configuring connections to IMS 3

IMS Universal drivers: WebSphere Application Server type-4 connections	3
Installing a type-4 IMS Universal Database resource adapter on WebSphere Application Server	5
Defining a connection factory for a type-4 IMS Universal Database resource adapter on WebSphere Application Server	5
Installing an EAR file that uses a type-4 IMS Universal Database resource adapter on WebSphere Application Server	6
IMS Universal drivers: WebSphere Application Server for z/OS type-2 connections	6
Configuring WebSphere Application Server for z/OS for a type-2 IMS Universal Database resource adapter	7
Installing a type-2 IMS Universal Database resource adapter on WebSphere Application Server for z/OS	8
Optional: set the WebSphere Application Server for z/OS classpath for applications that use a type-2 IMS Universal Database resource adapter	9
Defining a connection factory for a type-2 IMS Universal Database resource adapter on WebSphere Application Server for z/OS	10
Installing an EAR file that uses a type-2 IMS Universal Database resource adapter on WebSphere Application Server for z/OS	11
The IMS Universal drivers: CICS connections	12
Configuring CICS for the type-2 IMS Universal drivers	12
Running applications on CICS that use the type-2 IMS Universal drivers	13

Chapter 2. Classic Java APIs for IMS: configuring connections to IMS 15

Classic Java APIs for IMS: configuring CICS	15
Configuring CICS for the classic Java APIs for IMS	16

Running the IMS IVP application for the classic Java APIs for IMS and CICS	17
Running your applications on CICS with the classic Java APIs for IMS	18
Classic Java APIs for IMS: configuring DB2 for z/OS	19
Configuring DB2 for z/OS for the classic Java APIs for IMS	20
Running the IMS IVP applications for the classic Java APIs for IMS and DB2 for z/OS	21
Running your stored procedure with the classic Java APIs for IMS	24
Classic Java APIs for IMS: configuring WebSphere Application Server for z/OS	26
Configuring WebSphere Application Server for z/OS for the classic Java APIs for IMS	27
Running the IMS IVP applications in the WebSphere Application Server for z/OS Java environment	29
Running your applications on WebSphere Application Server for z/OS	32

Part 2. CPI Communications and APPC/IMS 35

Chapter 3. CPI Communications 37

CPI-C driven application programs	37
SAA resource recovery commit processing	37
Normal termination	37
Backout processing	38
Abnormal termination	38
Session failure	38
Return codes	38
System restart/resolve-in-doubt processing	39
CPI-C application program recovery	39
Programming requirements	40
Pseudonym files	40
RRS and distributed syncpoint/protected conversations	40
The two-phase commit protocol	41
Local-resource recovery versus distributed-resource recovery	42
IMS as a resource manager	43
Activating protected conversations	44

Chapter 4. Administering APPC/IMS and LU 6.2 devices 47

APPC/IMS overview	47
APPC/IMS application program interface	48
APPC/IMS application programs	49
Standard IMS application programs	49
MSC and standard IMS application programs	50
Modified IMS application programs	51
MSC and modified IMS application programs	51

CPI Communications driven application programs	52
Using the MOD name and LTERM interface	53
Establishing APPC/IMS	53
TP_Profile	54
APPC/MVS Administration utility (ATBSDFMU) example	55
Outbound LU specification	57
Outbound side information	57
PARMLIB member	58
APPC/MVS Timeout Service	59
APPC/MVS Error Extract Service	60
Initializing and changing LU 6.2 descriptors	60
Using MSC in an APPC/IMS environment	61
Recovering APPC transactions in an MSC environment	62
Recoverable versus nonrecoverable transactions	62
Local APPC transaction discardability versus nondiscardability	63
Transaction processing point of failure	63
Recoverability flows of LU 6.2 transactions	66
Planning for XRF and APPC.	67
Transaction retry characteristics.	68
Qualifying network LU names	68
Managing multiple LUs for a single IMS system	69
Reassigning an LU to another IMS system	70
DFSAPPC system service.	70
Message switching	70
Asynchronous output delivery	73
APPC transaction security	74

Part 3. Extended Terminal Option (ETO) 77

Chapter 5. Overview of the Extended Terminal Option 79

ETO terminology	79
Descriptores	81
ETO concepts.	83
Descriptores and exit routines	84
How descriptors are created and used	84
Summary of ETO implementation	84

Chapter 6. Administering the Extended Terminal Option 87

Planning for ETO	87
Identifying your requirements	88
ETO restrictions	88
Defining physical terminals	89
Identifying VTAM device types, screen sizes, and models	92
Planning a high-security environment with ETO	98
Planning for MFS	98
Planning user IDs	99
Planning user queue names	99
Planning operations	99
Planning for MSC support with ETO	100
Coding ETO descriptors.	100

Creating descriptors using the system definition process	101
Storing descriptors	101
Creating logon descriptors	101
Creating user descriptors	104
Creating MFS device descriptors	107
Creating MSC descriptors	109
Exit routines.	110
Starting ETO	110
Logging onto ETO terminals	111
Limiting dynamic logon to specific terminal types	111
Creating and reusing LTERM control blocks	112
Using default CINIT or BIND user data formats	112
Signing on and queue LTERM allocation	113
Providing signon data	113
Providing signon data for ISC, SLU-P, Finance, and output-only devices.	113
Signing on multiple times	113
Receiving DFS3649A, the signon required message	115
Receiving DFS3650I, the session status message	115
ETO terminal-LTERM relationship	115
How IMS determines which queues to allocate	116
Setting special processing modes	116
Printers with ETO	117
Direct printing	118
Associated printing	118
Defining your printers	119
Sharing printers using ETO.	120
Operator commands	120
System definition parameters for ETO	121
Setting DEADQ status time with the DLQT parameter	121
Autosignoff (ASOT)	122
Autologoff (ALOT)	123
Autosignoff and autologoff timer.	124
Autologon	125
Assigning output	125
Asynchronous output	125
Delivering output messages to non-originating terminals.	127
Inadvertent output data streams	127
Signing off	127
Logging off	128
Improving performance by deleting ETO control blocks	128
IDC0 Trace facility.	129
ETO and LU 6.1 (ISC) terminals	129
ETO and STSN terminals	130
SNA STSN terminal considerations	130
ETO and 3600/Finance and SLU P	130
/SIGN support for ETO STSN devices: ISC, Finance, and SLU P	131
Conversation mode and response mode with ETO	131

Part 4. External subsystem attach facilities 135

Chapter 7. DB2 Attach Facility 137

Preparing your system to use the DB2 Attach Facility	137
Managing how your Java dependent regions access DB2 for z/OS	137

Chapter 8. External Subsystem Attach Facility (ESAF) 139

What the external subsystem must provide	140
How external subsystems are specified to IMS	141
The basics of attach processing	142
Subsystem connections	142
Application call processing	145
Resource coordination	145
External subsystem command support	146
IMS services available to the ESAP	147

Chapter 9. Creating the external subsystem module table 149

DFSEMODL macro	149
DFSEWAL macro	152

Chapter 10. IMS External Subsystem Attach Facility processing 157

Loading the External Subsystem Attachment Package	157
Creating the EEVT control block	157
Loading external subsystem modules	158
Creating work areas for the ESAP	159
Initiating the external subsystem connection	159
Deferring the control region identify	160
Using the IMS Subsystem Startup Service	160
Establishing dependent region connections	161
Notify message	161
Application program request support	163
Language interface definition	163
Language interface entry points unique to external subsystems	163
Accessing multiple external subsystems	164
Resource recovery token	164
Terminating the external subsystem connection	165
Termination requested by the external subsystem	166
Dependent region connections	166
Explanation of stopped status	167

Part 5. IMS Connect and TCP/IP communications 169

Chapter 11. Overview of IMS Connect 171

IMS Connect client support	172
IMS Connect support for access to IMS DB	176
IMS Connect support for the IMS TM Resource Adapter	177
Overview of IMS Connect XML Conversion Support	178
IMS Connect support for z/OS Sysplex Distributor	179
Overview of IMS Connect security	179

Overview of defining and invoking IMS Connect	180

Chapter 12. Overview of IMS Connect exit routines 183

Overview of user message exit routines	184
Security and the IMS Connect user message exit routines	185
User-defined messages	185
Overview of function-specific exit routines	186
Macros that support IMS Connect exit routines	186
Exit interface blocks	188
XIB exit interface block for connections to IMS TM	188
XIBDS exit interface block for IMS TM data store information	189
XIB1 exit interface block for connections to IMS DB	191
XIBOD exit interface block for ODBM and IMS DB data store information	192

Chapter 13. IMS Connect support for IMSplex and shared queues 195

IMS Connect support for IMSplex	195
IMSplex support environment	195
Installing IMS Connect support for IMSplex	196
Retrieving ALTPCB output in a shared queues environment	197

Chapter 14. IMS Connect security support 199

IMS Connect support for RACF	199
Enabling RACF security checking in IMS Connect	200
IMS Connect default RACF user ID	200
IMS Connect security for clients of IMS DB	201
IMS Connect security exit routine	201
OTMA Resume TPIPE Security exit routine (DFSYRTUX)	202
HWSSMPL0 and HWSSMPL1 security actions	202
IMS Connect responses to errors on RACROUTE calls from the sample exits	204
IMS Connect password management	208
Changing RACF passwords by using client messages	208
Enabling mixed-case password support	209
IMS Connect support for RACF PassTicket	210
IMS Connect SSL connections	212
z/OS key management	214
SSL initialization	214
SSL default setup	218
Sample JCL for RACF-Managed SSL	218
Trusted-user support for IMS Connect messages	219
Specifying an OTMA ACEE aging value in the IMS Connect configuration member	220

Chapter 15. IMS Connect support for callout requests 223

Configuring user-written IMS Connect clients for synchronous callout requests	223

	Format of synchronous callout messages	225
	Retrieving synchronous callout requests with RESUME TPIPE	226
	Acknowledging receipt of synchronous callout messages	228
	Returning callout responses to IMS	229
	Returning an error response to IMS	230

Chapter 16. IMS Connect XML

message conversion 231

IMS Connect XML converters	232
Structure of the XML message	232
Message conversion example	234

Chapter 17. IMS Connect message

structures 237

IRM structures for IMS Connect client messages	237
Format of fixed portion of IRM in messages sent to IMS Connect.	238
Format of user portion of IRM for HWSSMPL0, HWSSMPL1, and user-written message exit routines	241
Output structure from client exit	248
Other IMS Connect structures	249
Message structures and IMS Connect user message exit routines	250
Input messages from client	251
Output message to client	251
Message structures	252
Examples of message structures in a simple interaction	262

Chapter 18. OTMA header fields used

by IMS Connect 265

OTMA message-control fields used by IMS Connect	265
OTMA state data fields used by IMS Connect	271
OTMA security data fields used by IMS Connect	281
OTMA user data fields used by IMS Connect.	283
Notes to OTMA header tables	289

Chapter 19. IMS Connect protocols 291

Transaction restrictions and limitations	291
Commit mode and synchronization level definitions	291
IMS Connect conversational support	292
OTMA conversational protocols	293
IMS Connect conversational protocols	295
Purging undeliverable commit-then-send output	299
Specifying the purge function for undeliverable commit-then-send output	300
When IMS purges undeliverable commit-then-send output	300
Rerouting commit-then-send output	301
Specifying the reroute function for commit-then-send output	302
Specifying a destination for rerouted output	302
When IMS reroutes commit-then-send output	303
Recoverable IMS transactions	304

Send-only protocol	305
Send-only with acknowledgment protocol	306
Send-only protocol with serial delivery protocol	307
Send-only protocol for synchronous callout responses.	307
Socket connections	308
Persistent sockets	308
Transaction sockets	309
Non-persistent sockets	309
Setting socket types for IMS TM clients.	309
Socket processing for transactions	310
Managing the number of sockets	311
Resolving duplicate client IDs	314
IMS Connect override for the z/OS TCP/IP	
KeepAlive interval	315
TCP/IP failures	316
IMS Connect timeout specifications	317
Timeout specifications for IMS DB clients	317
Timeout specifications for IMS TM clients	317
IMS Connect transaction expiration support	327
Setting a transaction expiration time with IMS Connect	328
Retrieval of output on OTMA tpipe hold queues	328
RESUME TPIPE/receive protocol.	329
Implementing asynchronous output support	332
Managing the retrieval of output messages	333
Retrieving output from alternate OTMA tpipe hold queues	339
Defining groups for shared asynchronous output.	340
Asynchronous output message flow	342
IMS Connect client call flows	342
IMS Connect dead letter queue (HWS\$DLQ)	346
Ping support for IMS Connect.	347

Chapter 20. IMS Connect two-phase

commit support 349

Overview of two-phase commit protocol	349
Distributed two-phase commit support.	350
Support for the IMS Universal drivers	351
Global (XA) transactions with the IMS Universal drivers	351
One-phase commit global transactions with the IMS Universal drivers	353
Support for the IMS TM Resource Adapter	356
Global (XA) transactions with IMS TM Resource Adapter	357
One-phase commit global transactions with IMS TM Resource Adapter	359
Two-phase commit support for IMS TM Resource Adapter Local Option	360

Chapter 21. Unicode considerations

for IMS Connect 363

Message translation	363

Chapter 22. TCP/IP settings for IMS Connect. 365

Part 6. IMS VTAM network administration 367

Chapter 23. Introduction to the IMS Transaction Manager network 369

IMS TM network overview	369
IMS Transaction Manager services	373
The Data Communication Control (DCCTL) environment.	375
Operating an IMS network	376
The shared-queues environment	377
Benefits of using shared queues	378
Required components of a shared-queues environment.	379
Overview of the Common Queue Server	380
Balancing sessions with generic resources	381
IMSplex terminal management	383
Benefits of managing resources with a resource structure	383
Shared TM resources	384
Resource name uniqueness	384
Resource type consistency	384
Fast Path expedited message handler	385

Chapter 24. Planning the network. 389

Planning for network administration	389
Documenting network and terminal requirements	389
IMS terminal network	390
Terminal connections to IMS	391
XRF terminal support	391
Logical terminals (LTERMs)	391
APPC/IMS and LU 6.2 terminal support	392
IMS messages and their scheduling	392
Message flow within the IMS online system	395
Conversational transactions.	396
Message switches	397
Designing logical terminal networks.	397
Logical-terminal chains	398
Logical-terminal queues	399
Separating input and output devices	399
Logical and physical terminal relationships	400
Master terminal	401
Master terminals in an XRF complex	401
NTO terminals	401
Resource modes and states	402
Terminal and user operating modes	402
Terminal and user states.	404
Resource status recovery	405
Planning for security	410
Authorizing transactions in a TM network.	411
Authorizing commands in a TM network	411
Transaction command security.	411
Password security	412
Security for APPC/IMS	412
Security for ETO	412
Planning for Fast Path terminals	413

Planning for the Extended Recovery Facility (XRF)	414
Using MNPS or USERVAR in XRF complexes	415
Terminals in an XRF complex	416
Specifying terminal support	418
Planning for Rapid Network Reconnect (RNR)	419
Specifying levels of support	420
Persistent Session Tracking	421
IMS shutdown and RNR	421
Using RNR with XRF or VGR	422
Terminal reconnect protocols	422
Signon security.	423

Chapter 25. Defining the network 425

Preparing for the operational network	425
Coordinating IMS definition and network definition.	426
Using IMS as a host subsystem	426
Defining VTAM nodes	426
Estimating VTAM storage requirements	427
Determining VTAM buffer pool values	427
Determining the NCP buffer pool values	427
Determining static and dynamic terminal signon requirements	427
Checking requirements for LOGON MODE tables	428
Specifying initial VTAM configurations.	430
Using SON/COS support in IMS.	430
Starting an IMS network	431
Session initiation	431
IMS transaction types and transaction states	432
Determining transaction states.	433
Restarting the network after remote takeover.	434
VTAM definition for Transport Manager Subsystem	435
Defining VTAM for Rapid Network Reconnect (RNR)	435
Defining the level of persistent support.	436
Defining the level of RNR support	436

Chapter 26. Editing and formatting IMS messages 437

Message Format Service	437
MFS components	438
Administering MFS	442
Advantages to using MFS	443
MFS control blocks	445
Creating MFS formats with SDF II	445
Basic edit.	447
IMS editing for Intersystem Communication (ISC)	448
Transparency option	449
Unprotected screen option	449
Bypassing MFS editing	449
Locking and unlocking the terminal keyboard	449
IMS sensitivity to nongraphic message data	450
Output message segment editing	450
Editing of input message segments by MFS	450
Editing of input message segments by basic edit	451
Controlling output devices	451
Small buffer devices	452
Controlling output	452
Using a printer component	453

Spooled output control	453
Using printer components of the IBM 3270	
Information Display System	454
Specifying candidate printers	454
Operational considerations	455
Sharing printers between systems	455

Part 7. Intersystem Communication (ISC) 457

Chapter 27. Overview of Intersystem Communication 459

Comparison of ISC and MSC	460
IMS facilities available to ISC	461
Sample system configurations	463
ISC between IMS and CICS.	466
Terminal device-dependent data	467
Passing CICS data to IMS	467

Chapter 28. VTAM facilities used for ISC connections 469

VTAM commands and indicators.	470
Using the VTAM application programming interface	471
Specifying logon modes when establishing a connection	471
Design considerations for secondary logical units	472

Chapter 29. IMS facilities affected by ISC 473

Editing messages	473
Issuing IMS commands from an ISC session	474
Effects on parallel sessions	474
Using IMS test mode for ISC VTAM sessions.	474
IMS control block storage on ISC parallel sessions	475
Relationship of ISC and IMS execution modes	475
External specification of execution modes	475
Internal definition of execution mode	476
Resultant processing mode during ISC VTAM communications	477
LTERM users (subpools) and components	478

Chapter 30. Designing communications using the ISC protocol. 481

Determining output protocols	481
Accessing existing application programs with ISC	482
Accessing programs that use MFS	482
Accessing programs that do not use MFS	483
Routing messages	484
Considerations for IMS-to-IMS ISC sessions	493
Statically defining an ISC node to IMS	496
Choosing parameters: system design considerations	498
COMM macro statement	498
NAME macro	499
SUBPOOL macro	499
TERMINAL macro	500
System definition summary	501

Chapter 31. ISC protocols for VTAM connections 503

Operating the network	503
Making IMS ready	503
Starting an IMS network for ISC	503
Shutting down an IMS network for ISC	503
Controlling the session (session control protocols)	504
Initiating an ISC VTAM session	504
Establishing a connection with an XRF complex	505
Binding the session	506
Resynchronizing sessions	507
Designing restart resynchronization procedures	509
Determining session synchronism using STSN	513
Performing the resynchronization	515
Completing session initiation	516
Running the session	516
Terminating an ISC VTAM session	516
Using STSN to resynchronize sessions	517
Primary-to-secondary flow matrix	518
Secondary-to-primary flow matrix	519
STSN command format	520
Handling IMS response mode or conversational output errors	523
Response mode errors	523
Conversational mode errors	523
Normal conversation termination extension with ISC.	524
Keeping half sessions synchronized	525
Sync points requested on input to IMS	525
Sync points requested on output by IMS	526
Sync point and response requirements	527
Sync-point indicators on messages	529
Data flow control protocol reference.	533
BID protocol.	533
Bracket and half-duplex protocol	533
CANCEL protocol.	541
Chaining protocol	542
CHASE protocol	542
ERP purging	542
LUSTATUS protocol	546
Paged messages ERP	549
Ready-to-receive protocol	549
RSHUT protocol	549
Selective receiver ERP	549
Sender ERP	554
Sense codes that IMS receives	557
Sense codes that IMS sends.	557
SIGNAL protocol	558
Symmetrical session shutdown for LU 6.1 (SBI and BIS)	558
Function management headers	559
Using FM headers to invoke ISC edit	561
Initiating a process: ATTACH FM header	561
Error recovery procedure FM header	562
Resetting the active process: RAP FM header	563
Requesting asynchronous transaction processing: SCHEDULER FM header	563
System message process (SYSMSG) and related FM headers	564

Chapter 32. Using MFS with ISC . . . 567

Activating MFS input formatting	567
Activating MFS output formatting for ISC	567
MFS Distributed Presentation Management (DPM) messages	568
MFS page delete function	568
MFS online error detection	569
The ATTACH and SCHEDULER FM headers under MFS	570
Data descriptor FM headers	571
Input data descriptor FM header	572
Output data descriptor FM header	572
Controlling demand-paged messages: QMODEL FM headers	572
Request (input) QMODEL FM headers	573
QGETN FM header	573
QGET FM header	574
QPURGE FM header	574
Reply (output) QMODEL FM headers	575
QXFR FM header	575
QSTATUS FM header	575
The RAP FM header under MFS	575

Chapter 33. FM header format

reference 577

ATTACH FM header format	577
ATTIU	579
ATTDSP	579
ATTDDBA	580
ATTDPN	581
ATTPRN	583
ATTRDPN and ATTRPRN	583
ATTDQN and ATTDQ	585
ATTACC	585
Data descriptor FM header formats	585
Error recovery procedure (ERP) FM header	587
QMODEL FM headers	587
QGET FM header format	587
QGETN FM header format	589
QPURGE FM header format	590
QSTATUS FM header format	590
QXFR FM header format	591
Reset attached process (RAP) FM header format	594
SCHEDULER FM header format	594
SYMSG process headers	595

Chapter 34. Examples using ISC edit

ATTACH parameters 597

ATTACH and SCHEDULER parameters with ISC edit	597
ATTACH parameters with the IMS SYMSG process	599
ATTACH and SCHEDULER parameters with IMS MFS	601

Chapter 35. How IMS and CICS use the ISC interface 607

Functions available to the ISC session	607

Overview of CICS synchronous and asynchronous processing for ISC	607
Functions available to an ISC VTAM session	609
General flow of CICS EXEC commands within a CICS application	612
CICS to IMS using SEND/RECEIVE EXEC commands	612
CICS to IMS using the SEND LAST EXEC command	614
IMS to CICS using the RECEIVE EXEC command	616
Coding asynchronous messages	617
CICS to IMS using the START/RETRIEVE EXEC commands	617
IMS to CICS using the RETRIEVE EXEC command	619
Commands that should not be used on an ISC session	620
Selecting appropriate CICS installation options for ISC	620
Coding CICS system definition options	620
Preparing CICS resource definition	621
Defining IMS-CICS LU 6.1 links	621
Defining compatible IMS and CICS nodes	622
System names	622
Number of sessions	622
Session names	623
Other session parameters	623
Defining multiple links to an IMS system	626
Defining CICS transactions for IMS-CICS ISC	628
Defining CICS backout in-doubt processing	628
Defining CICS transactions for asynchronous communication to IMS	629
Initiating and allocating a session from CICS	629
Other ways of initiating a session	630
Terminating a session from CICS	630
Designing CICS applications for ISC	630
Application-related concepts	631
Subsystem design: direct-control versus queued	631
Synchronous and asynchronous processing	631
Principal and alternate facility	634
CICS versus IMS conversation mode	634
Sending IMS commands from CICS	635
Sync points	636
Coding function management headers for CICS	638
ATTACH function management header	638
SCHEDULER function management header	641
Queue model function management headers	643
Data descriptor function management header	644
System message process (SYMSG) function management header	644
Error recovery procedure function management header	644
Recovery and restart concepts	645
Logical unit of work	645
Recovering outstanding message traffic after a failure	646
Handling transaction abends	649
Coding CICS applications for restart	650

Chapter 36. ISC data flow control examples 653

Non-MFS bracket and half-duplex protocol examples	653
MFS bracket and half-duplex protocol examples	655
MFS output examples	655
MFS input examples	658
SBI/BIS examples	660
Signal protocol example	663

Chapter 37. ISC error recovery procedure examples 665

Sender-detected error examples	665
Receiver-detected error examples	666

Chapter 38. Sample program for IMS-CICS ISC 669

Installation procedure	669
IMS sample program (DFSISC00)	669
Job control statements for the sample program	672
IMS system definition statements	672
MFS formats	673
Program specification block (PSB) generation for the sample program	674
Application control block (ACB) generation	674

Part 8. Multiple Systems Coupling (MSC) 675

Chapter 39. Overview of Multiple Systems Coupling 677

Multiple Systems Coupling concepts and terminology	677
MSC physical links	677
MSC logical links	679
MSC logical link paths	679
The MSC network and routing	681
Remote and local systems	681
Flow of data within multiple systems	681
Message routing	682
Routing path	683
Logical destinations	683
Input, destination, and intermediate systems	686
System identifiers (SYSIDs)	687
Routing messages with the destination name and SYSIDs	690
Remote LTERMs	691
Multiple Systems Coupling (MSC) directed routing	693
Remote destination verification	694

Chapter 40. Administering Multiple Systems Coupling 695

Design considerations for multiple systems	695
Minimizing resource consumption	695
Controlling the bandwidth of MSC links	696
Balancing resource demand	696
Planning for conversational processing	697

Routing exit routines with conversations	698
Remote destination verification for conversations	698
Saving truncated data in the SPA	698
Conversation termination	699
Abnormal conversation termination	700
Defining Multiple Systems Coupling resources	700
Local system definitions	700
Defining partner systems	701
Defining the physical link	701
Defining the logical link	704
Defining a logical path	705
Setting link priorities for remote transactions	705
Serial transaction processing in an MSC network	707
Specifying exit routines	707
How network definition is affected by multiple systems	709
Verifying transaction definitions across systems	709
Using the multiple systems verification utility	709
Verifying the system definition status online	710
Security considerations for MSC	711
Operations for Multiple Systems Coupling	712
MSC link statistics	712
Benchmark link activity	713
Determine your optimum MSC link type	713
Reset statistics regularly at system checkpoint	713
Adjust link buffer sizes to the size of the messages	714
Adjust logical link capacity for MSC bandwidth mode	714
Determining optimum MSC link buffer sizes	714
Use high-value link statistics to help diagnose MSC link problems	717
Monitoring and tuning multiple systems	718
Coordinating performance information	718
Reports generated by the IMS Monitor for MSC	719
Extracting multiple-system transaction statistics	719
Controlling the log merge	720
Interpreting the Transaction Analysis report	720
MSC and IMSplexes with shared queues	721
Message routing across MSC and IMSplex environments	721
Migrating from an MSC network to an IMSplex network	723
Managing remote transactions for APPC and OTMA when MSC and IMSplexes coexist	728
Avoiding pseudoabend U0830	734
VTAM Generic Resources (VGR) and MSC	735
TM and MSC Message Routing and Control user exit routine overview	735
The IMSplex affinity routing option of the DFSMSCE0 exit routine	738

Part 9. ODBA and DRA connections 741

Chapter 41. Accessing IMS databases with CICS 743

Coding considerations for PSBs	743
Using sequential buffering	744

CICS connected to DL/I.	744
Configuring CICS CCTL connections to IMS	
DBCTL systems	745
CICS tasks	745

Chapter 42. Accessing IMS databases through the ODBA interface 747

Creating the ODBA DRA start-up table.	747
Loading and running the ODBA and DRA modules in the z/OS application region	748
Binding application programs	748
Establishing and defining security	748
RAS security	749
Defining APSB security	750

Part 10. Open Transaction Manager Access (OTMA) 751

Chapter 43. Introduction to OTMA . . . 753

What is OTMA?	753
Capabilities of OTMA	754
Benefits of using OTMA.	755
Advantages of the OTMA protocol	757
How IMS messages flow in an OTMA environment	758
Basic OTMA message flow	758
Sample commit-then-send transaction processing flows	760
Using transaction pipes with OTMA	761
Differences in transaction pipes	763
Message flow using transaction pipes	763

Chapter 44. Enabling and using OTMA 767

Enabling OTMA	767
Summary of the OTMA configuration parameters	767
Defining the XCF group name.	768
Defining the OTMA XCF member name	768
Defining when OTMA starts up	769
Defining the level of OTMA security checking	769
Specifying synchronized tpipes for WebSphere MQ.	770
Enabling DFSYPRX0 member name override for OTMA clients	770
Specifying asynchronous delivery of program-to-program switch output messages.	770
OTMA descriptors.	771
OTMA client descriptors.	771
OTMA destination descriptors.	772
DFSOTMA descriptor	775
OTMA-supported exit routines	776
Using DFSYPRX0 and DFSYDRU0 OTMA exit routines to determine destination.	776
Administering IMS for OTMA.	779
IMS conversations and OTMA.	779
MSC and OTMA transactions	779
Fast Path and OTMA transactions	780
IMS restart processing and OTMA	780
XRF processing and OTMA	780
IMS Queue Control Facility and OTMA	781

Using shared queues with OTMA	782
IMS termination and OTMA	785
OTMA restrictions.	785
Managing system resources and OTMA	786
OTMA resource monitoring	786
Displaying the current transaction workload	787
Impact of OTMA message TIBs on storage	788
Message flood detection.	789
Administering OTMA tmembers	790
Tpipe resource impact	791
Automatic removal of idle tpipes.	793
IMS message queue data set size and OTMA	793
Buffer pool usage for OTMA	794
Dependent region occupancy and OTMA	794
Timeout interval for acknowledgments to OTMA messages	794
OTMA support for transaction expiration	797
OTMA security.	799
RACF security levels for OTMA	799
Specifying OTMA security	800
Securing messages on the asynchronous hold queue	804
General OTMA security considerations.	805
Using DL/I calls in an OTMA environment	806
OTMA program-to-program switch processing	807
OTMA single-stream program switch	807
OTMA program switch without ISRT to I/O PCB	808
OTMA program switch with express PCB	808
OTMA program switch to multiple programs	809
OTMA program switch for protected transactions	811
Other OTMA program switch considerations	811

Chapter 45. The OTMA client. 813

What is an OTMA client?	813
OTMA naming conventions	814
Messages sent by OTMA clients	815
Sending type-1 commands from an OTMA client	816
OTMA commit processing	817
Summary of OTMA commit processing.	818
Sample OTMA commit processing flows	819
Sample OTMA message flows.	824
Protecting transactions with OTMA	829
Initiating protected transactions from an OTMA client	829
Processing protected transactions in IMS	830
Client/server resynchronization with OTMA.	830
Assumptions for OTMA resynchronization	831
Recoverable OTMA transactions	831
Unrecoverable OTMA transactions	832
Summary results of IMS transactions and commands	832
OTMA resynchronization protocol	834
Sample OTMA resynchronization message flow	837
Sample OTMA resynchronization messages	839
Managing commit-then-send output.	840
Purging commit-then-send asynchronous output	840
Rerouting commit-then-send asynchronous output.	841

	Timeout for acknowledgments of	
	commit-then-send output	842
	Sharing asynchronous commit-then-send output:	
	the OTMA super member function	843

	Chapter 46. OTMA support for callout requests	845
	Callout requests from IMS application programs	845
	Synchronous callout requests	846
	Asynchronous callout request	847
	Implementing the asynchronous callout function	848
	IMS TM Resource Adapter and asynchronous	
	callout requests.	849
	SOAP Gateway and asynchronous callout	
	requests	850
	IMS application programs and the asynchronous	
	callout function.	851

	Chapter 47. OTMA message prefix	855
	Message-control information section.	855
	Explanation of OTMA message-control	
	information fields	859
	State data section	867
	Server-Available and Client-Bid commands	867
	SRVresynch command	869
	REQresynch command	870
	REPresynch command	870
	TBresynch command	871
	Transaction and callout messages.	871
	Server state protocol command	876
	Resume output for Tpipe	878
	Resume output for the hold queue for tpipe	879
	Cancel resume output for tpipe hold queue	
	request	880
	No messages on tpipe hold queue	880
	Security data section	880
	Explanation of OTMA security data fields.	881
	User data section	882
	Explanation of OTMA user data fields	883
	Application data section.	883
	Sample OTMA messages	884

	Chapter 48. OTMA Callable Interface	887
	OTMA C/I initialization.	889
	OTMA C/I security	889
	OTMA C/I restrictions	890
	Timing out OTMA C/I sessions after	
	otma_send_receive API calls for CM1 transactions	890

	Chapter 49. OTMA architected transaction attributes	893

	Part 11. SLU P and Finance Communication	897

	Chapter 50. Overview of SLU P and Finance Communication.	899
	The IMS-SLU P network.	899

	System configuration	900
	SLU P and Finance workstations	900
	System controller application program	900
	Writing the controller application program with	
	MFS and XRF	901
	Considerations for controller application	
	programs for XRF systems	901
	Converting controller application programs from	
	Finance to SLU P	902
	VTAM facilities used	902
	VTAM commands and indicators used with SLU P	903
	Request-recovery command	904
	Change-direction indicator	905
	Establishing connection and specifying logon	
	modes.	905
	Establishing connection with the XRF complex	906
	Bracket and send/receive management.	907

	Chapter 51. IMS facilities used for SLU P and Finance	909
	Component definition	909
	LTERM naming	909
	Output component selection	909
	Input component determination	910
	Terminal-response mode.	910
	Defining a workstation for terminal-response mode	912
	Output messages sent while in a between-brackets	
	state	913
	Designing for output messages sent while in	
	between-brackets state	914
	IMS Message Format Service	914
	Designing MFS for the workstation environment	914
	MID/MOD chaining	915
	MFS output formatting for the SLU P system	915
	MFS message recovery	915
	MFS control functions (Finance)	915
	MFS control functions (SLU P)	916
	MFS paging and BID options	916
	Display screen protection for finance stations.	917
	Extended output component protection (SLU P)	918
	Input and output editing options (SLU P)	919
	Use of responses or brackets to acknowledge	
	recoverable input	920
	Message recovery	921
	Message resynchronization	922
	Finance and SLU P in an XRF complex	923
	Fast Path messages with Finance and SLU P	923
	Fast Path output messages (Finance)	923
	Fast Path output messages (SLU P)	924
	Fast Path message resynchronization	925

	Chapter 52. Network operation for SLU P and Finance	927
	Starting an IMS network	927
	Making IMS ready	927
	Session initiation (starting workstations)	927
	Session-initiation transmission sequence	928
	Controller application program involvement in	
	message resynchronization	929
	Design considerations	929

Sequence number management	929
Set-and-Test-Sequence-Numbers (STSN)	930
Suspending output from IMS	934
Session termination	934
Orderly termination	935
Immediate termination	936
Shutting down an IMS network (SLU P)	936
SLU P messages	937
Send/receive and bracket protocol	937

Chapter 53. SLU P message protocols 939

General format of input function management headers (Finance)	939
Input message descriptor byte (Finance)	940
General format of input function management headers (SLU P)	940
Input message descriptor bytes (SLU P)	941
Input component identification (SLU P)	942
Input bracketing protocol	942
Activating MFS input formatting for Finance workstations.	942
Output messages	943
MFS Distributed Presentation Management output (SLU P)	945
General format of output function management headers (Finance)	946
Output message descriptor byte (Finance)	946
Output component ID byte (Finance)	946
MFS data bytes (Finance)	947
General format of output function management headers (SLU P)	947
Output message descriptor bytes (SLU P)	947
MFS data bytes (SLU P)	948

Output bracketing protocol	949
Activating MFS output formatting for SLU P	949
Response requests (Finance)	950
Response requests (SLU P)	950
Input response requirements	951
Output response requirements.	951
IMS transaction types	952
Recoverable-inquiry transactions	952
Irrecoverable-inquiry transactions	953
Verifying IMS receipt of irrecoverable messages	953
IMS message switches	954
IMS commands.	954
VTAM commands and indicators.	954
MFS control requests	954
Error handling	955
IMS-detected errors	955
Controller or station-detected errors.	956
VTAM logical unit status (LUSTATUS) command	957
VTAM ready-to-receive (RTR) command	958
VTAM CANCEL command.	958
VTAM request-recovery command	959

Part 12. Appendixes 961

Notices 963

Programming interface information	965
Privacy policy considerations	966

Bibliography. 967

Index 969

About this information

These topics describe how to administer IMS[™] communications and connections: CPI Communications and APPC/IMS, facilities for attaching to external subsystems, IMS Extended Terminal Option (ETO), IMS Connect, IMS Universal driver connections, Intersystem Communication (ISC), Multiple Systems Coupling (MSC), IMS Open Database Access (ODBA) and database resource adapter (DRA) interfaces, IMS Open Transaction Manager Access (OTMA), SLU P and Finance communication systems, TCP/IP communications, and VTAM[®] networking.

This information is available in IBM[®] Knowledge Center at www.ibm.com/support/knowledgecenter.

Prerequisite knowledge

You should first read *IMS Version 11 System Administration*. Its introductory chapters, which cover planning activities, the IMS environments, and administration concepts, is particularly useful as a background for this book.

If you use APPC/IMS, you must be familiar with APPC/MVS in order to correctly define APPC/MVS configurations. For more information about APPC/MVS, see *CPI Communications Reference*.

If you implement ISC sessions between IMS and CICS[®], you should understand the information in *CICS Transaction Server for z/OS CICS Intercommunication Guide*.

If you are connecting to IMS by using TCP/IP connections, you should understand z/OS[®] TCP/IP, as documented in *z/OS Communications Server: IP Configuration Guide*.

You can learn more about z/OS by visiting the “z/OS basic skills education” topics in IBM Knowledge Center.

You can gain an understanding of basic IMS concepts by reading *An Introduction to IMS*, an IBM Press publication. An excerpt from this publication is available in the IBM Knowledge Center.

IBM offers a wide variety of classroom and self-study courses to help you learn IMS. For a complete list of courses available, go to the IMS home page at ibm.com/ims and link to the Training and Certification page.

IMS function names used in this information

In this information, the term HALDB Online Reorganization refers to the integrated HALDB Online Reorganization function that is part of IMS Version 11, unless otherwise indicated.

IMS provides an integrated IMS Connect function, which offers a functional replacement for the IMS Connect tool (program number 5655-K52). In this information, the term IMS Connect refers to the integrated IMS Connect function that is part of IMS Version 11, unless otherwise indicated.

How new and changed information is identified

New and changed information in most IMS library PDF publications is denoted by a character (revision marker) in the left margin. The first edition (-00) of *Release Planning*, as well as the *Program Directory* and *Licensed Program Specifications*, do not include revision markers.

Revision markers follow these general conventions:

- Only technical changes are marked; style and grammatical changes are not marked.
- If part of an element, such as a paragraph, syntax diagram, list item, task step, or figure is changed, the entire element is marked with revision markers, even though only part of the element might have changed.
- If a topic is changed by more than 50%, the entire topic is marked with revision markers (so it might seem to be a new topic, even though it is not).

Revision markers do not necessarily indicate all the changes made to the information because deleted text and graphics cannot be marked with revision markers.

How to read syntax diagrams

The following rules apply to the syntax diagrams that are used in this information:

- Read the syntax diagrams from left to right, from top to bottom, following the path of the line. The following conventions are used:
 - The >>--- symbol indicates the beginning of a syntax diagram.
 - The ---> symbol indicates that the syntax diagram is continued on the next line.
 - The >--- symbol indicates that a syntax diagram is continued from the previous line.
 - The --->< symbol indicates the end of a syntax diagram.
- Required items appear on the horizontal line (the main path).

▶▶—required_item—————▶▶


- Optional items appear below the main path.

▶▶—required_item—┐optional_item┘—————▶▶


If an optional item appears above the main path, that item has no effect on the execution of the syntax element and is used only for readability.

▶▶—required_item—┐optional_item┘—————▶▶


- If you can choose from two or more items, they appear vertically, in a stack. If you *must* choose one of the items, one item of the stack appears on the main path.




If choosing one of the items is optional, the entire stack appears below the main path.




If one of the items is the default, it appears above the main path, and the remaining choices are shown below.



- An arrow returning to the left, above the main line, indicates an item that can be repeated.



If the repeat arrow contains a comma, you must separate repeated items with a comma.



A repeat arrow above a stack indicates that you can repeat the items in the stack.

- Sometimes a diagram must be split into fragments. The syntax fragment is shown separately from the main syntax diagram, but the contents of the fragment should be read as if they are on the main path of the diagram.



fragment-name:



- In IMS, a b symbol indicates one blank position.
- Keywords, and their minimum abbreviations if applicable, appear in uppercase. They must be spelled exactly as shown. Variables appear in all lowercase italic letters (for example, *column-name*). They represent user-supplied names or values.
- Separate keywords and parameters by at least one space if no intervening punctuation is shown in the diagram.

- Enter punctuation marks, parentheses, arithmetic operators, and other symbols, exactly as shown in the diagram.
- Footnotes are shown by a number in parentheses, for example (1).

Accessibility features for IMS Version 11

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in z/OS products, including IMS Version 11. These features support:

- Keyboard-only operation.
- Interfaces that are commonly used by screen readers and screen magnifiers.
- Customization of display attributes such as color, contrast, and font size.

Keyboard navigation

You can access IMS Version 11 ISPF panel functions by using a keyboard or keyboard shortcut keys.

For information about navigating the IMS Version 11 ISPF panels using TSO/E or ISPF, refer to the *z/OS TSO/E Primer*, the *z/OS TSO/E User's Guide*, and the *z/OS ISPF User's Guide Volume 1*. These guides describe how to navigate each interface, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

Related accessibility information

Online documentation for IMS Version 11 is available in IBM Knowledge Center.

IBM and accessibility

See the *IBM Human Ability and Accessibility Center* at www.ibm.com/able for more information about the commitment that IBM has to accessibility.

How to send your comments

Your feedback is important in helping us provide the most accurate and highest quality information. If you have any comments about this or any other IMS information, you can take one of the following actions:

- Click the **Feedback** link at the bottom of any IBM Knowledge Center topic.
- Send an email to imspubs@us.ibm.com. Be sure to include the book title and the publication number.

Changes to the IMS library

Changes to the IMS library for IMS Version 11 include minor title changes to all the publications, merging four sets of publications, a change to how *IMS Messages and Codes* are delivered, and a reorganization of *IMS Version 11 Exit Routines*. The library is no longer provided in BookManager format. In addition, as of the May, 2012, a master index is no longer provided.

The titles of all PDF publications in the IMS Version 11 library have been changed so that the publication titles match the contents view (navigation tree) titles in IBM Knowledge Center. For example, the *IMS Diagnosis Guide* has been renamed to *IMS Diagnosis*. In addition, the titles for the three *IMS Command Reference* PDF volumes have been renamed so that their titles better reflect the contents of each publication, as follows:

- *IMS Version 11 Commands, Volume 1: IMS Commands A-M*
- *IMS Version 11 Commands, Volume 2: IMS Commands N-V*
- *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*

Four sets of publications have been streamlined:

- To eliminate redundancy and group relevant application programming planning information with associated application programming tasks, the *IMS Application Programming Planning Guide* has been merged with the *IMS Application Programming Guide*, in a new title: *IMS Version 11 Application Programming*.
- To group diagnosis reference information with associated diagnosis tasks, the *IMS Diagnosis Reference* has been merged with the *IMS Diagnosis Guide*, in a new title: *IMS Version 11 Diagnosis*.
- To group all system administration tasks, regardless of whether the tasks relate to using a single IMS or an IMSplex, the *IMS IMSplex Administration Guide* has been merged with the *IMS System Administration Guide*, in a new title: *IMS Version 11 System Administration*.
- To group system definition reference information with associated system definition tasks, the *IMS System Definition Reference* has been merged with the *IMS System Definition Guide*, in a new title: *IMS Version 11 System Definition*.

IMS Version 11 Exit Routines is now organized by the call interface and available services for each group of exit routines, as follows:

- IMS system exits, database manager exits, and transaction manager exit routines are grouped under “IMS control region exit routines.” These exit routines use IMS interfaces and callable services, as described in the topic “Overview of exit routines for IMS.”
- Exit routines that run in an address space with BPE support, including those for CQS, BPE-based DBRC, and most CSL components, are grouped under “Base Primitive Environment-based exit routines.” BPE-based exit routines use the interfaces and services provided by the Base Primitive Environment, as described in the topic “BPE user-supplied exit routine interfaces and services.”
- The CQS exit routines are now divided into two groups: those that run in the CQS address space and have access to BPE services are grouped under “BPE-based CQS user-supplied exit routines,” and those that execute in a CQS client address space are grouped under “CQS client exit routines.”

- Standard DBRC exit routines are grouped under “IMS system exit routines,” while the variants that run in BPE are group under “BPE-based DBRC exit routines.”
- CSL SCI exit routines are now divided into two categories: those that run in the CSL SCI address space are grouped under “BPE-based SCL SCI user exit routines,” and those that run in the address space of a registered IMSplex member are grouped under “CSL SCI IMSplex member exit routines.”
- Exit routines that have specific guidelines, including those for IMS Connect, CQS clients, and CSL SCI IMSplex members, are described in the sections for those routines.

Part 1. Configuring external Java environment connections

IMS provides two generations of Java™ drivers and resource adapters: the IMS Universal drivers, which is the most recent generation, and the drivers and resource adapters provided by the classic Java APIs for IMS, which are the older generation.

Recommendation: Because the IMS Universal drivers are built on industry standards and open specifications, and provide more flexibility and improved support for connectivity, data access methods, and transaction processing options, use the IMS Universal drivers to support your Java applications that access IMS.

If you are migrating from the classic Java APIs for IMS to the IMS Universal drivers, you can have both generations of the IMS DB resource adapters installed on WebSphere® Application Server simultaneously. After you have switched your applications to use the IMS Universal Database resource adapters, you no longer need the imsDBJCA.rar of the classic Java APIs for IMS.

Related concepts:

 Transaction types and programming interfaces supported by the IMS Universal Database resource adapter (Application Programming)

 IMS solutions for Java development overview (Application Programming)

Chapter 1. IMS Universal drivers: configuring connections to IMS

The IMS Universal drivers can run in z/OS and distributed environments, including, WebSphere Application Server for z/OS, and WebSphere Application Server for distributed platforms.

The IMS Universal drivers include the following adapters and drivers:

- The IMS Universal Database resource adapters, which take advantage of Java Platform, Enterprise Edition (Java EE) services.
- The IMS Universal JDBC driver, which support SQL calls that directly access your IMS data.
- The IMS Universal DL/I driver, which provides calls that are similar to DL/I in a Java programming interface.

When running in a distributed environment on a server such as WebSphere Application Server for distributed platforms, or in a remote z/OS environments on a server such as WebSphere Application Server for z/OS, the IMS Universal drivers connect to IMS using a type-4 connection architecture, which supports TCP/IP communications and socket management.

Note: Support for type-2 connectivity from WebSphere Application Server for z/OS was added to IMS and the IMS Universal drivers by IMS APARs PK99686 and PM14766.

When running locally on the same logical partition (LPAR) as IMS, the IMS Universal drivers connect to IMS by using a type-2 connection architecture, which supports direct communication with IMS through the IMS Open Database Access (ODBA) and IMS database resource adapter (DRA) interfaces.

WebSphere Application Server supports all of the IMS Universal drivers in both distributed and z/OS environments.

Related concepts:

 Distributed and local connectivity with the IMS Universal drivers (Application Programming)

Related reference:

 Programming with the IMS Universal drivers (Application Programming)

 Comparison of IMS Universal drivers programming approaches for accessing IMS (Application Programming)

IMS Universal drivers: WebSphere Application Server type-4 connections

Java applications that run on WebSphere Application Server can access IMS databases by using the type-4 connectivity provided by the IMS Universal drivers. The type-4 connectivity of the IMS Universal drivers enables Java application programs to access IMS databases from a wide variety of distributed and mainframe environments, either in stand-alone mode or under an application server, with or without XA support for global transactions.

The following figure provides an overview of a configuration that uses the type-4 connectivity of an IMS Universal Database resource adapter to connect to IMS from WebSphere Application Server for distributed platforms.




Figure 1. A WebSphere Application Server EJB application using a type-4 IMS Universal Database resource adapter

The following procedures apply to both WebSphere Application Server for distributed platforms and WebSphere Application Server for z/OS.

Related concepts:

Transaction types and programming interfaces supported by the IMS Universal Database resource adapter (Application Programming)

Installing a type-4 IMS Universal Database resource adapter on WebSphere Application Server

You must install an IMS Universal Database resource adapter before it can be used to access IMS databases from WebSphere Application Server.

Prerequisites:

- Copy the RAR files for the IMS Universal drivers to storage that is accessible to WebSphere Application Server.

To install the IMS Universal Database resource adapter:

1. In the WebSphere Application Server administrative console, click **Resources > Resource Adapters > Resource Adapters**.
2. Click **Install RAR**. A dialog opens for installing the resource adapter.
3. Enter the path to the RAR files. The path can be to a local location or to a location on the server.
 - If your RAR file is located on your local workstation, select **Local path** and browse to find the file. For example: `C:\install_directory\imsudbLocal.rar`
 - If your RAR file is located on your server, select **Server path** and specify the fully qualified path to the file.
For example: `path_prefix/usr/lpp/ims/ims11/imsjava/rar/imsudbLocal.rar`
4. Click **Next**. The configuration panel opens.
5. Click **OK**. The IMS Universal Database resource adapter is listed.
6. In the messages box, click **Save**. The save page is displayed.
7. Click **Save** to update the master repository with your changes.

Related reference:

 Comparison of IMS Universal drivers programming approaches for accessing IMS (Application Programming)

Defining a connection factory for a type-4 IMS Universal Database resource adapter on WebSphere Application Server

After you install an IMS Universal Database resource adapter, you define the connection factory by using the WebSphere Application Server administrative console.

To define the connection factory for the IMS Universal Database resource adapter:

1. In the left frame of the WebSphere Application Server administrative console, click **Resources > Resource Adapters > Resource Adapters**.
2. Click the name of the IMS Universal Database resource adapter that you chose when you installed the adapter.
3. Under Additional Properties, click **J2C connection factories**.
4. Click **New**.
5. Type the following information:
 - Name**: the name for the connection factory, for example, PhonebookCF
 - JNDI Name**: the JNDI name for the connection factory that is unique within this server, for example, PhonebookCF
6. Click **Apply**. The data source is listed in the J2C Connection Factories.
7. Under Additional Properties, click **Custom Properties**.
8. Specify values defined by your installation for each of the following properties by clicking the property name and, in the configuration pane that opens,

specifying the value for the property in the **Value** field. After specifying the value for a property, click OK to return to the list of properties.

See Connecting using the IMS Universal Database resource adapter in a managed environment (Application Programming) for a list of connection properties.

9. In the messages box, click **Save**.
10. Restart WebSphere Application Server.

Related concepts:

 [IMS Connect definition and tailoring \(System Definition\)](#)

Installing an EAR file that uses a type-4 IMS Universal Database resource adapter on WebSphere Application Server

After installing a type-4 IMS Universal Database resource adapter and configuring the connection factory, you must install the client application program, or EAR file, on WebSphere Application Server.

To install the EAR file on the WebSphere Application Server:

1. In the left pane of the WebSphere Application Server administrative console, click **Applications > Install New Application**.
2. Enter the path to the EAR file or locate the EAR file by clicking the **Browse** button. Click **Next**.
3. In the “Selection installation options” panel, accept the defaults and click **Next**.
4. In the “Map modules to servers” panel, accept the defaults and click **Next**.
5. In the “Summary” panel, verify that the options are correct and click **Next**. A series of messages are issued that indicate that the application is being installed.
6. After the message that indicates that your application was installed successfully, click **Save**.

After your application is installed, it is ready to be started.

Related reference:

 [Programming with the IMS Universal drivers \(Application Programming\)](#)

IMS Universal drivers: WebSphere Application Server for z/OS type-2 connections

When WebSphere Application Server for z/OS and IMS are on the same logical partition (LPAR), Java applications running in WebSphere Application Server for z/OS can access IMS databases by using the type-2 connectivity provided by the IMS Universal Database resource adapters.

Note: Type-2 connectivity support for WebSphere Application Server for z/OS was added to IMS and the IMS Universal drivers through the service process by the IMS APARs PK99686 and PM14766.

The type-2 connectivity of the IMS Universal Database resource adapters provides local, non-TCP/IP access to IMS databases.

To deploy an application that uses the type-2 connectivity, you must install one of the IMS Universal Database resource adapters in WebSphere Application Server for z/OS and configure the IMS Open Database Access (ODBA) interface.

IMS Universal Database resource adapter type-2 connectivity supports both bean-managed bean methods and container-managed bean methods. Optionally, transaction applications can be managed by the z/OS Resource Recovery Services local option.

The following figure shows an EJB application that is accessing IMS data. The database requests are passed to a type-2 IMS Universal Database resource adapter, which converts the requests to DL/I calls. The IMS Universal Database resource adapter passes these calls to ODBA, which uses the IMS database resource adapter (DRA) interface to access the DL/I region in IMS.




Figure 2. A WebSphere Application Server for z/OS EJB application using a type-2 IMS Universal Database resource adapter

Configuring WebSphere Application Server for z/OS for a type-2 IMS Universal Database resource adapter

To use a type-2 IMS Universal Database resource adapter to access IMS DB from WebSphere Application Server for z/OS, you first must configure WebSphere Application Server for z/OS.

Prerequisites:

- If you are running WebSphere Application Server for z/OS V6.1 or V7.0 ensure that the fix pack that contains WebSphere Application Server APAR PK89274 has been applied to your system.
- If your transaction application is managed by the z/OS Resource Recovery Services local option, ensure that z/OS Resource Recovery Services is enabled.

To configure WebSphere Application Server for z/OS to access IMS databases:

1. Set the Java Virtual Machine custom property **com.ibm.ws.runtime.component.ResourceMgr.postBindNotify** to true. To set this custom property:
 - a. In the administrative console, click **Servers > Application servers > *server_name***.
 - b. In the Server Infrastructure section select **Java and process management > Process Definition**.
 - c. Select **Servant**.
 - d. Under Additional Properties, click **Java virtual machine**.
 - e. Under Additional Properties, click **Custom Properties**.
 - f. Set to true the custom property named **com.ibm.ws.runtime.component.ResourceMgr.postBindNotify**. If the custom property is not present in the list of already defined custom properties, create it and set it to true.
 - g. Save the changes.
 - h. Restart the server.
 - i. To query the custom properties of the corresponding Connection Factory, use the `getProperty` method exposed by the Connection Factory MBean.
2. If not already done, create an ODBA startup table. The ODBA startup table module name can be from 5 to 8 bytes long and must conform to the following naming convention:
 - Bytes 1-3 must be “DFS”
 - Bytes 4-7 are the 1- to 4-byte ID
 - The final byte must be the character number “0”

For example, both DFS10 and DFSIMSA0 are valid module names for an ODBA startup table.

Recommendation: Use the IMS ID as the 1- to 4-byte ID.

3. If not already done, link the ODBA startup table into a load library.
4. Update the JCL for WebSphere Application Server for z/OS by adding to the STEPLIB the following data sets:
 - The load library that contains the ODBA startup table and the ODBA runtime code.
 - The SDFSJLIB data set. This data set contains the DFSCLIB member.
5. Note the ODBA name, which is defined by the MBR parameter. You will need to know bytes 4-7, which are usually the IMS system ID, when you install the data source.

Related concepts:

 Database resource adapter (DRA) (System Programming APIs)

Related tasks:

Chapter 42, “Accessing IMS databases through the ODBA interface,” on page 747

Installing a type-2 IMS Universal Database resource adapter on WebSphere Application Server for z/OS

After you configure WebSphere Application Server for z/OS to have access to IMS databases, you must install the type-2 IMS Universal Database resource adapter on WebSphere Application Server for z/OS.

Prerequisite:

- Perform the steps described in “Configuring WebSphere Application Server for z/OS for a type-2 IMS Universal Database resource adapter” on page 7.
- Copy the RAR files for the IMS Universal drivers to storage that is accessible to WebSphere Application Server for z/OS.

To install the type-2 IMS Universal Database resource adapter:

1. In the WebSphere Application Server for z/OS administrative console, click **Resources > Resource Adapters > Resource Adapters**. A list of resource adapters is displayed.
2. Click **Install RAR**. A panel is displayed for installing the resource adapter.
3. In the “Install RAR” panel:
 - Under **Scope**, select a node
 - Under **Path**, enter the path to the RAR file or locate the RAR file by clicking the **Browse** button
4. Click **Next**. A configuration panel opens.
5. In the “Configuration” panel under **General properties > Native library path**, enter the path to the directory that contains the libT2DLI.so file and click **OK**. For example, /usr/lpp/ims11/imsjava/lib.
 The libT2DLI.so file must have the proper read and execute permissions in Unix Systems Services. Also the SDFSJLIB must be included in the STEPLIB for the WebSphere servant region.
 After you click **OK**, the IMS Universal Database resource adapter that you installed is listed.
6. In the messages box, click **Save**. A page is displayed that asks you if you want to synchronize the changes with the nodes.
7. Click **OK** to update the master repository with your changes.

Optional: set the WebSphere Application Server for z/OS classpath for applications that use a type-2 IMS Universal Database resource adapter

Your application can include the IMS database metadata class (the DLIDatabaseView subclass that is generated by the IMS Enterprise Suite Explorer for Development) or the metadata class can be stored elsewhere.

If your application does not include the metadata class, you must set the WebSphere Application Server for z/OS classpath to point to the IMS database metadata class that is used by the application.

One way to set the classpath is to add these files to the classpath of your IMS Universal Database resource adapter.

To add the required files to the classpath of an IMS Universal Database resource adapter:

1. From the WebSphere Application Server for z/OS administrative console, click **Resources > Resource Adapters**. A list of resource adapters is displayed.
2. Click the name of your IMS Universal Database resource adapter. A configuration dialog is displayed.
3. In the class path text box, add the location of the metadata class.
4. Click **OK**.
5. In the messages box, click **Save**. The save page is displayed.

6. Click **Save** to update the master repository with your changes.

Related reference:

 Generating the runtime Java metadata class (Application Programming)

Defining a connection factory for a type-2 IMS Universal Database resource adapter on WebSphere Application Server for z/OS

The DataSource facility is a factory for connections to a physical data source, or database. A data source is registered with a naming service based on the Java Naming and Directory (JNDI) API. DataSource objects have properties that pertain to the actual data source that an application needs to access.

Requirement: You must use the DataSource facility, which replaces the DriverManager facility, because the DriverManager facility is not supported by the Java EE Connection Architecture Specification.

To install the data source for your application:

1. In the left frame of the WebSphere Application Server for z/OS administrative console, click **Resources > Resource Adapters > Resource Adapters**. A list of resource adapters is displayed.
2. Click the name of the IMS Universal Database resource adapter that you chose when you installed the adapter. A configuration dialog is displayed.
3. Under Additional Properties, click **J2C connection factories**.
4. Click **New**. A configuration dialog is displayed.
5. Under General Properties, type the following information:
Name: the name for the connection factory, for example, PhonebookCF.
JNDI Name: the JNDI name for the connection factory that is unique within this server, for example, PhonebookCF.
6. Click **OK**. The new connection factory is listed in the table of resources that you can administer.
7. Click the name of the data source that you just installed.
8. Under Additional Properties, click **Custom Properties**. The properties are listed in a table.
9. Specify values defined by your installation for each of the following properties by clicking the property name and, in the configuration pane that opens, specifying the value for the property in the **Value** field. After specifying the value for a property, click OK to return to the list of properties.

DatastoreName

Enter bytes 4-7 of the DRA startup table module name (usually the IMS system ID). For more information about the DRA startup table, see “Configuring WebSphere Application Server for z/OS for a type-2 IMS Universal Database resource adapter” on page 7.

DriverType

Set the driverType to either of the following values:

- 2 Specifies a local transaction model in which a unit of work is scoped to a particular connection. Multiple connections can have independent units of work associated with each.

Application programs can issue local commit and rollback calls through either the JDBC Connection interface or the CCI LocalTransaction interface.

DriverType=2 does not support the UserTransaction interface.

Container-managed bean methods require the following properties in the EJB Deployment Descriptor:

- In the **Bean** tab, specify the following properties under the **LocalTransaction** heading:
 - **Boundary:** BeanMethod
 - **Resolver:** ContainerAtBoundary
 - **Unresolved action:** Rollback
- In the **Assembly** tab, set the transaction scope to NotSupported.

2_CTX

Specifies a global scope transaction model in which a unit of work can span multiple bean methods.

Application programs can use the UserTransaction interface for explicit commit and rollback calls.

Application programs cannot issue local commit and rollback calls through either the JDBC Connection interface or the CCI LocalTransaction interface.

When 2_CTX is specified, use the default properties of the EJB Deployment Descriptor.

MetadataURL

Enter the location of the database metadata representing the target IMS database. The MetadataURL property is the fully qualified name of the Java metadata class generated by the IMS Enterprise Suite DLIModel utility plug-in. The URL must be prefixed with class:// (for example, class://com.foo.BMP255DatabaseView).

10. In the messages box, click **Save**. The save page is displayed.
11. Click **Save** to update the master repository with your changes.
12. Restart the server.

Installing an EAR file that uses a type-2 IMS Universal Database resource adapter on WebSphere Application Server for z/OS

This topic describes how to deploy an application on WebSphere Application Server for z/OS.

Prerequisite: Perform the steps described in “Defining a connection factory for a type-2 IMS Universal Database resource adapter on WebSphere Application Server for z/OS” on page 10.

To install your application:

1. In the left pane of the WebSphere Application Server for z/OS administrative console, click **Applications > New Application > New Enterprise Application**.
2. Enter the path to the EAR file or locate the EAR file by clicking the **Browse** button. Click **Next**.
3. In the “Preparing for the application installation” panel, accept the default and click **Next**.
4. In the “Selection installation options” panel, accept the defaults and click **Next**.

5. In the “Map modules to servers” panel, accept the defaults and click **Next**.
6. In the “Summary” panel, verify that the options are correct and click **Next**. A series of messages are issued that indicate that the application is being installed.
7. After the message that indicates that your application was installed successfully, click **Save**.

The IMS Universal drivers: CICS connections

Java applications that run on IBM CICS Transaction Server for z/OS can access IMS databases by using the type-2 connectivity provided by the IMS Universal drivers. Other than the Java layer, access to IMS from a Java application is the same as for a non-Java application.

Note: Type-2 connectivity support will be delivered through the IMS service process.

CICS supports the following IMS Universal drivers:

- IMS Universal JDBC driver
- IMS Universal DL/I driver

The following figure shows a JCICS application that is accessing an IMS database by using the IMS database resource adapter (DRA) interface and the type-2 connectivity of an IMS Universal driver.




Figure 3. CICS application using a type-2 IMS Universal driver

Configuring CICS for the type-2 IMS Universal drivers

To run Java applications in a CICS environment that access IMS databases through a type-2 IMS Universal driver, you must install the type-2 IMS Universal driver in the IBM CICS Transaction Server for z/OS subsystem.

Prerequisite: Load the install files for the type-2 IMS Universal driver in a path that the CICS subsystem can access.

To configure CICS for a type-2 IMS Universal driver:

1. Build the DFST2CIC DLL. To build a DFST2CIC DLL:
 - a. Modify the Makefile, which is in the *pathprefix*/usr/lpp/ims/ims11/imsjava/cics directory, by changing both occurrences of *your.pdse.loadlib* to the data set that will store the DFST2CIC module.
 - b. Set the CICS library to the `_CXX_LXSYSLIB` environment variable by issuing the following commands from the UNIX System Services prompt:

```
export _CXX_LXSYSLIB=CEE.SCEEBND2:CBK.SCCNOBJ:$ _CXX_LXSYSLIB
export _CXX_LXSYSLIB=CICSVS.TS320.CICS.SDFHLOAD:$ _CXX_LXSYSLIB
export _CXX_LXSYSLIB=IMS.SDFSRESL:$ _CXX_LXSYSLIB
export _CXX_LXSYSLIB=SYS1.CSSLIB:$ _CXX_LXSYSLIB
```
 - c. Run the Makefile by issuing the following command from the UNIX System Services prompt: `make`. In the data set that you specified in the Makefile, member DFST2CIC is created.
 - d. Add the following data sets to the STEPLIB statement in your CICS startup job:
 - The data set that you specified in the Makefile that now contains DFST2CIC.
 - The data set that contains DFSCLIBU.
 - SDFHAUTH and SDFJAUTH, which are required for Java support.
 - If it is not already specified in a library on the MVS™ LNKSTnn concatenation, SCEERUN2. Specify SCEERUN2 in the same location on the STEPLIB statement as SCEERUN. SCEERUN2 is the Language Environment® run-time library and is required to support the IBM JVM.
2. Modify the CICS environment UNIX System Services/HFS file, DFHJVMPR, that contains the JVM profile.
 - a. Add a CLASSPATH_SUFFIX variable, which adds libraries to the end of the middleware classpath.
 - b. To the CLASSPATH_SUFFIX variable, add the path to the IMS Universal drivers that your installation requires as follows:

```
CLASSPATH_SUFFIX=pathprefix/usr/lpp/ims/ims11/imsjava/imsudb.jar
```
 - c. Update the LIBPATH variable so that it contains the path to the file libT2DLI.so as follows:

```
LIBPATH_SUFFIX=pathprefix/usr/lpp/ims/ims11/imsjava/lib
```

Related reading: For detailed information about CICS system definition, see the *CICS Transaction Server for z/OS CICS System Definition Guide*.

Running applications on CICS that use the type-2 IMS Universal drivers

To run a CICS application program that accesses IMS DB through a IMS Universal driver, you must perform several steps.

To run a Java application in CICS that accesses IMS DB through a type-2 IMS Universal driver, complete the following steps.

1. Start IMS DB and CICS.
2. Turn off the uppercase translation feature of CICS by entering: CEOT NOUCTRAN
3. Define a program that can run the Java application (JVM class).

- | 4. Define a transaction that can run the program.
- | 5. Install the program that you defined to run the Java application (JVM class).
- | 6. Install the transaction that you defined to run the Java application (JVM class).

Chapter 2. Classic Java APIs for IMS: configuring connections to IMS

The classic Java APIs for IMS provide a Java API for Java application development, as well as a subset of SQL calls for accessing IMS databases.

Recommendation: The following information applies to the classic Java APIs for IMS, which are no longer recommended. Instead, use the IMS Universal drivers to develop Java applications that access IMS. The IMS Universal drivers, which are built on industry standards and open specifications, provide more flexibility and improved support for connectivity, data access methods, and transaction processing options.

Related concepts:

 Overview of the IMS Java dependent regions (Application Programming)

Related tasks:

Chapter 1, “IMS Universal drivers: configuring connections to IMS,” on page 3

Related reference:

 Programming with the IMS Universal drivers (Application Programming)

 Programming with the classic Java APIs for IMS (Application Programming)

Classic Java APIs for IMS: configuring CICS

Java applications that run on CICS Transaction Server for z/OS can access IMS databases by using the classic Java APIs for IMS. Other than the Java layer, access to IMS from a Java application is the same as for a non-Java application.

The following figure shows a JCICS application that is accessing an IMS database using DRA and classic Java APIs for IMS.




Figure 4. CICS application using the classic Java APIs for IMS

Configuring CICS for the classic Java APIs for IMS

To run Java applications that access IMS databases in a CICS environment, you must have CICS Transaction Server for z/OS installed.

Prerequisite: Install the Java class libraries for IMS DB.

To configure CICS for the Java class libraries for IMS DB:

1. Build the CICSPSB DLL. To build a CICSPSB DLL:
 - a. Modify the Makefile by changing both occurrences of your.pdse.loadlib to the data set that will store the CICSPSB module. The Makefile is in the *pathprefix/ims/ims11/imsjava/classic/lib/cics* directory.
 - a. Set the CICS library to the `_CXX_LSYSLIB` environment variable by issuing the following command from the UNIX System Services prompt: `export _CXX_LSYSLIB=CICS.SDFHLOAD:IMS.SDFSRESL:$_CXX_LSYSLIB`
 - a. Run the Makefile by issuing the following command from the UNIX System Services prompt: `make`. In the data set that you specified in the Makefile, two data set members named CICSPSB and DFSCLIB are created.
 - b. Add the following data sets to the STEPLIB statement in your CICS startup job:

The data set that you specified in the Makefile that now contains CICSPSB and DFSCLIB.

SDFJAUTH, which is required for Java support.

If it is not already specified in a library on the MVS LNKSTnn concatenation, SCEERUN2. Specify SCEERUN2 in the same location on the STEPLIB statement as SCEERUN. SCEERUN2 is the Language Environment run-time library and is required to support the IBM JVM.

2. Modify the CICS environment UNIX System Services/HFS file that contains the JVM profile. The name of the UNIX System Services/HFS file is defined by your installation:
 - a. Add a TMPPREFIX variable, which adds libraries to the beginning of the middleware path.
 - b. To the TMPPREFIX variable, add the path to the Java and XML class libraries that your installation requires as follows:


```
TMPPREFIX=pathprefix/ims/ims11/imsjava/classic/imsjavaBase.jar:
pathprefix/ims/ims11/imsjava/classic/imsJDBC.jar:
pathprefix/ims/ims11/imsjava/classic/imsXQuery.jar
```
 - c. Update the LIBPATH variable so that it contains the path to the file libJavTDLI.so as follows: LIBPATH=pathprefix/ims/ims11/imsjava/classic/lib

For information about CICS system definition, see the *CICS Transaction Server for z/OS CICS System Definition Guide*.

Related reference:

 Software requirements (Release Planning)

Running the IMS IVP application for the classic Java APIs for IMS and CICS

After you configure CICS to run Java applications that access IMS databases, verify that the classic Java APIs for IMS are installed correctly and that CICS is configured correctly by running CICSIVP, the IMS installation verification program (IVP) for the CICS Java environment.

Prerequisites:

- “Configuring CICS for the classic Java APIs for IMS” on page 16
- Ensure that the base IMS IVPs have been run. These IVPs prepare the DBD for the IVP database (named IVPDB2) and load the IVP database. They also prepare the Java application PSB (named DFSIVP37), build ACBs, and prepare other IMS control blocks required by the CICSIVP application. For information about how to run the base IMS IVP procedures, see *IMS Version 11 Installation*.

The necessary files for the IVP applications are included in the IMS Java On Demand features FMID.

To run the CICSIVP application on CICS:

1. Create an HFS file named dfjjvmpr.props that contains the following class path: `ibm.jvm.shareable.application.class.path=/pathprefix/ims/ims11/imsjava/classic/samples.jar` If you need to debug your application, you can also add the JVM debug options in this file.
2. In the UNIX System Services/HFS file that contains the JVM profile, add:


```
JVMPROPS=path/dfjjvmpr.props
STDOUT=path
STDERR=path
```
3. Start IMS DB and CICS.
4. Turn off the uppercase translation feature of CICS by entering CEOT NOUCTRAN. By default, everything you type on the CICS terminal is converted to uppercase characters. However, the `samples.jar` file and `path` contain lowercase letters that must remain in lowercase letters.

5. Define a program that can run the CICSIVP application (JVM class):
 - a. From the CICS terminal, enter: CEDA DEFINE PROGRAM
 - b. In the list of program attributes, type the following information:


```

PROGram      ==> cicsivp
Group        ==> ivp
COncurrency  ==> Threadsafe
JVM          ==> Yes
JVMClass     ==> samples.ivp.cics.CICSIVP
JVMProfile   ==> The name of the UNIX System Services/HFS file
               that contains the JVM profile
          
```
 - c. Press F3 to return to the main CICS terminal.
6. Define a transaction that can run the program:
 - a. From the CICS terminal, enter: CEDA DEFINE TRANSACTION
 - b. In the list of transaction attributes, type the following information:


```

TRANSAction  ==> civp
Group        ==> ivp
PROGram      ==> cicsivp
          
```
 - c. Press F3 to return to the main CICS terminal.
7. Install the program that runs the CICSIVP application:
 - a. From the CICS terminal, enter: CEDA INSTALL
 - b. In the list of program attributes, type the following information:


```

PROGram      => cicsivp
Group        => ivp
          
```
 - c. Press F3 to return to the main CICS terminal.
8. Install the transaction that runs the CICSIVP application:
 - a. From the CICS terminal, enter: CEDA INSTALL
 - b. In the list of transaction attributes, type the following information:


```

TRANSAction  => civp
Group        => ivp
          
```
 - c. Press F3 to return to the main CICS terminal.
9. Run the transaction by entering: civp.
 If the IVP was successful, it displays IVP PASSED.
 If the IVP was not successful, it displays IVP FAILED or IVP INCOMPLETE.
 The STDOUT file contains the results of the individual tests that are completed by the IVP. The STDOUT file is a UNIX System Services/HFS file. The name of the STDOUT file is specified on the STDOUT= parameter in the UNIX System Services/HFS file that contains the JVM profile for the CICS environment.

Running your applications on CICS with the classic Java APIs for IMS

Configuring IMS and CICS to support a CICS application program that accesses IMS DB through the classic Java APIs for IMS requires you to perform several steps.

Prerequisite: Perform the steps described in “Running the IMS IVP application for the classic Java APIs for IMS and CICS” on page 17.

To run a Java application that accesses IMS DB from CICS through the classic Java APIs for IMS complete the following steps.

1. Modify the HFS dfjvmp.r.props file to set the `ibm.jvm.shareable.application.class.path=` parameter to the path of the application. The location of the dfjvmp.r.props file is specified by the JVMPROPS variable in the CICS JVM profile.
2. Start IMS DB and CICS.
3. Turn off the uppercase translation feature of CICS by entering: CE0T NOUCTRAN
4. Define a program that can run the Java application (JVM class).
5. Define a transaction that can run the program.
6. Install the program that you defined to run the Java application (JVM class).
7. Install the transaction that you defined to run the Java application (JVM class).

Related tasks:

 Programming Java applications for CICS with the classic Java APIs for IMS (Application Programming)

Classic Java APIs for IMS: configuring DB2 for z/OS

You can write Java stored procedures for DB2® for z/OS that access IMS databases through the classic Java APIs for IMS.

To deploy a Java stored procedure on DB2 for z/OS, you must configure the classic Java APIs for IMS, ODBA, and DRA.

The following figure shows a DB2 for z/OS stored procedure using the classic Java APIs for IMS, ODBA, and DRA to access IMS databases.




Figure 5. DB2 for z/OS stored procedure using the Java class libraries for IMS DB

Configuring DB2 for z/OS for the classic Java APIs for IMS

Configuring access for DB2 for z/OS stored procedures to IMS DB by using the classic Java APIs for IMS involves creating a data set pointed to by the JAVAENV DD statement, specifying the parameters of the ENVAR keyword, and setting the DB2 for z/OS environment variables in UNIX System Services.

Prerequisite: Installing the classic Java APIs for IMS.

To configure DB2 for z/OS for the classic Java APIs for IMS:

1. Create a data set with the following attributes. The JAVAENV DD statement points to this data set.
 - Organization: PS
 - Record format: VB
 - Record length: 1028
 - Block size: 6144
2. In the data set that is pointed to by the JAVAENV DD statement, add the ENVAR keyword with following parameters:

JAVA_HOME

The HFS directory of the JVM.

DB2_HOME

The HFS directory of the JDBC driver for DB2 for z/OS.

CLASSPATH

The HFS directories of the client application Java class files. Do not

specify the CLASSPATH parameter if you specify the client application Java class files in the stored procedure definition.

LIBPATH

The HFS directory of the file libJavTDLI.so.

TMSUFFIX

The HFS directories of the Java and XML class libraries that your installation requires:

```
TMSUFFIX=pathname/ims/ims11/imsjava/classic/imsjavaBase.jar
:pathname/ims/ims11/imsjava/classic/imsJDBC.jar
:pathname/ims/ims11/imsjava/classic/imsXQuery.jar
```

The following example shows a sample of the contents of the JAVAENV data set:

```
ENVAR("CLASSPATH=/usr/lpp/ims/ims11/imsjava/classic/samples.jar",
      "DB2_HOME=/usr/lpp/db2/db28",
      "JAVA_HOME=/usr/lpp/J1.4",
      "LIBPATH=/usr/lpp/ims/ims11/imsjava/classic/lib",
      "TMSUFFIX=pathname/ims/ims11/imsjava/classic/imsjavaBase.jar
:pathname/ims/ims11/imsjava/classic/imsJDBC.jar
:pathname/ims/ims11/imsjava/classic/imsXQuery.jar"
```

3. Set DB2 for z/OS environment variables in UNIX System Services by issuing the following commands:

```
export SQLJ_HOME=location of DB2 SQLJ driver (for example /usr/lpp/db2/db2810)
export JDBC_HOME=location of DB2 JDBC driver (for example /usr/lpp/db2/db2810)
export JAVA_HOME=location of SDK1.3 (for example /usr/lpp/java/J1.3)

export DB2SQLJPROPERTIES=/path/db2sqljjdbc.properties (file created later)

export CLASSPATH=$JDBC_HOME/classes/db2jdbcclasses.zip
export CLASSPATH=$CLASSPATH:$SQLJ_HOME/classes/db2sqljruntime.zip
export CLASSPATH=$CLASSPATH:$SQLJ_HOME/classes/db2sqljclasses.zip
export CLASSPATH=$CLASSPATH:pathname/ims/ims11/imsjava/classic/imsjavaBase.jar
export CLASSPATH=$CLASSPATH:pathname/ims/ims11/imsjava/classic/imsJDBC.jar
export CLASSPATH=$CLASSPATH:$JAVA_HOME/lib:.

export LIBPATH=$SQLJ_HOME/lib:$JDBC_HOME/lib
export LIBPATH=$JAVA_HOME/lib:$LIBPATH

export LD_LIBRARY_PATH=.:$SQLJ_HOME/bin:$JDBC_HOME/lib
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$JAVA_HOME/lib

export PATH=$SQLJ_HOME/bin:$PATH

export STEPLIB=yourDB2HLQ.DSNEXIT:yourDB2HLQ.SDSNLOAD:
export STEPLIB=yourDB2HLQ.SDSNLOAD2:yourDB2HLQ.SDSNLINK:$STEPLIB
```

Running the IMS IVP applications for the classic Java APIs for IMS and DB2 for z/OS

To verify that DB2 for z/OS is configured correctly and that the classic Java APIs for IMS are installed correctly, run the IMS installation verification program (IVP) applications for the DB2 for z/OS Java environment on DB2 for z/OS.

Prerequisites:

- “Configuring DB2 for z/OS for the classic Java APIs for IMS” on page 20
- Ensure that the base IMS IVPs have been run. These IVPs prepare the DBD for the IVP database, named IVPDB2, and load the IVP database. They also prepare the Java application PSB (named DFSIVP37), build ACBs, and prepare other IMS control blocks required by the IMS IVP applications for the DB2 for z/OS Java environment.

The IMS IVP applications for the DB2 for z/OS Java environment consist of two programs, delivered as part of the IMS Java On Demand features FMID:

- The Java application DB2IvpClient, which runs under UNIX System Services
- The stored procedure DB2IvpStoredProcedure, which runs in a WLM-managed address space

To run the IMS IVP applications for the DB2 for z/OS Java environment on DB2 for z/OS:

1. In the JAVAENV data set, modify the CLASSPATH= parameter to *pathprefix/ims/ims11/imsjava/classic/samples.jar*.
2. Edit the IMS-provided V81AWLM procedure, depending on your version of DB2 for z/OS, as follows (if IMS.SDFSRESL does not contain the DRA startup table, add that data set to the DFSRESLB DD statement):

```
//V81AWLM PROC RGN=OM,APPLENV=,
//          DB2SSN=,NUMTCB=
/** Define the V81AWLM procedure parameters here on in the service policy.
//IEFPROC EXEC PGM=DSNX9WLM,REGION=&RGN,TIME=NOLIMIT,
//          PARM='&DB2SSN,&NUMTCB,&APPLENV'
//STEPLIB DD DISP=SHR,DSN=CEE.SCEERUN
/** DB2 Library
//          DD DISP=SHR,DSN=yourDB2HLQ.SDSNLOAD
//          DD DISP=SHR,DSN=yourDB2HLQ.SDSNLOD2
//          DD DISP=SHR,DSN=yourDB2HLQ.RUNLIB.LOAD
/** DBRM library
//          DD DISP=SHR,DSN=yourHLQ.SDSNDBRM
//DFSRESLB DD DISP=SHR,DSN=IMS.SDFSRESL
//JAVAENV DD DISP=SHR,DSN=data set with ENVAR settings
//JSPDEBUG DD SYSOUT=*
//CEEDUMP DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
```

3. Create a new service policy for WLM. You can define the V81AWLM procedure parameters in the service policy or you can modify the procedure itself.
4. Define the procedure V81AWLM to RACF®.
5. Start DB2 for z/OS, the WLM-managed address space, and IMS DB.
6. Define the stored procedure to DB2 for z/OS by running the following job (*Your_WLM_Environment_Name* must match the APPLENV= parameter of the V81AWLM procedure):

```
//CREATIVP JOB 'name',
//          MSGCLASS=H,TIME=3,
//          USER=SYSADM,PASSWORD=XXXXXXXX,
//          MSGLEVEL=(1)
//CREATJSP EXEC PGM=IKJEFT01,DYNAMNBR=20
//STEPLIB DD DISP=SHR,DSN=DB2HLQ.DSNEXIT
//          DD DISP=SHR,DSN=DB2HLQ.SDSNLOAD
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
DSN SYSTEM(DB2_Subsystem_Name)
  RUN PROGRAM(DSNTIAD) PLAN(DSNTIA81) -
  LIB('DB2HLQ.RUNLIB.LOAD') -
  PARM('RC0')
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
CREATE PROCEDURE IVPStoredProc(VARCHAR (50) IN, VARCHAR(800) OUT)
  FENCED
  NO SQL
  LANGUAGE JAVA
  DYNAMIC RESULT SET 0
  EXTERNAL NAME 'samples.ivp.db2.DB2IvpStoredProcedure.execute'
```

```

PARAMETER STYLE JAVA
COLLID DSNJDBC
WLM ENVIRONMENT Your_WLM_Environment_Name;
GRANT EXECUTE ON PROCEDURE IVPStoredProc TO PUBLIC;

```

7. Create a DB2 for z/OS plan that runs the client program by running the following job:

```

//BNDIVPCL JOB , 'YOUR NAME',
//          MSGCLASS=H, TIME=3,
//          USER=SYSADM, PASSWORD=XXXXXXXX,
//          MSGLEVEL=(1)
//BINDCLNT EXEC PGM=IKJEFT01, DYNAMNBR=20
//STEPLIB DD DISP=SHR, DSN=DB2HLQ.DSNEXIT
//          DD DISP=SHR, DSN=DB2HLQ.SDSNLOAD
//DBRMLIB DD DISP=SHR, DSN=DB2HLQ.SDSNDBRM
//SYSTSPRT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSTSIN DD *
DSN SYSTEM(DB2ID)
  BIND PACKAGE (DSNJDBC) MEMBER(DSNJDBC1) ISOLATION(UR) -
  ACTION(REPLACE) VALIDATE(BIND)
  BIND PACKAGE (DSNJDBC) MEMBER(DSNJDBC2) ISOLATION(CS) -
  ACTION(REPLACE) VALIDATE(BIND)
  BIND PACKAGE (DSNJDBC) MEMBER(DSNJDBC3) ISOLATION(RS) -
  ACTION(REPLACE) VALIDATE(BIND)
  BIND PACKAGE (DSNJDBC) MEMBER(DSNJDBC4) ISOLATION(RR) -
  ACTION(REPLACE) VALIDATE(BIND)
  BIND PLAN(DB2IVPCL) KEEP DYNAMIC(YES) ACTION(REPLACE) -
    PKLIST(DSNJDBC.DSNJDBC1, -
            DSNJDBC.DSNJDBC2, -
            DSNJDBC.DSNJDBC3, -
            DSNJDBC.DSNJDBC4)
  RUN PROGRAM(DSNSTEP2) PLAN(DSNSTEP81) -
    LIB('DB2HLQ.RUNLIB.LOAD')
END
//SYSTSIN DD *
GRANT EXECUTE ON PLAN DB2IVPCL TO PUBLIC;
/*
//

```

8. In UNIX System Services, in the directory that you specified by the export DB2SQLJPROPERTIES command, create the file db2sqljjdbc.properties that contains the following:

```

DB2SQLJSSID=yourDB2ID
DB2SQLJPLANNAME=DB2IVPCL
DB2SQLJATTACHTYPE=RRSAF
DB2SQLJDBRMLIB=DB2HLQ.SDSNDBRM

```

9. Run the client application by issuing following command from UNIX System Services: `java samples.ivp.db2.DB2IvpClient IMSID`

The IVP displays the results of the tests that it performs.

If the IVP was successful, it displays IVP PASSED.

If the IVP was not successful, it displays IVP FAILED or IVP INCOMPLETE.

Fix any errors and rerun the IVP.

Related concepts:

 [IMS installation verification program \(IVP\) overview \(Installation\)](#)

Running your stored procedure with the classic Java APIs for IMS

Running a Java application that accesses IMS DB through the classic Java APIs for IMS on DB2 for z/OS requires modifications to IMS data sets and procedures, z/OS and RACF settings, and DB2 for z/OS.

Prerequisite: Perform the steps described in “Running the IMS IVP applications for the classic Java APIs for IMS and DB2 for z/OS” on page 21.

To run your Java application that accesses IMS DB through the classic Java APIs for IMS on DB2 for z/OS:

1. In the JAVAENV data set, modify the CLASSPATH= parameter to point to your application files. If your application files are in JAR files, include the JAR file names. If the application files are not in JAR files, do not include the file names.
2. Edit the IMS-provided V81AWLM procedure as follows (if IMS.SDFSRESL does not contain the DRA startup table, add that data set to the DFSRESLB DD statement):

```
//V81AWLM PROC RGN=0M,APPLENV=,
//          DB2SSN=,NUMTCB=
/** Define the V81AWLM procedure parameters here on in the service policy.
//IEFPROC EXEC PGM=DSNX9WLM,REGION=&RGN,TIME=NOLIMIT,
//          PARM='&DB2SSN,&NUMTCB,&APPLENV'
//STEPLIB DD DISP=SHR,DSN=CEE.SCEERUN
/* DB2 Library
//          DD DISP=SHR,DSN=yourDB2HLQ.SDSNLOAD
//          DD DISP=SHR,DSN=yourDB2HLQ.SDSNLOAD2
//          DD DISP=SHR,DSN=yourDB2HLQ.RUNLIB.LOAD
/* DBRM library
//          DD DISP=SHR,DSN=yourHLQ.SDSNDBRM
//DFSRESLB DD DISP=SHR,DSN=IMS.SDFSRESL
//JAVAENV DD DISP=SHR,DSN=data set with ENVAR settings
//JSPDEBUG DD SYSOUT=*
//CEEDUMP DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
```

3. Create a new service policy for WLM. You can define the V81AWLM procedure parameters in the service policy or you can modify the procedure itself.
4. Define the procedure V81AWLM to RACF.
5. Start DB2 for z/OS, the WLM-managed address space, and IMS DB.
6. Define the stored procedure to DB2 for z/OS by running the following job (*Your_WLM_Environment_Name* must match the APPLENV= parameter of the V81AWLM procedure):

```
//name JOB , 'name',
//      MSGCLASS=H,TIME=3,
//      USER=user,PASSWORD=XXXXXXX,
//      MSGLEVEL=(1)
//CREATJSP EXEC PGM=IKJEFT01,DYNAMNBR=20
//STEPLIB DD DISP=SHR,DSN=DB2HLQ.DSNEXIT
//          DD DISP=SHR,DSN=DB2HLQ.SDSNLOAD
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
DSN SYSTEM(DB2_Subsystem_Name)
RUN PROGRAM(DSNTIAD) PLAN(DSNTIA81) -
LIB('DB2HLQ.RUNLIB.LOAD') -
PARM('RC0')
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
```

```

CREATE PROCEDURE StoredProcName(... IN,... OUT)
FENCED
NO SQL
LANGUAGE JAVA
DYNAMIC RESULT SET 0
EXTERNAL NAME 'package.StoredProcedure.targetMethod'
PARAMETER STYLE JAVA
COLLID DSNJDBC
WLM ENVIRONMENT Your_WLM_Environment_Name;
GRANT EXECUTE ON PROCEDURE StoredProcName TO PUBLIC;

```

7. Create a DB2 for z/OS plan that runs the client program by running the following job:

```

//name      JOB  ,'name',
//          MSGCLASS=H,TIME=3,
//          USER=user,PASSWORD=XXXXXXX,
//          MSGLEVEL=(1)
//BINDCLNT EXEC PGM=IKJEFT01,DYNAMNBR=20
//STEPLIB DD  DISP=SHR,DSN=DB2HLQ.DSNEXIT
//          DD  DISP=SHR,DSN=DB2HLQ.SDSNLOAD
//DBRMLIB DD  DISP=SHR,DSN=DB2HLQ.SDSNDBRM
//SYSTSPRT DD  SYSOUT=*
//SYSPRINT DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//SYSTSIN DD  *
DSN SYSTEM(DB2ID)
  BIND PACKAGE (DSNJDBC) MEMBER(DSNJDBC1) ISOLATION(UR)  -
  ACTION(REPLACE) VALIDATE(BIND)
  BIND PACKAGE (DSNJDBC) MEMBER(DSNJDBC2) ISOLATION(CS)  -
  ACTION(REPLACE) VALIDATE(BIND)
  BIND PACKAGE (DSNJDBC) MEMBER(DSNJDBC3) ISOLATION(RS)  -
  ACTION(REPLACE) VALIDATE(BIND)
  BIND PACKAGE (DSNJDBC) MEMBER(DSNJDBC4) ISOLATION(RR)  -
  ACTION(REPLACE) VALIDATE(BIND)
  BIND PLAN(plan_name) KEEPYNAMIC(YES) ACTION(REPLACE) -
    PKLIST(DSNJDBC.DSNJDBC1, -
            DSNJDBC.DSNJDBC2, -
            DSNJDBC.DSNJDBC3, -
            DSNJDBC.DSNJDBC4)
  RUN PROGRAM(DSNSTEP2) PLAN(DSNSTEP81) -
    LIB('DB2HLQ.RUNLIB.LOAD')
END
//SYSIN DD *
GRANT EXECUTE ON PLAN plan_name TO PUBLIC;
/*
//

```

8. In UNIX System Services, in the directory that you specified by the export DB2SQLJPROPERTIES command, create the file db2sqljjdbc.properties that contains the following:

```

DB2SQLJSSID=yourDB2ID
DB2SQLJPLANNNAME=plan_name
DB2SQLJATTACHTYPE=RRSAF
DB2SQLJDBRMLIB=DB2HLQ.SDSNDBRM

```

9. Run the client application.

Related tasks:

 Programming Java applications in DB2 for z/OS stored procedures with the classic Java APIs for IMS (Application Programming)

Classic Java APIs for IMS: configuring WebSphere Application Server for z/OS

When WebSphere Application Server for z/OS and IMS are on the same LPAR (logical partition), you can write Java applications that run on WebSphere Application Server for z/OS and access IMS databases through the IMS DB resource adapter of the classic Java APIs for IMS.

Recommendation: The following information applies to the classic Java APIs for IMS, which are no longer recommended. Instead, use the IMS Universal drivers to develop Java applications that access IMS. The IMS Universal drivers, which are built on industry standards and open specifications, provide more flexibility and improved support for connectivity, data access methods, and transaction processing options.

To deploy an application that will use the classic Java APIs for IMS on WebSphere Application Server for z/OS, you must install the IMS DB resource adapter on WebSphere Application Server for z/OS, and configure both IMS the Open Database Access (ODBA) and the database resource adapter (DRA) interfaces.

The following figure shows an EJB application that is accessing IMS data. JDBC or IMS hierarchical database interface for Java calls are passed to the IMS DB resource adapter, which converts the calls to DL/I calls. The IMS DB resource adapter passes these calls to ODBA, which uses the DRA to access the DL/I region in IMS.




Figure 6. A WebSphere Application Server for z/OS EJB application using IMS DB resource adapter

Related tasks:

"IMS Universal drivers: WebSphere Application Server for z/OS type-2 connections" on page 6

Configuring WebSphere Application Server for z/OS for the classic Java APIs for IMS

To use the IMS DB resource adapter of the classic Java APIs for IMS with WebSphere Application Server for z/OS, you need to complete three steps: configure WebSphere Application Server for z/OS, install the IMS DB resource adapter, and install the custom service on WebSphere Application Server for z/OS.

Prerequisites:

Install the classic Java APIs for IMS.

Ensure that z/OS Resource Recovery Services is enabled.

Recommendation: The following information applies to the classic Java APIs for IMS, which are no longer recommended. Instead, use the IMS Universal drivers to develop Java applications that access IMS. The IMS Universal drivers, which are built on industry standards and open specifications, provide more flexibility and improved support for connectivity, data access methods, and transaction processing options.

The IMS DB resource adapter supports the 64-bit addressing mode of WebSphere Application Server for z/OS.

Requirements:

If you need to access IMS databases from WebSphere Application Server on a non-z/OS platform, you need to install and configure both the distributed version of WebSphere Application Server in the distributed environment and WebSphere Application Server for z/OS on the z/OS platform.

Configuring WebSphere Application Server for z/OS to access IMS

To use the IMS DB resource adapter to access IMS DB from WebSphere Application Server for z/OS, you first must configure WebSphere Application Server for z/OS to access IMS databases using ODBA.

ODBA uses the database resource adapter (DRA) interface to access IMS databases.

To configure WebSphere Application Server for z/OS to access IMS databases:

1. If not already done, create a DRA startup table. The DRA startup table module name must have the following naming convention:
 - Bytes 1-3: "DFS"
 - Bytes 4-7: 1- to 4-byte ID
 - Byte 8: "0"

Recommendation: Use the IMS ID as the 1- to 4-byte ID.

2. If not already done, link the DRA startup table into a load library.
3. Update the JCL for WebSphere Application Server for z/OS by adding to the STEPLIB the following data sets:

- The load library that contains the DRA startup table and the ODBA runtime code
 - The SDFSJLIB data set. This data set contains the DFSCLIB member.
4. Note the DRA name, which is defined by the MBR parameter. You will need to know bytes 4-7, which are usually the IMS system ID, when you install the data source.

Related concepts:

 Database resource adapter (DRA) (System Programming APIs)

Related tasks:

Chapter 42, “Accessing IMS databases through the ODBA interface,” on page 747

Installing the IMS DB resource adapter on WebSphere Application Server for z/OS

After you configure WebSphere Application Server for z/OS to have access to IMS databases, you must install the IMS DB resource adapter on WebSphere Application Server for z/OS.

Prerequisite: Perform the steps described in “Configuring WebSphere Application Server for z/OS to access IMS” on page 27.

To install the IMS DB resource adapter:

1. From the WebSphere Application Server for z/OS administrative console, click **Resources > Resource Adapters > Resource Adapters**. A list of resource adapters is displayed.
2. Click **Install RAR**. A dialog for installing the resource adapter is displayed.
3. Select **Server path** and type the path to the imsDBJCA.rar file:
pathprefix/ims/ims11/imsjava/classic/imsDBJCA.rar
4. Click **Next**. A configuration dialog is displayed.
5. Type the following information:
 - **Name:** A name for the resource adapter
 - **Classpath:** The Java library for IMS consists of the following JAR files. The path and file names are as follows:
 - *pathprefix/ims/ims11/imsjava/classic/imsJDBC.jar*
 - *pathprefix/ims/ims11/imsjava/classic/imsDBJCA.jar*
 - *pathprefix/ims/ims11/imsjava/classic/imsjavaBase.jar*
 - *pathprefix/ims/ims11/imsjava/classic/imsXQuery.jar*

The imsXQuery.jar file is required only if you intend to use XQuery.
6. Click **OK**. The IMS DB resource adapter is listed.
7. In the messages box, click **Save**. The save page is displayed.
8. Click **Save** to update the master repository with your changes.

Installing the custom service on WebSphere Application Server for z/OS

When WebSphere Application Server for z/OS is started, the custom service initializes the ODBA environment. When the server is stopped, the custom service terminates the ODBA environment. After a server is started, every application that is running in the server uses the initialized ODBA environment.

Prerequisite: Perform the steps described in “Installing the IMS DB resource adapter on WebSphere Application Server for z/OS.”

To install the custom service:

1. Modify the WebSphere Application Server for z/OS server.policy file, which is in the properties directory (for example, *installdir*/profiles/default/properties), by adding the following code:

```
grant codeBase "file:/pathprefix/ims/ims11/imsjava/classic/-" {  
    //Allows the IMS DB resource adapter and the custom service to read and  
    //write environment properties  
    permission java.util.PropertyPermission "*", "read, write";  
  
    //Allows the IMS DB resource adapter and the custom service to use  
    //the JavTDLI load library during runtime.  
    permission java.lang.RuntimePermission "loadLibrary.JavTDLI";  
};
```

2. In the left frame of the WebSphere Application Server for z/OS administrative console, click **Servers**, and then click **Application servers**. A list of application servers is displayed.
3. Click the name of the server on which you want to deploy your custom service.
4. Under Server Infrastructure, click **Administration**, and then click **Custom Services**. A list of custom services is displayed.
5. Click **New**. A configuration dialog is displayed.
6. Select the **Enable service at server startup** check box. If you do not select this check box, the custom service is not invoked when you start the server.
7. Type the following information:
 - **Classname:** com.ibm.connector2.ims.db.IMSJdbcCustomService
 - **Display Name:** A name for the custom service
 - **Classpath:** The path to the directory that contains imsDBJCA.jar and libJavTDLI.so. For a standard installation the path is '*pathprefix*/ims/ims11/imsjava/classic'
8. Click **OK**. The custom service is listed.
9. In the messages box, click **Save**. The save page is displayed.
10. Click **Save** to update the master repository with your changes.
11. Restart the server in order for the custom service to take effect.

Running the IMS IVP applications in the WebSphere Application Server for z/OS Java environment

Run the IMS IVP applications for the WebSphere Application Server for z/OS Java environment on WebSphere Application Server for z/OS to ensure that you have configured IMS and WebSphere Application Server for z/OS properly.

Prerequisites:

- Complete the steps in “Configuring WebSphere Application Server for z/OS for the classic Java APIs for IMS” on page 27 before running the IMS IVP applications.
- Ensure that the base IMS IVPs have been run. The IMS IVPs prepare the DBD for the IVP database, named IVPDB2, and load the IVP database. They also prepare the Java application PSB (named DFSIVP37), build ACBs, and prepare other IMS control blocks that are required by the IMS IVP applications for the WebSphere Application Server for z/OS Java environment.

The necessary files for the IVP applications are included in the IMS Java On Demand features FMID.

To run the IMS IVP applications in the WebSphere Application Server for z/OS environment:

1. Install the data source of the IMS IVP application as a new J2C connection factory. The JNDI name is `imsjavaIVP`. The database view name is `samples.ivp.DFSIVP37DatabaseView`.
2. Install the IMS IVP application EAR file. The file is *pathprefix/ims/ims11/imsjava/classic/samples/ivp/was/imsjavaIVP.ear*.
3. Test the IVP. The URL is `http://host_IP_address:port/IMSJavaIVPWeb/IMSJavaIVP.html`.

Related concepts:

 [IMS installation verification program \(IVP\) overview \(Installation\)](#)

Installing the data source for the IMS IVP application on WebSphere Application Server for z/OS

The `DataSource` facility is a factory for connections to a physical data source, or database. A data source is registered with a naming service based on the Java Naming and Directory (JNDI) API. `DataSource` objects have properties that pertain to the actual data source that an application needs to access.

Requirement: You must use the `DataSource` facility, which replaces the `DriverManager` facility, because the `DriverManager` facility is not supported by the Java EE Connection Architecture Specification.

To install the data source for the IMS IVP application:

1. In the left frame of the WebSphere Application Server for z/OS administrative console, click **Resources**, and then click **Resource Adapters**. A list of resource adapters is displayed.
2. Click the name of IMS DB resource adapter that you chose when you installed the adapter. A configuration dialog is displayed.
3. Under Additional Properties, click **J2C connection factories**.
4. Click **New**. A configuration dialog is displayed.
5. Type the following information:
 - Name:** name for the data source
 - JNDI Name:** `imsjavaIVP`
6. Click **OK**. The data source is listed in the J2C Connection Factories.
7. Click the name of the data source that you just installed.
8. Under Additional Properties, click **Custom properties**. Three properties are listed in a table.
9. Click **DatabaseViewName**. A configuration dialog is displayed.
10. In the **Value** field, type: `samples.ivp.DFSIVP37DatabaseView`
11. Click **OK**. The properties table displays the `DatabaseViewName` value that you just entered.
12. Click **DRAName**. A configuration dialog is displayed.
13. In the **Value** field, type bytes 4-7 of the DRA startup table module name (usually the IMS system ID).
14. Click **OK**. The properties table displays the DRA name that you just entered.
15. In the messages box, click **Save**. The save page is displayed.
16. Click **Save** to update the master repository with your changes.
17. Restart the server in order for the custom service to take effect.

Related tasks:

“Configuring WebSphere Application Server for z/OS to access IMS” on page 27

Installing the IMS IVP application on WebSphere Application Server for z/OS

This topic describes how to deploy the IMS IVP application on WebSphere Application Server for z/OS.

Prerequisite: You must complete the steps in “Installing the data source for the IMS IVP application on WebSphere Application Server for z/OS” on page 30 before installing the IMS IVP application.

To install the IMS IVP application:

1. From the WebSphere Application Server for z/OS administrative console, click **Applications**, and then click **Install New Application**. A dialog for installing the application is displayed.
2. Select **Remote file system** and type the path to IMSJavaIVP.ear:
pathprefix/ims/ims11/imsjava/classic/samples/ivp/was/imsjavaIVP.ear
3. Click **Next**.
4. Accept the defaults and click **Next**. The Install New Application wizard is started.
5. Click **Step 7** to accept the installation defaults. Depending on your specific server configuration, you might have to use the wizard to change some default values.
6. Verify that the options are correct, and then click **Finish**. A message is displayed that indicates first that the application is being installed, and then that the installation was successful.
7. Click **Save to Master Configuration**. The save page is displayed.
8. Click **Save** to update the master repository with your changes.

Testing the IMS IVP application on WebSphere Application Server for z/OS

This topic describes how to test the IMS IVP application on WebSphere Application Server z/OS.

Prerequisite: Perform the steps described in “Installing the IMS IVP application on WebSphere Application Server for z/OS.”

To test the IMS IVP application:

1. From the WebSphere Application Server for z/OS administrative console, click **Applications**, and then click **Enterprise Applications**. The application IMSJava IVP is listed with a red X, which indicates that the application is stopped.
2. Select **IMSJava IVP**.
3. Click **Start**. The application IMSJava IVP is listed with a green arrow, which indicates that the application is started.
4. Open a web browser.
5. Type the web address:
http://host_IP_address:port/IMSJavaIVPWeb/IMSJavaIVP.html

An input web page opens.

6. Click **Run the IVP**. If WebSphere Application Server for z/OS is configured properly, the IVP displays "The IVP was SUCCESSFUL" and the results of checks performed by the IVP. If WebSphere Application Server for z/OS is not

configured properly, the IVP displays "The IVP was NOT SUCCESSFUL" and the results of checks performed by the IVP.

Running your applications on WebSphere Application Server for z/OS

To run your application on WebSphere Application Server for z/OS, you must complete three basic steps.

Prerequisite: Perform the steps described in "Running the IMS IVP applications in the WebSphere Application Server for z/OS Java environment" on page 29.

To run your applications on WebSphere Application Server for z/OS:

1. Set the WebSphere Application Server for z/OS classpath to point to the IMS database metadata class.
2. Install the data source for the application as a J2C connection factory.
3. Install the application.

Related reference:

 Programming with the classic Java APIs for IMS (Application Programming)

Setting the WebSphere Application Server for z/OS classpath

Your application can include the IMS database metadata class (DLIDatabaseView subclass) or the metadata class can be stored elsewhere. If your application does not include the metadata class, you must set the WebSphere Application Server for z/OS classpath to the location of the IMS database metadata class that is used by the application.

One way to set the classpath is to add these files to the IMS DB resource adapter classpath.

To add the required files to the IMS DB resource adapter classpath:

1. From the WebSphere Application Server for z/OS administrative console, click **Resources**, and then click **Resource Adapters**. A list of resource adapters is displayed.
2. Click the name of the IMS DB resource adapter. A configuration dialog is displayed.
3. Click **OK**.
4. In the messages box, click **Save**. The save page is displayed.
5. Click **Save** to update the master repository with your changes.

For more information about the IMS database metadata class, see *IMS Version 11 Application Programming APIs*.

Installing the data source for your application on WebSphere Application Server for z/OS

The DataSource facility is a factory for connections to a physical data source, or database. A data source is registered with a naming service based on the Java Naming and Directory (JNDI) API. DataSource objects have properties that pertain to the actual data source that an application needs to access.

Requirement: You must use the DataSource facility, which replaces the DriverManager facility, because the DriverManager facility is not supported by the Java EE Connection Architecture Specification.

To install the data source for your application:

1. In the left frame of the WebSphere Application Server for z/OS administrative console, click **Resources**, and then click **Resource Adapters**. A list of resource adapters is displayed.
2. Click the name of the IMS DB resource adapter that you chose when you installed the adapter. A configuration dialog is displayed.
3. Under Additional Properties, click **J2C connection factories**.
4. Click **New**. A configuration dialog is displayed.
5. Type the following information:
 - Name**: name for the data source
 - JNDI Name**: the JNDI name for the data source.
6. Click **OK**. The data source is listed in the J2C Connection Factories.
7. Click the name of the data source that you just installed.
8. Under Additional Properties, click **Custom Properties**. Three properties are listed in a table.
9. Click **DatabaseViewName**. A configuration dialog is displayed.
10. Optional: In the **Value** field, type the fully-qualified DLIDatabaseView subclass name. If you do set the subclass name, you must either create a data source for every PSB an application accesses, or you must override the DLIDatabaseView subclass name in the DataSource object that is within the application by calling the setDatabaseView method and providing the fully-qualified name of the subclass. If you do not set the subclass name, you need to create a data source only for each IMS. In the application, define the DLIDatabaseView subclass name in the DataSource object by calling the setDatabaseView method and providing the fully-qualified name of the subclass.
11. Click **OK**. The properties table displays the DatabaseViewName value that you just entered.
12. Click **DRAName**. A configuration dialog is displayed.
13. In the **Value** field, type bytes 4-7 of the DRA startup table module name (usually the IMS system ID).
14. Click **OK**. The properties table displays the DRA name that you just entered.
15. In the messages box, click **Save**. The save page is displayed.
16. Click **Save** to update the master repository with your changes.
17. Restart the server in order for the custom service to take effect.

Related tasks:

“Configuring WebSphere Application Server for z/OS to access IMS” on page 27

Installing your application on WebSphere Application Server for z/OS

This topic describes how to deploy an application on WebSphere Application Server for z/OS.

Prerequisite: Perform the steps described in “Installing the data source for your application on WebSphere Application Server for z/OS” on page 32.

To install your application:

1. From the WebSphere Application Server for z/OS administrative console, click **Applications**, and then click **Install New Application**. A dialog for installing a new application is displayed.
2. Type the path to the EAR file.

3. Click **Next**.
4. Accept the defaults and click **Next**. The Install New Application wizard is started. Step 1, "Provide options to perform the installation," is displayed.
5. Click **Step 4: Provide JNDI Names for Beans**.
6. In the **JNDI Name** field, type the path to the EJB home interface.
7. Click **Step 7: Summary**. The options that you specified are displayed.
8. Verify that the options are correct, and then click **Finish**. A message is displayed that indicates first that the application is being installed, and then that the installation was successful.
9. Click **Save to Master Configuration**. The Save page is displayed.
10. Under **Save to Master Configuration**, click **Save**.

Related concepts:

 Java application development for IMS (Application Programming)

Part 2. CPI Communications and APPC/IMS

These topics introduce CPI-Communications and APPC/IMS. The topics discuss how CPI-Communications driven application programs function and how to administer APPC/IMS and use APPC/IMS with the CPI Communications interface to build CPI application programs.

Chapter 3. CPI Communications

This topic introduces CPI Communications driven application programs and distributed Syncpoint protected conversations.

CPI-C driven application programs

A CPI Communications driven application program can use IMS-managed resources in two ways: by using the SQL calls to access DB2 for z/OS through the IMS External Subsystem (ESS) Attach Facility and by using the APSB call to allocate IMS resources.

If you use the SQL calls to access DB2 for z/OS through the IMS External Subsystem (ESS) Attach Facility and the DB2 for z/OS resource translation table (RTT) is not used, the initial DB2 for z/OS plan name is the application program name. After the APSB call, the DB2 for z/OS plan name is the PSB name specified in the APSB call.

You can use the following SAA resource recovery calls when you want an application program to commit or back out changes to IMS or DB2 for z/OS resources:

- Use the Commit call (SRRCMIT) to commit changes.
- Use the Backout call (SRRBACK) to back out changes.

SAA resource recovery commit processing

An application program tells IMS to commit changes to database resources by issuing the SAA resource recovery call, SRRCMIT.

By issuing the SRRCMIT call, an application program tells IMS to commit changes to database resources:

- Issue the SRRCMIT call when the application program updates any IMS resources or accesses DB2 for z/OS resources.
- Reissue the SRRCMIT call after making any subsequent changes to an IMS or DB2 for z/OS resource.
- Issue the SRRCMIT call before terminating your application program.

If the application program terminates with any uncommitted changes to IMS resources, IMS attempts to commit these changes. Dangling conversations are deallocated abnormally.

When you issue the SRRCMIT call, IMS gets control and generates an internal sync-point call (if the conversation was not allocated with SYNCLVL=SYNCPT). All database changes are committed. All messages inserted to alternate PCBs (program control blocks) are sent to their final destination.

Normal termination

In IMS, normal termination occurs when an application program terminates without abending. For a CPI Communications driven application program, an implicit commit occurs.

Definition: An *implicit commit* occurs when the application program does not issue an explicit call to commit the current transaction, but one of the following occurs:

- The application program terminates normally.

Backout processing

Transaction updates can be backed out for a variety of reasons.

Transaction updates are backed out when any of the following occurs:

- Backout is issued by the application program.
- The application program terminates abnormally.
- The application program terminates normally, but the implicit commit fails.
- IMS resources are inflight during IMS system restart.

The backout consists of the following actions:

- All database updates are backed out.
- All messages inserted to non-express alternate PCBs are discarded.
- All messages inserted to express PCBs that have not been enqueued are discarded.
- The APPC/MVS verb ATBCMTP TYPE=ABEND is issued.¹

The application program tells IMS that a backout is required by issuing SRRBACK or by terminating abnormally.

Abnormal termination

When your application program abends, IMS backs out to the last IMS sync point.

IMS considers an application program to have terminated abnormally if either of the following occurs:

- The implicit commit fails.
- The application program abends.

Session failure

If any LU 6.2 session fails during the conversation, you can choose to end the application program or continue processing.

IMS TM is not involved and places no restrictions on your choice of committing or backing out updates. The application programmer makes this decision.

Because IMS TM is not informed of the session failure, it takes no action. The normal processing rules for commit and backout apply.

Return codes

Your application program receives return codes from IMS on the SAA resource recovery SRRCMIT and SRRBACK calls.

Your application program can receive the following return codes:

1. Issuing the verb ATBCMTP causes all LU 6.2 conversations associated with this TPI to terminate with CM_DEALLOCATE_ABEND.

RR_OK

The backout or commit operation completed successfully. All protected resources if backed out have been returned to their previous sync point; if they have been committed, they have advanced to a new sync point, and all changes made during the logical unit of work have been made permanent.

RR_PROGRAM_STATE_CHECK

A non-CPI Communications driven IMS application program issued an SAA resource recovery Commit call. No commit or backout has been performed.

RR_BACKED_OUT

A resource manager voted “no” during sync point processing. The sync point was initiated by the SAA resource recovery Commit call. The resource state is backed out for all resources.

System restart/resolve-in-doubt processing

After a system failure, a key part of restart processing is known as *resolve-in-doubt* processing. If the system fails, IMS determines whether to perform resolve-in-doubt processing for IMS-protected resources.

Examples of IMS-protected resources are:

- IMS DB databases
- DB2 for z/OS databases
- IMS TM message-queue messages

If the IMS system fails before the transaction completes phase one of the two-phase commit process (sync point), IMS backs out during IMS restart. Backout includes transactions that were processing at the time of failure.

If the transaction completes phase one of the commit process, resolve-in-doubt processing can take place during IMS restart. If only IMS resources are affected, commit processing occurs. If DB2 for z/OS resources are involved, resolve-in-doubt processing occurs between IMS and DB2 for z/OS.

No transactions using explicit CPI Communications driven application programs are preserved across an IMS restart.

CPI-C application program recovery

No recovery processing exists for application programs using the explicit CPI Communications driven interface.

IMS discards all CPI Communications driven transactions at restart regardless of their state at the time of failure. Application program designers should be aware of the SAA resource recovery resynchronization functions and consider the impact on their application program designs.

The application program should provide full integrity by issuing a SAA resource recovery Commit or Backout call for session failures. Application programs that require recovery assistance must be standard DL/I application programs.

Related reference:

 CALL statement (Application Programming APIs)

Programming requirements

The calls that initiate implicit sync point (DL/I GU to the message queue, CHKP, and SYNC) are invalid for CPI Communications driven application programs, and receive status AD (function parameter invalid). The CPI Communications driven application program activates IMS sync point processing by issuing the SRRCMIT and SRRBACK calls.

If you allocate a conversation with SYNCLVL=NONE or SYNCLVL=CONFIRM, include module DFSCPIR0 with your application program in the bind step. Including this module allows your application program to resolve the external references for SRRCMIT and SRRBACK.

No language-unique programming concerns exist in IMS for the SAA resource recovery interface.

Pseudonym files

APPC/IMS uses APPC/MVS services to provide SAA resource recovery support. APPC/MVS does not provide SAA resource recovery pseudonym files. However, you can create your own pseudonym files.

Related reading: For sample pseudonym files, see *SAA CPI Resource Recovery Reference*. These sample pseudonym files include examples on how to define working storage in the different languages.

Briefly, programmers of different languages need to define the following:

- IMS TM C programmers need to define z/OS as their operating system.
- IMS TM COBOL programmers must define their buffers in working storage.
- IMS TM FORTRAN programmers must define EXTERNAL statements for SRRCMIT and SRRBACK.

RRS and distributed syncpoint/protected conversations

z/OS Resource Recovery Services is the sync point manager, coordinating the update and recovery of multiple protected resources. RRS controls how and when protected resources are committed by coordinating with the resource managers, such as IMS, that have registered with RRS.

RRS supports the Common Programming Interface for Resource Recovery (CPI-RR), an element of the SAA CPI that specifies resource recovery and coordinates recovering local and distributed resources.

Definitions:

- A *protected resource* is a set of local or distributed data that is updated in a synchronized and controlled manner. In the APPC environment, a protected resource is a resource that is updated in an allocated conversation in which SYNCLVL=SYNCPT has been specified.
- A *resource manager* is a product, such as IMS, that owns protected data resources that are updated in an APPC conversational environment in which SYNCLVL=SYNCPT has been specified. IMS acts as a resource manager for DL/I data, Fast Path data, and the message queues.

The three participants in resource recovery include:

RRS (sync point manager)

Resource manager (such as IMS or DB2 for z/OS)

Application program

The following figure shows the three participants in the resource recovery process, and their interaction.




Figure 7. Participants in resource recovery

The two-phase commit protocol

The *two-phase commit protocol* is a process involving z/OS Resource Recovery Services (RRS) and the resource manager that ensures that updates made to a set of resources by an application program are either all made or none made.

The application program decides whether to commit its changes to the resources; this commit is made to RRS, which polls all of the resource managers as to the feasibility of the commit call. Each resource manager votes whether to commit the updates. This is called *phase 1* of the two-phase commit protocol.

After RRS has gathered the votes, *phase 2* begins. If all votes are to commit the updates, then the phase 2 action is to commit; otherwise, phase 2 results in a backout of the updates. System failures, communication failures, resource manager failures, or application program failures are not barriers to completing the two-phase commit protocol.

Definitions:

- A *unit of recovery* is a unit of work that spans one commit (synchronization) point to the next commit point.
- Units of recovery are termed *inflight* between the time they are created (or the previous sync point) until the resource manager votes to commit the updates. If the resource manager fails while units of recovery are inflight, the resource manager backs out all of the database updates on the subsequent start.
- Units of recovery are termed *indoubt* between the time when the resource manager votes to commit the updates and the time when the sync-point

manager calls the resource manager to do the commit. If IMS fails while units of recovery are indoubt, IMS holds the database updates until they are resolved.

Local-resource recovery versus distributed-resource recovery

In a *local-resource recovery environment*, the recovery participants reside on the same z/OS system. In a *distributed-resource recovery environment*, the recovery participants and the updated resources are scattered across multiple systems.

In a distributed-resource recovery environment, the APPC/PC (APPC/protected conversation) resource manager is used to provide the communications for the sync-point calls to remote systems.

The following figure illustrates how a distributed recovery environment operates.




Figure 8. Distributed resource recovery

IMS as a resource manager

A resource manager controls a protected resource.

In general, a resource manager does the following:

- Provides an application programming interface (API) to allow application programs to access its resources
- Logs changes to data before making the changes permanent
- Logs unit of work status
- Participates in the commit or backout actions for updated resources
- Contains recovery mechanisms to restore data to a previous state

For its own resources, IMS is both the sync-point manager and the resource manager.

Within the two-phase commit protocol, IMS must do each of the following:

- Register with z/OS Resource Recovery Services (RRS) as a resource manager
- Participate in the sync-point process
- Express interest in the unit of work
- Recover its unit of work status after an outage

Registration

A component of RRS provides registration services so that IMS can identify itself as a resource manager. By registering, IMS is provided a set of services to aid in maintaining resource consistency associated with the protected conversation.

IMS registers each time a control region is started for a DB/DC active system on a z/OS system with the recovery platform support. In an XRF environment, the active system registers during its restart. The alternate system registers at the time of takeover. No registration or participation occurs from an RSR tracking system.

Expressing interest

In addition to registering and supplying resource manager exit routines for specific stages of the two-phase commit protocol, IMS must also express interest in participating in the two-phase commit process for a particular unit of recovery.

Resolution of incomplete interests

In the event of an IMS or z/OS outage, during the IMS restart, the incomplete UR expressions of interest must be resolved.

RRS maintains unit of recovery information (such as identifier, state, and resource manager private data), which RRS presents to the restarting resource managers that previously expressed interest.

Sync-point participation

After IMS successfully registers and restarts, it supplies addresses of its exit routines to RRS. Several exit routines (such as prepare, commit, and backout) represent specific points in the two-phase commit protocol, which IMS can call to participate in the process.

Activating protected conversations

z/OS uses a construct with z/OS Resource Recovery Services (RRS) called a context.

Definition: A *context* is the entity for which resource managers perform services, to which they allocate resources and lock ownership, and in which they can express interest in participating in the protocol to ensure that the resource is updated in an orderly manner.

The type of context that the resource manager creates, owns, and manipulates is called the *private context*. A resource manager can create a context on behalf of another resource manager. RRS uses the private context to identify an application program's unit of work to maintain information for the resource manager concerning which of their resources are associated with the unit of work.

APPC as the communications manager

When APPC is the communications manager, RRS support is activated when a conversation is allocated with SYNCLVL=SYNCPT. This type of conversation is a protected conversation.

When SYNCLVL=SYNCPT is specified, APPC acquires a private context on behalf of IMS. IMS provides its resource manager name to APPC in its identity call. APPC provides the private context to IMS as the message header. IMS, using this context, then assumes the role of a participant in the two-phase commit process with the sync-point manager, RRS.

In addition to the SYNCLVL=SYNCPT, the keyword ATNLOSS=ALL must be specified in the VTAM definition file for whichever LUS need to be enabled for protected conversations.

Using OTMA with protected conversations

In an OTMA environment, OTMA is not a resource manager registered with RRS. The process remains an inter-process protocol between a server (IMS) and a number of clients (application programs). Therefore, OTMA cannot obtain a private context token to pass to IMS, as APPC does. The client-adaptor code that uses OTMA is responsible for obtaining and owning a private context, and for providing the context ID. In messages passed between the partners, the context-ID field contains the context token (if it is a protected conversation).

When IMS finds the context-ID in the message, IMS assumes the role of a participant in the two-phase commit process, as it does in the APPC environment.

XRF and protected conversations

Running protected conversations (using RRS with either APPC/PC or OTMA) in an IMS-XRF environment does not guarantee that the alternate system can resume and resolve any unfinished work started by the active system. A failed resource manager must re-register with its original RRS system if the RRS system is still available when the resource manager restarts. Only if the RRS on the active system is not available can an XRF alternate system register with another RRS in the sysplex and obtain the incomplete unit of recovery data of the failing active system.

Recommendation: Because IMS retains indoubt units of recovery until they are resolved, switch back to the active system as soon as possible to obtain the unit of recovery information and to resolve and complete all the work of the resource managers.

RSR and protected conversations

Active systems tracked in an RSR environment by a remote system can process protected conversations with APPC/PC or OTMA, although it is necessary to resolve indoubt units of recovery using commands after a takeover. Most likely, the remote site is not part of the active sysplex, and no means exists for the new IMS system to acquire unfinished unit of recovery information from RRS. IMS provides commands to investigate protected conversation work and to resolve it if necessary.

Chapter 4. Administering APPC/IMS and LU 6.2 devices

This topic introduces APPC/IMS and describes how to administer APPC/IMS and LU 6.2 devices.

APPC/IMS overview

APPC/IMS, a part of IMS TM, lets you use the CPI Communications interface to build CPI application programs.

APPC/IMS allows distributed and cooperative processing between IMS and systems that have implemented APPC as shown in the following figure. APPC/IMS delivers support for APPC with facilities provided with APPC/MVS. (The APPC/IMS interface is provided by APPC/MVS and supports the CPI Communications interface. IMS TM supports the CPI resource recovery interface.) APPC/IMS supports the CPI resource recovery Commit (SRRCMIT) and Backout (SRRBACK) calls for IMS-managed local resources. These protected resources include:

- IMS TM message-queue messages
- IMS DB databases
- DB2 for z/OS databases

APPC/IMS also supports the existing IMS DL/I application programming interface (API) enabling application programs to use LU 6.2 communications without the function of the CPI Communications interface. This allows most existing applications to continue to function with the APPC/IMS support of LU 6.2.




Figure 9. APPC support for IMS

Definitions:

- Within the context of administering IMS TM, “transaction programs,” “applications,” “application programs,” and “programs” are synonymous.
- Within the context of administering APPC/IMS and LU 6.2 devices, “APPC application programs” are synonymous with “LU 6.2 application programs.”
- “LU 6.2 transactions” are those that originate from an LU 6.2 application program.

Recommendations: For APPC/IMS, do the following:

- Schedule your IMS standard or modified application programs entered from LU 6.2 remote systems using MSC. Be aware that CPI-C driven application programs cannot have transactions that execute on remote systems.
- Define your APPC/IMS LUs for use by APPC/MVS, as well as by any APPC application program.
- Use the LTERM and the MOD name in the first segment of your message. Use the LTERM to change the destination for your output to a non-LU 6.2 device. Use the MOD name to format error messages.
- Use a network-qualified LU name. You do not need to have unique names for the LUs on different systems.

IMS dependent regions are automatically defined to APPC as subordinate address spaces of the IMS Scheduler. An IMS BMP cannot be defined as an ASCH controlled application. It may use explicit conversation services through the IMS base LU.

IMS manages the APPC/IMS message buffers automatically; no definition is necessary. No special considerations are needed for EMH.

APPC/IMS application program interface

APPC/IMS has two distinct application program interfaces (APIs): the implicit and explicit interfaces. The same application program can use both APIs.

Implicit API

The *implicit API* is an extension of the IMS standard DL/I API (call xxxTDLI). It allows IMS application programs to communicate with LU 6.2 application programs without being sensitive to LU 6.2 protocols and without requiring the programmer to have any knowledge of LU 6.2. APPC/IMS provides functions not normally available to an LU 6.2 application program: message queuing, and automatic asynchronous message delivery and recovery. Existing IMS transactions use the implicit API to communicate with APPC.

Implicit API messages are placed on the IMS message queues or the Fast Path expedited message handling (EMH) buffers for Fast Path transactions. The originating IMS determines whether to mark the input messages as discardable or nondiscardable.

When the implicit API is used, IMS issues all required CPI Communications calls. The application program interacts strictly with the IMS message queues or the Fast Path EMH buffers.

Explicit API

The *explicit API* is the CPI Communications API and is available to any IMS application program. The application program makes calls to APPC using the CPI Communications interface without using IMS. These CPI calls are handled directly by APPC/MVS. Messages sent or received by the CPI Communications interface are not stored on the IMS message queues or in the EMH buffers, and these messages are not available for transaction restart. No IMS-provided functions are involved for these messages.

Alternatively, you can also use z/OS the ATBxxxx calls of the APPC/MVS TP services. For information on using these calls, see *z/OS MVS Programming: Writing Transaction Programs for APPC/MVS*.

APPC/IMS application programs

APPC/IMS has three different types of application programs: standard, modified, and CPI Communications driven.

The application programs are defined as:

Standard

No explicit use of CPI Communications facilities.

Modified

Uses the I/O PCB to communicate with the original input terminal. Uses CPI Communications calls to allocate new conversations and to send and receive data.

CPI Communications driven

Uses CPI Communications calls to receive the incoming message and to send a reply on the same conversation. Uses the DL/I APSB call to allocate a PSB to access IMS databases and alternate PCBs.

You can schedule your standard and modified application programs locally and remotely using MSC. The logic for local application programs differs from the logic for remote application programs. In the following topics, the differences are described.

Standard IMS application programs

Standard IMS application programs use the existing IMS call interface.

Application programs that use the IMS standard API can take advantage of the LU 6.2 protocols. Standard IMS application programs use a DL/I GU call to trigger a sync point and to get the incoming transaction. These standard IMS application programs also use DL/I ISRT calls to generate output messages to the same or different terminals, which can be LU 6.2 terminals. (A non-message-driven BMP is considered a standard IMS application program when it does not use the explicit API.) The identical program can work correctly for both LU 6.2 and non-LU 6.2 terminal types. IMS generates the appropriate calls to APPC/MVS services.

IMS provides the following services for standard IMS application programs:

- Receives incoming transaction from an LU 6.2 application program
- Calls the Input Message Routing exit routine
- Schedules transactions into local and remote IMS dependent regions
- Provides necessary transaction recoverability
- Provides necessary transaction rollback and retry
- Provides integration of IMS-controlled conversation flows with database updates during sync point for all APPC Sync_Level options (NONE, CONFIRM, SYNCPT)
- Provides all needed LU 6.2 calls to APPC/MVS services
- Sends either synchronous or asynchronous output to an LU 6.2 application program
- Keeps asynchronous output on IMS message queue until successful transmission

- Allocates new LU 6.2 conversations for messages inserted to alternate PCBs using the DL/I ISRT call

Existing application programs that are sensitive to a terminal's hardware characteristics, such as cursor position or MFS format names, might need to be changed to communicate with LU 6.2 application programs.

Restrictions:

1. If a LU 6.2 synchronous conversation implicit transaction initializes other transactions (program-to-program switch), an express PCB can not be used to do the ISRT. An express PBC causes race conditions to occur and the output may randomly return to the inputting terminal on a new asynchronous conversation with TPNAME DFSASYNC. The original conversation may not be deallocated.
2. If a transaction initializes more than one child transaction, which in turn may initialize other transactions, and one of the child transactions provides the response, the result is unpredictable.

Depending on the execution sequence of these transactions the LU can receive a DFS2082 message with the response sent to the default TP name DFSASYNC or the LU receives the response and no DFS2082 message is issued.

MSC and standard IMS application programs

When an APPC application program enters an IMS transaction that executes on a remote IMS, an LU 6.2 conversation is established between the APPC application program and the local IMS.

The local IMS is considered the partner LU of the LU 6.2 conversation. The transaction is then queued on the remote transaction queue of the local IMS. From this point on, the transaction goes through normal MSC processing. After the remote IMS executes the transaction, the output is returned to the local IMS, and is then delivered to the originating LU 6.2 application program.

The originating (local) IMS provides the following services:

- Receives incoming transaction from an LU 6.2 application program
- Calls the Input Message Routing exit routine
- Queues the transaction to its remote transaction queue
- Sends the transaction across the MSC link
- Receives the transaction response
- Sends either synchronous or asynchronous output to an LU 6.2 application program

The remote IMS provides the following services for the remote standard application program:

- Receives the incoming transaction from the partner IMS (originating or intermediate IMS) over the MSC link
- Schedules transactions into dependent regions
- Commits database changes at sync point
- Provides necessary transaction recoverability
- Provides necessary transaction rollback and retry
- Keeps transaction output on the IMS message queue until the transmission is successful

- Returns the transaction output to the local IMS over the MSC link

Restriction: MSC is not supported if the originating LU 6.2 conversation is allocated with SYNCLVL=SYNCPT.

Modified IMS application programs

Modified IMS application programs use a DL/I GU call to retrieve the incoming transaction and to trigger a sync point.

These modified IMS application programs also use DL/I ISRT calls to generate output messages to the same or different terminals, which can be LU 6.2 terminals.² Unlike standard IMS application programs, modified IMS application programs use CPI Communications calls to allocate new conversations, and to send and receive data. IMS has no direct control of these CPI Communications conversations.

Modified IMS transactions are indistinguishable from standard IMS transactions until program execution. In fact, the same application program can be a “standard IMS” application on one execution, and a “modified IMS” application on a different execution. The distinction is simply whether the application program has used CPI Communications resources.

IMS provides the following services for modified IMS application programs:

- Receives incoming transactions from LU 6.2 application programs
- Schedules transactions into local and remote dependent IMS regions
- Appropriate transaction recoverability before transaction scheduling
- Provides integration of IMS-controlled conversation flows with database updates during sync point for APPC Sync_Level options (NONE, CONFIRM, SYNCPT)
- Provides all necessary LU 6.2 calls to APPC/MVS services for IMS-controlled LU 6.2 conversations
- Sends either synchronous or asynchronous output to LU 6.2 application programs
- Keeps asynchronous output on the IMS message queue until successful send occurs
- Allocates new LU 6.2 conversations for any messages inserted to alternate PCBs using the DL/I ISRT calls

IMS does not provide any services for conversations explicitly allocated by the application program. Explicitly allocated conversations need to be deallocated if a program abend occurs.

MSC and modified IMS application programs

When an APPC program enters an IMS transaction that executes on an MSC remote IMS, an LU 6.2 conversation is established between the APPC program and the local IMS.

The local IMS is considered the partner LU of the LU 6.2 conversation. The transaction is then queued on the remote transaction queue of the local IMS. From this point on, the transaction goes through normal MSC processing. After the remote IMS executes the transaction, the output is returned to the local IMS, and then delivered to the originating LU 6.2 program.

2. A non-message-driven BMP is considered a modified standard IMS application program when it uses the explicit API.

The originating (local) IMS provides the following services:

- Receives incoming transaction from an LU 6.2 application program
- Calls the Input Message Routing exit routine
- Queues the transaction to its remote transaction queue
- Sends the transaction across the MSC link
- Receives the transaction response
- Sends either synchronous or asynchronous output to an LU 6.2 application program

The remote IMS provides the following services for the remote modified application program:

- Receives the incoming transaction from the partner IMS (originating or intermediate system) over the MSC link
- Schedules transactions into dependent regions
- Appropriate transaction recoverability before transaction scheduling
- Commits database changes at sync point
- Provides necessary transaction recoverability
- Provides necessary transaction rollback and retry
- Keeps transaction output on the IMS message queue until the transmission is successful
- Returns the transaction output to the local IMS over the MSC link

Restriction: MSC is not supported if the originating LU 6.2 conversation is allocated with SYNCLVL=SYNCPT.

CPI Communications driven application programs

CPI Communications driven application programs are defined only in the APPC/MVS TP_Profile data set; they are not defined to IMS.

The CPI Communications driven application program definition is dynamically built by IMS when a transaction is presented for scheduling by APPC/MVS based on the APPC/MVS TP_Profile definition after IMS restart. The definition is keyed by TP name. APPC/MVS manages the TP_Profile information.

When a CPI Communications driven transaction program requests a PSB, the PSB must already be defined to IMS by using the APPLCTN macro during system definition and by generating the appropriate PSBs and ACBs with the Program Specification Block (PSB) generation utility and the Application Control Blocks Maintenance utility when APPLCTN PSB= is specified. When APPLCTN GPSB= is specified, generating PSBs and ACBs is not required.

CPI Communications driven application programs must use CPI Communications calls to accept the incoming conversation and to send a reply on the same conversation. The DL/I GU call is not used to retrieve the initiating transaction from the LU 6.2 application program. No IMS resources are allocated when the application program is scheduled. Instead, the application program can use the DL/I APSB call to allocate a PSB that provides access to IMS databases and to alternate PCBs. A CPI Communications driven application program can send messages to other terminals (either LU 6.2 or non-LU 6.2) or other IMS transactions (either local or remote) by inserting to an alternate PCB, after allocating the appropriate PSB. Both the explicit and implicit API can be used on the same application program.

IMS provides the following services for CPI Communications driven application programs:

- Schedules the transaction.

IMS does not receive input before scheduling. It does not interact with the conversation at any time other than to possibly reject the inbound allocate request. If IMS rejects the inbound allocate request, the transaction is not scheduled.

- Provides sync point of local resources.
- Schedules PSB if called by application program.
- Processes calls to alternate or database PCB made by the application program.

Related concepts:

“RRS and distributed syncpoint/protected conversations” on page 40

Chapter 3, “CPI Communications,” on page 37

 Designing an application for APPC (Application Programming)

Using the MOD name and LTERM interface

Your LU 6.2 application program can use an interface to emulate MFS.

For example, the application program can use the MOD name to communicate with IMS to specify how an error message could be formatted. For non-LU 6.2 application programs, the MOD name is given to the MFS formatting modules in IMS; for LU 6.2 application programs, the MFS modules are not called, and the MOD name is given to the LU 6.2 Edit exit routine (DFSLUEE0) as a parameter. The LU 6.2 Edit exit routine can do whatever the programmer specifies with the MOD name, such as format an error message.

Your LU 6.2 application program uses the LU name to send data to an LU 6.2 application program. However, if you want to send data to a non-LU 6.2 device such as a printer, you can use the LTERM instead of the LU name.

The Initialization exit routine (DFSINTX0) can be used to create a user table of MOD names that you might want to use for formatting messages, and LTERMs that you might want to use as printers. This user table can be used by DFSLUEE0 to find the appropriate MOD name or LTERM for your application program.

LU 6.2 application programs can send both the LTERM and the MOD name in the first segment of the message. The LU 6.2 Edit exit routine (DFSLUEE0) checks the contents of the first message segment. Based on the information it finds in a user table, the exit routine decides whether to return the LTERM and the MOD name to IMS. IMS saves the LTERM and the MOD name in the I/O PCB. For formatting output, IMS provides the address of the MOD name in the first segment of the message to the LU 6.2 Edit exit routine (DFSLUEE0). For changing the destination to a non-LU 6.2 device, IMS provides the LTERM in the first segment of the message to the LU 6.2 Edit exit routine (DFSLUEE0). The Initialization exit routine (DFSINTX0) can be used to create the user table. This exit routine must pass the address of the user table to IMS, and IMS passes the address to DFSLUEE0.

Establishing APPC/IMS

Before activating APPC/IMS, an IMS system definition is needed to specify 390 as the third parameter of the SYSTEM keyword for the IMSCTRL macro.

CPI Communications driven application programs and LU 6.2 application programs cannot be defined in a system definition. LU 6.2 application programs are only defined to VTAM.

Start APPC/IMS by specifying APPC=Y on the IMS startup parameter. The default is APPC=N. When 'N' is specified, a connection to APPC/MVS services is not established during IMS initialization. When 'Y' is specified, IMS establishes a connection with APPC/MVS during IMS initialization. The /START APPC command overrides APPC=N.

TP_Profile

The TP_Profile is a VSAM data set owned by APPC/MVS and maintained by the APPC/MVS Administration utility (ATBSDFMU) or by the administrator using TSO/ISPF dialogs. The purpose of the TP_Profile entries is to provide attribute information for TP names.

CPI Communications driven application programs must be defined in the APPC/MVS TP_Profile. IMS system-defined transaction codes can optionally be defined in a TP_Profile. The existence of an IMS definition (in the IMS GEN or by online change) causes the transaction to be considered a standard DL/I or modified-standard application.

The TP_Profile, an APPC/MVS resource, contains definitions of transaction program names (TPNs) and their characteristics. You can define a TP_Profile to schedule an IMS transaction program that uses a transaction code that is different from the TPN.

IMS uses the TP_Profile to establish transaction scheduling characteristics for CPI Communications driven application programs. Based on the IMS dependent section of the APPC/MVS TP_Profile definition, IMS dynamically defines these characteristics when a transaction is presented for scheduling after restart by APPC/MVS.

CPI Communications driven transaction programs are defined only in the TP_Profile. The definition is by TPN. The same TPN can be defined differently for different LU names by using a different TP_Profile data set. The LU name is associated with an IMS.

The default TP_Profile data set name is SYS1.APPCTP. The LUADD TPDATA option in the SYS1.PARMLIB (APPCPMxx) member specifies the TP_Profile data set name used for this LU.

Use TP_Profile dialog or the APPC/MVS administration utility (ATBSDFMU) to define a TP_Profile.

Example: The following figure is an example of the IMS-specific area of the TP_Profile definition (Panel 1).

```

----- IMS_TP_Profile Panel -----

TP Name . . . : INQUIRY_Part

Transaction Code . . . . . PART
Security Type . . . . . _____ (NONE,CHECK,FULL, default=CHECK)

CPI Communications Driven Options
  Transaction Class . . . _____ (Range 1 - 999, default=1)
  Maximum Regions . . . . _____ (Range 0 - 999, default=1)

Comments . . . (Optional 1 to 10 lines)
> TP_PROFILE Created 10/8/91 <
> Access IMS Sample Parts DATABASE via. program DFSSAM02 <
> _____ <
> _____ <
> _____ <
> _____ <
> _____ <
> _____ <
> _____ <
> _____ <
> _____ <

PF01 = Help    PF03 = Exit    PF12 = Cancel    Enter = Accept

```

Figure 10. IMS-specific TP_Profile panel 1

To maintain IMS TP_Profiles using ISPF, do each of the following:

1. Enter TSO ICQASRM0 from the TSO command line of the TSO/E to start the ISPF TP_Profile System Data Facility Maintenance Utility from a TSO user ID. If this utility is not available, contact your z/OS system programmer.
2. Enter S next to the TP_Profile selection and the TP_Profile data set name specified on the TPDATA keyword on the LUADD statement for your IMS LU. (The LUADD statement is in the APPCTPxx PARMLIB member, where xx is the APPC suffix.)
3. A list of TP_Profiles is displayed. Select A to add a new TP_Profile or E to edit an existing TP_Profile. If you are adding a TP_Profile, you must supply a scheduler name. This name was set at IMS installation time. The recommended name is IMS .
4. After the general TP_Profile characteristics are supplied, the ISPF editor panel is displayed. Enter DFSTPROF on the command line to display the IMS TP_Profile Maintenance panel.
5. Supply IMS scheduler-dependent characteristics. Press enter to save your changes or PF3 or PF12 to cancel your changes. You can press PF1 for online help on fields supplied in this panel.

The TP_Profile name is not available on all releases of TSO/E, so a value of “Not Available” is displayed. This does not suggest that a problem exists.

Related reading: For more information about this utility, see *z/OS MVS Programming: Writing Transaction Programs for APPC/MVS*.

APPC/MVS Administration utility (ATBSDFMU) example

The following example is an APPC/MVS Administration utility (ATBSDFMU) entry.

```

TPADD TPSCHED_EXIT(DFSTPPE0)
      TPNAME(INQUIRY_PART)
      SYSTEM
      ACTIVE(YES)
      TPSCHED_DELIMITER(##)

```



```

TRANCODE=PART
CLASS=1
MAXRGN=1
                                CPUTIME=0
##

```

In this example, the IMS section starts with TRANCODE=PART. The other control statements are shown for completeness.

The IMS TP_Profile parsing module, DFSTPPE0, performs validity checking and parses the input data in IMS. This module should be loaded into the STEPLIB data set of the step that adds the TP_Profile. The APPC/MVS Administration utility (ATBSDLFNU) requires that STEPLIB be APF authorized.

The following five keywords are used to add an IMS section to the TP_Profile entry. The keyword-parameter sets must be separated by one or more blanks. The keyword-parameter sets must be specified between columns 1-72. An asterisk (*) in column 1 indicates a comment.

TRANCODE= 1 - 8 characters

Name of the IMS transaction code associated with this TP name consisting of alphanumeric or '#', '\$', '@'. IMS translates the TP name to the TRANCODE. IMS scans for valid characters (00640 character set). If invalid characters exist, IMS uses the default transaction code, IMSTRAN, instead of a transaction code with the non-00640 characters.

CLASS= 1 - 999

Specifies the class used for scheduling. The default value is 1.

Recommendation: Define CPI transactions with a different message class from that used for non-CPI transactions. IMS handles all CPI transactions as priority zero within the transaction class.

MAXRGN= 0 - 999

Restricts the number of dependent regions that this CPI Communications driven transaction program can use. The default value is 1.

RACF=NONE, CHECK, or FULL

RACF=NONE causes IMS to call the Transaction Authorization exit routine (DFSCTRNO).

RACF=CHECK causes IMS to call RACF for security checking when IMS receives a transaction (using RCLASS of TIMS or CIMS), but does not clone the security environment into the dependent region when the transaction executes.

RACF=FULL clones the security environment to the dependent region at execution time. Specifying this parameter and issuing the IMS command /SECURE APPC PROFILE enables APSB SAF Security for this CPI-C application program.

CPUTIME= 0 - 1440

Specifies the number of CPU seconds that the CPI-C program is allowed to use. If it exceeds the limit, it is terminated with ABENDU0240. This time limit protects against program loops, which locks resources from other applications. The default is 0, which is no limit.

You can use the TP_Profile entry in two ways:

- To specify an IMS transaction code that is defined in the IMS. The CLASS and MAXRGN parameters in the TP_Profile are ignored and the transaction values in

IMS remain unchanged. The TP_Profile entry provides mapping for a 64-character TPN into an 8-character transaction code.

- To specify an IMS transaction code that is not defined in the IMS. The IMS transaction code is a CPI Communications driven transaction, and is used as the load module name of the scheduled application program and the dynamically built transaction name.

When a TP_Profile is not defined, IMS uses the first 8 bytes of the TPN translated to the IMS character set as the transaction code.

The allocate request is rejected if the transaction code is not valid.

Related reading: For more information about using the APPC/MVS Administration utility (ATBSDFMU), see *z/OS MVS Planning: APPC/MVS Management*.

Outbound LU specification

You can specify a defined APPC LU as the outbound LU.

The default setting for defined APPC LUs is BASE LU. Changing an outbound LU is useful because, when the outbound LU has status of disabled, IMS does not allocate the APPC conversation.

The outbound LU must be defined in the APPCPMxx member of the SYS1.PARMLIB library. To specify an LU as the outbound LU, use the OUTBND= parameter in the DFSDCxxx PROCLIB member. You can set the outbound LU by using the /CHANGE APPC OUTBND command. However, a restart sets the outbound LU to the value in the DFSDCxxx member, if specified. If it is not specified, the outbound LU is set to BASE LU.

Outbound side information

APPC/MVS *side information* supplies destination information, such as the name of the partner program, the name of the LU at the partner node, and the logon mode name.

CPI Communications provides a way to use system-defined values for these required fields; these system-defined values are the side information. This information can be used by an IMS application program allocating (establishing) an APPC conversation using CPI Communications, an IMS LU 6.2 descriptor, a DL/I change call (CHNG), or a DFSAPPC message switch.

System programmers supply and maintain the side information for CPI Communications programs.

Side information is accessed by a *symbolic destination name*. The symbolic destination name, called *sym_dest_name* within the context of administering IMS TM, corresponds to an entry in the side information file containing the following information:

partner LU name

Shows the name of the LU where the partner program is located. This LU name is any name for the remote LU that is recognized by the local LU for allocating a conversation. An example is a USERVAR name.

This LU name can be a 17-byte network-qualified LU name.

logon mode name

Used by LU 6.2 to designate the properties for the session that will be allocated for the conversation. The properties include the class of service to be used on the conversation. The network administrator defines a set of mode names used by the local LU to establish sessions with its partners. The system programmer uses one of these values in a side table entry. An invalid mode name prevents a conversation from being allocated.

TP name

Transaction program (TP) name specifies the name of the remote application program.

IMS and z/OS do not accept blank *sym_dest_name* values on the Initialize_Conversation call.

The default name for the side information file is SYS1.APPCSI. Define this file in the SYS1.PARMLIB(APPCLMxx) as shown in the following example.

```
SIDEINFO
  DATASET(SYS1.APPCSI)
```

The destination name, partner LU name, mode name, and TP name can be defined using the APPC/MVS Administration utility (ATBSDFMU) as shown in the following example.

```
SIADD
  DESTNAME(DESTX)
  TPNAME(LU62USER_TPX)
  MODENAME(APPCMODE)
  PARTNER_LU(APPCLUX)
```

Related reading: For more information on APPC calls, see *CPI Communications Specification*.

PARMLIB member

The APPC address space uses the APPCLMxx member of SYS1.PARMLIB. Define IMS as a local APPC component LU that is controlled by the APPC address space.

The scheduler name is the same IMSID used in the IMSCTRL macro. When IMS identifies to APPC, it passes its IMSID as the scheduler name SCHED(IMS1) in APPC member APPCLMxxx. The following is an example of the APPCLMxx member:

```
LUADD
  ACBNAME(IMSLU62)
  SCHED(IMS1)
  BASE
  TPDATA(SYS1.APPCTP)
  TPLEVEL(SYSTEM)
```

For XRF add:

```
USERVAR=uservar_name ALTLU=luname
```

The LUADD option keywords are defined below:

ACBNAME=local LUNAME of IMS

SCHED=IMS id

BASE*mandatory parameter*

TPDATA(*TP_Profile dataset name*)

TPLEVEL(*system*) **suggested value**

USERVAR=(*uservar_name*)

ALTLU=(*LUNAME*)

Related reading: For more information on these keywords, see z/OS MVS Planning: APPC/MVS Management.

Communication between VTAM and its application programs requires an ACB (application control block) whose name must be identically defined in the SYS1.VTAMLST APPL statement and in the APPCPMxx LUADD statement ACBNAME parameter.

APPC manages the IMS ACB. When IMS identifies to APPC, APPC gives IMS the name of the APPC-managed ACB name (LUNAME). The APPC LUNAME is not defined in IMS, because IMS does not manage the ACB. The entries in the SYS1.PARMLIB member APPCPMxx include both the IMS scheduler name (IMSID) and the LUNAME ACBNAME(xxxxxxx) that ties an IMS to an LUNAME.

This ACBNAME must be different from the ACBNAME used by IMS for non-LU 6.2 terminals. APPC/MVS expects its LUs to be defined as VTAM resources so that these LUs can access the SNA network. A VTAM application program (APPL) definition macro must be coded for each APPC/MVS LU. LU 6.2 application programs are only defined to VTAM, not to IMS. The SYS1.VTAMLST member example follows:

```
IMSLU62 APPL ACBNAME=IMSLU62
        APPC=Y
        ...
```

APPC/MVS Timeout Service

Using APPC/MVS Timeout Service, you can indicate the maximum time interval an application waits before terminating a conversation and regaining control from APPC/MVS callable services.

When APPC/MVS does not respond to an APPC call, due to a network delay for instance, the dependent region hangs and the caller cannot regain control.

The timeout feature is activated at startup by specifying the APPCIOT=(*mmmm:ss,mmmm:ss*) parameter in the DFSDCxxx member of the IMS.PROCLIB data set. The APPC time-out values are specified in minutes (*mmmm*) and seconds (*ss*). Valid values for *mmmm* are 00 to 1440. Valid values for *ss* are 00 to 59. If APPCIOT=00, there is no time-out detection. When a transaction is terminated due to timeout, messages DFS1965E and DFS1959E are sent to the MTO terminal and the z/OS console. The timeout value can be changed using the /CHANGE command.

For synchronous APPC conversations, if APPC timeout is active, then IMS uses ATBSTO6 service (SET_TIMEOUT_VALUE) to set the timeout value for each conversation.

For asynchronous APPC conversations, if APPC timeout is active, then IMS sets the timeout value when the conversation gets allocated. In either case, the timeout value is active until the conversation is deallocated, which occurs, in the case of IMS conversational transactions, when the IMS conversation ends.

Common Programming Interface Communications (CPI-C) transactions are not automatically supported by APPC/MVS Timeout service, but can exploit APPC/MVS Timeout service using ATBSTO5 service provided the proper coding is supplied.

Related Reading:

- For more information on programming MVS services, see *z/OS MVS Programming: Writing Transaction Programs for APPC/MVS*.

Related reference:

 DFSDCxxx member of the IMS PROCLIB data set (System Definition)

APPC/MVS Error Extract Service

Whenever an APPC/MVS service call returns an unexpected return code, IMS issues APPC/MVS Error Extract Service call ATBEES3 with a DFS1995E prefix.

Related reading: For more information on ATB return codes, see:

- *z/OS MVS System Messages, Vol 3 (ASB-BPX)*
- *z/OS MVS Dump Output Messages*

Initializing and changing LU 6.2 descriptors

LU 6.2 descriptors allow the system programmer to specify an LTERM that associates an output destination with an LU 6.2 application program. This allows the system programmer to change the application program's destination using alternate PCBs to LU 6.2 application programs, without requiring any application program coding changes.

LU 6.2 descriptors are optional, but they are required if you want to dynamically create queue control blocks and define processing options.

The application program uses an LTERM name as a symbolic destination; only the system programmer needs to be aware of the actual term associated with this name.

The LU 6.2 descriptor entry contains:

- APPC/MVS side information entry name; this parameter can be omitted
- APPC conversation type (BASIC or MAPPED)
- APPC Sync_Level options (NONE, CONFIRM)
- LTERM name
- LU name of the destination of the LU 6.2 application program (overrides side information); this can be a network-qualified LU name up to 17 bytes in length
- The name of the local LU that IMS uses to allocate the outbound conversation, specified in the OUTBND parameter

- TP name to be scheduled (overrides side information)
- VTAM mode table entry to be used (overrides side information)

Here's an example of an LU 6.2 descriptor:

```
U L62TERM1 LUNAME=L62IMS1 TPNAME=CPICTRN1 MODE=L62MDE02
U L62TERM1 SYNCLEVEL=N OUTBND=MYLU02
```

Do not code a parameter and leave it blank (such as SIDE=b), or an error message is issued. Instead, omit the parameter completely.

These LU 6.2 descriptor LTERMs are only for output and are never used by IMS as an LTERM name associated with an input message. DFSAPPC is an IMS-reserved name for the message switch facility.

The LU 6.2 descriptors are built as specified in IMS PROCLIB member DFS62DTx during IMS initialization. They can be added, deleted, or changed without restarting IMS. You can specify any number of descriptors. If an error occurs, the z/OS system console and the IMS JOBLIB record the error messages. IMS initialization continues regardless of any errors during descriptor initialization.

To add descriptors while IMS is running, you must first define the LU 6.2 descriptors in PROCLIB member DFS62DTx. Load the LU 6.2 descriptor from the IMS PROCLIB using the /START DESC command. To delete descriptors, use the /DELETE DESC command. To change descriptors, use the /CHANGE DESC command.

Related reading: For more information on coding these parameters, see *IMS Version 11 System Definition*.

Using MSC in an APPC/IMS environment

APPC/IMS uses the services of APPC/MVS and MSC to provide the communication interface for an MSC configuration.

Together, MSC and APPC/IMS allow:

- LU 6.2 programs to use the TP name of an IMS remote standard application program or an IMS remote modified application program. (The transaction is sent to the remote IMS and executes. The transaction's reply is sent across the MSC links to the local IMS and then on to the LU 6.2 application program.)
- A message switch to a remote logical terminal (LTERM) through the DFSAPPC System Service.
- Use of DFSAPPC for sending IMS remote transactions and data.
- Immediate or deferred program-to-program switching to an MSC-routed remote application program.

CPI Communications driven application programs cannot include transactions that execute on remote IMS systems.

All IMS transaction types except Fast Path are supported: conversational, nonconversational, response mode, and nonresponse mode.

IMS adds a prefix to the LU 6.2 message when the message is sent over an MSC link. The minimum size of this prefix is 480 bytes. The buffer sizes defined for MSC links should be large enough to hold at least one complete message. Valid MSC buffer sizes are 1024 bytes to 65536 bytes.

To change the input message destination to any IMS local or remote destination after a message is received but before it is processed, use the TM and MSC Message Routing and Control user exit (DFSMSCE0).

Definitions: Using MSC with APPC/IMS requires you to understand the terminology used for the different MSC systems:

- The *originating system (local)* is the system from which the LU 6.2 program enters the IMS transaction.
- The *remote system* is the system in which the remote transaction executes.
- The *intermediate system* is the IMS that routes messages between the local and remote systems.

At any time, any of these three systems can receive LU 6.2 transactions.

Related concepts:

Chapter 39, "Overview of Multiple Systems Coupling," on page 677

Recovering APPC transactions in an MSC environment

The recoverability of an IMS LU 6.2 transaction depends on whether the message is recoverable, irrecoverable, discardable, or nondiscardable, and when an error occurs.

You can determine the recoverability of APPC messages in an MSC environment. Resource failures affect recovery.

To recover APPC transactions in an MSC environment, analyze the types of failures that can occur. How you handle the error depends on the following:

- The resource that fails: Was it an LU 6.2 session failure, an IMS failure, an application program failure, or an MSC link failure?
- Transaction mode: Was it recoverable or irrecoverable?
- Transaction type: Was it local or remote?
- LU 6.2 conversation mode: Was it asynchronous or synchronous?

You are in control of recovery by the way you define the transaction. The information in the following topics highlights pertinent facts, and then points you to other areas in the IMS library where the subjects are explained in greater depth.

Recoverable versus nonrecoverable transactions

By coding the INQUIRY= keyword on the TRANSACT macro, you tell IMS the recovery status of a transaction. Non-inquiry mode transactions are recoverable; inquiry-mode transactions are not recoverable unless the RECOVER parameter is specified on the TRANSACT macro.

Recoverable transactions are recovered across any IMS failure, shutdown, or restart unless a COLDSTART, COLDSYS, or COLDCOMM restart is performed.

You must define remote transactions with identical recoverability attributes on the local system where the LU 6.2 session originates and on the remote system where the transaction is processed by the application program. You do not need to define the transaction on any intermediate IMS.

Message switches (messages from one LTERM to another) are always recoverable.

Related concepts:

Local APPC transaction discardability versus nondiscardability

The LU 6.2 protocol that you choose for sending a transaction to IMS and the transaction mode (recoverable or irrecoverable) you choose determine if a local APPC transaction is discardable or nondiscardable.

IMS discards a local APPC transaction when it is any of the following:

- A CPI Communications driven application program (without SYNCLVL=SYNCPT specified)
- It is defined as inquiry-only and nonrecoverable
- It is the result of synchronous input from the LU 6.2 application program
- It uses the APPC Sync_Level option NONE

Otherwise, the transaction is nondiscardable. IMS recovers nondiscardable transactions whenever possible; it never recovers discardable transactions.

Transaction processing point of failure

The point of failure in the processing of a transaction also affects its recoverability.

For example, when a local or remote transaction processes and reaches a commit point (sync point), IMS recovers the output response (from the log) even if you have defined the transaction as irrecoverable. Local APPC discardable transactions reach a commit point after IMS sends the output response message to the inputting APPC application program. In this situation, IMS has no output response message to recover or discard if a failure occurs after the commit point. If IMS has queued the transaction on an MSC link, IMS recovers the transaction across link failures.

A message can be either recoverable or irrecoverable, and either discardable or nondiscardable, according to the type of failure that might occur. The descriptions in this topic show you what happens to your transaction when the LU 6.2 session, MSC link, local IMS, intermediate IMS, remote IMS, or application program fail. This information assumes that you can recognize where a failure has occurred and what you need to do to recover.

Recovering transactions after an LU 6.2 session failure

If an LU 6.2 session fails while IMS is receiving an input message, IMS discards the message.

If IMS receives the complete message, processing depends on whether the conversation is:

- CPI-C or not CPI-C
- Synchronous or asynchronous
- Local or remote

CPI-C transaction

If an LU 6.2 session fails while processing a CPI-C transaction, the application program can choose to end the conversation or to continue processing. IMS TM is not involved, and places no restrictions on the choice of committing or backing out updates. The application program makes the decision. Because IMS TM does not

know about the session failure, it takes no action. The normal processing rules for commit and backout apply. IMS does not recover the LU 6.2 conversation.

Related reading: For information on designing your CPI-C LU 6.2 application program, see *IMS Version 11 Application Programming*.

Not a CPI-C transaction

If an LU 6.2 session fails while a local IMS is sending transaction output that is not CPI-C to the LU 6.2 program, and the conversation is synchronous, IMS calls the Message Control/Error exit routine to determine whether to abort and back out, or to continue processing. The default action is to stop the transaction and discard the output message (this is the mode of operation for all protected conversations; that is, conversations allocated using SYNCLVL=SYNCPT). If the conversation is asynchronous, IMS does not call the Message Control/Error exit routine, but queues the output on the message queue to the TP name of DFSASYNC.

Related reading: For information on coding the Message Control/Error exit routine, see *IMS Version 11 Exit Routines*.

Remote APPC transaction

If an LU 6.2 session fails while processing a remote APPC transaction, IMS recovers the output message if it has been enqueued on the local system's MSC link. If the transaction has not at least reached the point of being enqueued on the MSC link, IMS discards it. IMS discards the transaction regardless of the recoverability mode and regardless of whether the LU 6.2 conversation is synchronous or asynchronous. IMS does not call the Message Control/Error exit routine:

- If the transaction is asynchronous, when the output from a remote transaction for a failed LU 6.2 session returns from the remote system to the originating system, IMS sends the response asynchronously to the LU 6.2 application program by using the DFSASYNC TP name.
- If the transaction is synchronous, when the output from a remote transaction for a failed LU 6.2 session returns from the remote system to the originating system, IMS calls the Message Control/Error exit routine to either discard or re-route the transaction output. The default action is to discard the output.

Related reading: For information on using DFSASYNC in your application program, see *IMS Version 11 Application Programming*.

Related concepts:

Chapter 3, "CPI Communications," on page 37

Recovering transactions after an MSC link failure

When an MSC link failure occurs, IMS always recovers all messages, including IMS transactions and responses that are enqueued, about to be enqueued, or being sent across an MSC link. This recoverability is guaranteed, regardless of whether the message is en route to a local, intermediate, or remote MSC system.

The recoverability is not affected by the transaction mode (recoverable or irrecoverable) or the discardable or nondiscardable characteristics of the LU 6.2 protocol used to send the transaction to the local IMS.

Link failures can delay messages from other IMS systems and cause the synchronous LU 6.2 conversation to wait longer than expected for the response.

Related tasks:

 Restarting a logical link (Operations and Automation)

Recovering transactions after a local IMS failure

IMS discards local APPC transactions across a local IMS failure if they meet the discardability criteria. Otherwise, IMS does not discard local APPC transactions, because they are nondiscardable.

If you define your remote APPC transactions as inquiry-type transactions and do not specify RECOVER on the TRANSACT macro definition in the local IMS, IMS does not recover them after a local IMS failure. Otherwise, IMS recovers all recoverable, nonconversational transactions.

After the local IMS sends the transaction message to an intermediate or remote IMS, a local IMS failure has no effect on your transaction's recoverability. The transaction continues to its destination and is processed. When the remote IMS sends the transaction response to the originating IMS after the failure, IMS sends the response to its destination asynchronously through the default transaction program name (TPN) DFSASYNC. The LU 6.2 application programmer needs to plan for this situation.

Related concepts:

“Local APPC transaction discardability versus nondiscardability” on page 63

 Designing an application for APPC (Application Programming)

Recovering transactions after a remote IMS failure

When a remote IMS failure occurs, based on the recoverability attributes of an APPC transaction in a remote IMS, IMS recovers the transaction if it is queued for processing or is being processed in the remote IMS.

Recoverable transactions are recovered; irrecoverable transactions are not. After the transaction reaches a commit point, IMS recovers the output response message regardless of the recovery attributes of the transaction. The discardable and nondiscardable characteristics of the APPC conversation in the originating IMS have no bearing on the transaction's recoverability in the remote IMS across a remote IMS failure.

IMS recovers transactions that are en route to the remote IMS (meaning the transaction message is still en route in the local or intermediate IMS) when the remote IMS fails, regardless of the transaction's recoverability characteristics. After the failure, the remote IMS receives and processes the transactions that were en route at the time of the failure.

Recovering transactions after an intermediate IMS failure

IMS always recovers all messages en route to or from an intermediate IMS that are queued in the intermediate IMS regardless of the recoverability characteristics of the transaction or message.

If the intermediate IMS restarts with either a COLDSTART, COLDCOMM, or COLDSYS, the messages are lost.

Recovering transactions after an application program failure

Transactions sent to IMS from LU 6.2 application programs are processed in the same way as non-LU 6.2 initiated transactions during application program failures.

If an application program fails before reaching a commit point while processing a local or remote transaction from an LU 6.2 device, IMS backs out all messages except those that were inserted to an alternate express PCB and committed with a PURG call. If the failure occurs after the transaction reaches a commit point, IMS recovers everything. If the failing application program's input message was received from another application program (program-to-program switch), this prior application program's processing is still committed (as is true for non-LU 6.2 application programs).

Recoverability flows of LU 6.2 transactions

This topic contains four lists that show synchronous and asynchronous transaction flows, and shows when the transactions are recoverable, irrecoverable, discardable, or nondiscardable.

The following list shows the flow of a transaction sent from an LU 6.2 synchronous conversation to a local IMS.

1. LU 6.2 program: ALLOC LU=*IMS LU name*
2. LU 6.2 program: SEND to local IMS
3. LU 6.2 program: RECEIVE_AND_WAIT
4. Local IMS receives the transaction
5. Transaction is enqueued
6. Transaction executes (if application fails before reaching commit point, message is discarded)
7. Message inserted to I/O PCB
8. Local IMS sends output (Message Control/Error exit routine receives control if LU 6.2 session fails here)
9. Commit point
10. LU 6.2 program: DEALLOCATE

The following list shows the flow of a transaction sent from an LU 6.2 asynchronous conversation to a local IMS:

1. LU 6.2 program: ALLOC LU=*IMS LU name*
2. LU 6.2 program: SEND to local IMS
3. LU 6.2 program: DEALLOCATE
4. Local IMS receives the transaction
5. Transaction is enqueued
6. Transaction executes (if application fails before reaching commit point, message is discarded)
7. Message inserted to I/O PCB
8. Commit point
9. Local IMS: ALLOCATE with TPN=DFSASYNC
10. Local IMS sends output
11. Local IMS: DEALLOCATE

The following list shows the flow of a transaction sent from an LU 6.2 synchronous conversation to a remote IMS:

1. LU 6.2 program: ALLOC LU=*IMS LU name*
2. LU 6.2 program: SEND to local IMS
3. LU 6.2 program: RECEIVE_AND_WAIT

4. Local IMS receives the transaction
5. Transaction is enqueued on remote queue and MSC link (message is recoverable across MSC link failure after enqueue on MSC link)
6. Local IMS sends message across MSC link to remote IMS
7. Remote IMS receives the transaction
8. Transaction executes
9. Output message inserted to I/O PCB
10. Remote IMS enqueues output message to MSC link (message is recoverable across MSC link failure after enqueue on MSC link)
11. Commit point
12. Remote IMS sends output message across MSC link to local IMS
13. Local IMS receives output message
14. Local IMS enqueues output message for LU 6.2 program
15. Local IMS sends output message to LU 6.2 program (Message Control/Error exit routine receives control if LU 6.2 session fails here)
16. LU 6.2 program: DEALLOCATE

The following list shows the flow of a transaction sent from an LU 6.2 asynchronous transaction to a remote IMS:

1. LU 6.2 program: ALLOC LU=*IMS LU name*
2. LU 6.2 program: SEND to local IMS
3. LU 6.2 program: DEALLOCATE
4. Local IMS receives the transaction
5. Transaction is enqueued on remote queue and MSC link (message is recoverable across MSC link failure after enqueue on MSC link)
6. Local IMS sends message across MSC link to remote IMS
7. Remote IMS receives the transaction
8. Remote IMS enqueues the transaction
9. Transaction executes
10. Output message inserted to I/O PCB
11. Remote IMS enqueues output message to MSC link (message is recoverable across MSC link failure after enqueue on MSC link)
12. Commit point
13. Remote IMS sends output message across MSC link to local IMS
14. Local IMS receives output message
15. Local IMS enqueues output message for LU 6.2 program
16. Local IMS: ALLOCATE with TPN=DFSASYN
17. Local IMS sends output message to LU 6.2 program (Message Control/Error exit routine receives control if LU 6.2 session fails here)
18. Local IMS: DEALLOCATE

Planning for XRF and APPC

All LU 6.2 conversation messages that are sent to a failing IMS are broken at the time of an XRF takeover. New conversations must be allocated on the new active IMS after takeover processing completes. XRF support for LU 6.2 application programs is Class 3 service.

The application program uses the USERVAR parameter as the LU name parameter on the ALLOCATE verb. This allows the same LU name to allocate the conversation on the new active IMS. This is the same concept as the LU usage for non-LU 6.2 IMS terminals, but a different USERVAR name is used. At takeover time, the alternate IMS issues a VTAM MODIFY USERVAR command to redirect the USERVAR from the failed active IMS to the alternate IMS. Also, a VTAM VARY TERM command is issued against the failed active IMS to terminate any conversations held during takeover. Then VTAM propagates the USERVAR changes to other VTAM instances as needed.

Any APPC transactions that are enqueued to a remote IMS are handled by the MSC protocols. After takeover in the local IMS, the queued irrecoverable messages are lost. If the same message is queued to a remote IMS, it is recovered when the local IMS takes over. Processing continues and the reply is sent to the active IMS.

VTAM, z/OS, Tivoli® NetView® for z/OS, and any other product requirements should be carefully reviewed. APPC/IMS requires correct operation of these other products.

Transaction retry characteristics

IMS retries certain abend conditions.

Some examples of these conditions that are retried are:

- Deadlock
- Lock reject
- Fast Path retry conditions

These retry conditions still apply to standard DL/I application programs even if they receive their messages from an LU 6.2 application program. If an abend that can be retried occurs, IMS issues an APPC ATBEXAI call to APPC/MVS to determine if any established conversations exist. If a conversation has been allocated, the abend is not retried and the application program is terminated.

If any CPI Communications resource is being used by the application program, the abend condition cannot be retried. Thus, CPI Communications driven application programs and modified IMS application programs that have allocated an LU 6.2 conversation before the abend occurred can never be retried. This prohibition on retry is necessary, because IMS cannot control the state of CPI Communications resources. IMS supports the DL/I INIT STATUS GROUPB call for CPI Communications driven application programs, but not for modified IMS application programs that have allocated an LU 6.2 conversation before the deadlock is detected.

Qualifying network LU names

You can use the same name for LUs on different systems by adding the network identifier. This eliminates the need to have unique names for every LU on every system in your complex. You can use the network-qualified LU name as the name of the partner LU to allocate remote LU 6.2 conversations and sessions.

A network-qualified LU name consists of a one- to eight-character network identifier of the originating system, followed by a period and the LU name. A network-qualified LU name must be enclosed in single quotes for commands.

Example: /DISPLAY LUNAME'NETID1.LUAPPC2'

Use a network-qualified LU name when transmitting data to a remote destination. If no network identifier is present, IMS allows z/OS to determine the destination.

The parameter ALL for either the network identifier or the LU name cannot be substituted in a network-qualified LU name in commands.

Example: /DISPLAY LUNAME'NETID1.ALL'

The LU name in the LU 6.2 descriptor can be network qualified. The network-qualified LU name is optional on commands that support the LUNAME keyword.

APPC/MVS uses the name of LU 6.2 network-qualified LUs to allocate conversations. APPC/MVS strips off the network ID and passes the 8-byte LU name to VTAM. Without APPC/MVS support for network-qualified names, the LU name must be unique for the different networks.

The ALL parameter for either the network identifier or the LU name cannot be substituted in a network-qualified LU name in a command.

The LU name in the LU 6.2 descriptor can be network qualified.

Related reading: For information on using network-qualified names in commands, see *IMS Version 11 Commands, Volume 1: IMS Commands A-M*.

Managing multiple LUs for a single IMS system

An IMS system can be represented on a network by more than one LU name. If more than one LU name exists for a single IMS system, you might need to specify which LU should process asynchronous outbound messages.

When more than one LU is defined for an IMS system, one of the LUs serves as the default, or *base*, LU. Normally, the base LU handles all allocation requests for outbound conversations. However, in some cases the remote partner LU might expect the outbound conversation to come from an LU other than the base LU, you can control which LU handles outbound conversation allocation requests. There are three methods for controlling which LU handles outbound conversation allocation requests:

- By specifying an LU name in the OUTBND parameter of the LU 6.2 descriptor in the DFS62DTx member of the IMS.PROCLIB data set. IMS routes outbound messages sent to the LU 6.2 descriptor by using the local LU name specified in the OUTBND parameter.

Use the /CHANGE DESC command to change the local LU for a LU 6.2 descriptor.

Use the /DISPLAY DESC command to display the local LU and other descriptor specifications.

- By specifying an LU name in the OUTBND parameter of the DFSDCxx startup procedure. When an LU name is specified in the OUTBND parameter of the DFSDCxx startup procedure, the specified LU serves as the default LU for all asynchronous outbound conversation messages, regardless of which LU received the original inbound conversation.

- By specifying the local LU option APPCLLU=Y in the DFSDCxx startup procedure. When the local LU option is specified, IMS routes asynchronous outbound conversations through the LU that received the original inbound conversation.
- By overriding any LU name specifications in the LU 6.2 Edit exit routine (DFSLUEE0).

Reassigning an LU to another IMS system

You can reassign an LU from one IMS system to another by using MVS commands.

To reassign an LU to another IMS system:

1. Delete the LU from its current IMS system by issuing the MVS command SETAPPC LUDEL.
2. Redefine the LU on the new IMS system by issuing the MVS command SETAPPC LUADD,ACBNAME=*luname*,SCHED=*new_IMS*,NQN.

DFSAPPC system service

DFSAPPC is an IMS system service for exchanging messages between LU 6.2 application programs (LU 6.2 to LU 6.2), and between LU 6.2 application programs and IMS-managed LTERMs. Message delivery is asynchronous; messages are held on the IMS message queue until they are delivered.

LU 6.2 application program can use DFSAPPC to send messages to IMS-managed LTERMs. Use the LTERM option of the DFSAPPC service to select this capability.


Message switching

Message switching is part of the implicit APPC interface and allows IMS terminals and LU 6.2 application programs to exchange messages. Messages routed to an LU 6.2 application program initiate LU 6.2 application programs.

When using DFSAPPC, the remote device can choose to route a message using either the LTERM or LU 6.2 TPN option. Messages sent with the LTERM option are directed to IMS-managed local or remote LTERMs.³ Messages sent without the LTERM option are sent to the specified LU 6.2 application program.

The message format for DFSAPPC is shown in the following figure.

3. If the LTERM is associated with an LU 6.2 destination, the message is sent as if an LU 6.2 application program had been explicitly selected.



Notes:

- 1 A mandatory blank is required between DFSAPPC and the options.
- 2 Use blanks anywhere within the DFSAPPC options except within keywords or values. Use commas as delimiters between keyword-parameter sets along with or in place of blanks. However, because the TP name character set allows commas, at least one blank must be used to terminate the TPN value.
- 3 You can specify either the LTERM= or the TPN= option, but not both. Only use the other keyword options when you specify the TPN= option.
- 4 Use the IMS default values for the DFSAPPC options to establish an LU 6.2 conversation with a partner program when the values are not provided by another source. If the DFSAPPC service is coded without specifying any options, use IMS default LU 6.2 conversation characteristics.

Figure 11. DFSAPPC message format

The DFSAPPC option keywords are defined as follows in order of occurrence (except for the keywords following TPN=, which are listed alphabetically):

LTERM=

The 1- to 8-character LTERM option is the name of an IMS LTERM. Messages sent with the LTERM option are directed to an IMS-managed local or remote LTERM. If the LTERM is associated with the LU 6.2 descriptor, it is treated as if an LU 6.2 application program has been explicitly selected. LTERM names can contain uppercase alphabetic, numeric, and national characters ('@', '\$', '#'). When LTERM is specified, other keywords cannot be specified.

TPN=

The 1- to 64-character TPN option is the partner's transaction program name used with the logical unit name to establish an LU 6.2 conversation with a partner program. Because the TP name character set allows commas, at least one blank must be used to terminate the TPN value.

TP names can contain any character from the 00640 character set except a blank. The 00640 character set, documented in the *CPI Communications Specification* includes uppercase and lowercase letters A through Z, numerals 0-9, and 20 special characters.

When the TPN and SIDE options are specified, the TPN name overrides the TP name contained in the side information entry.

Although DFSAPPC allows the use of the 00640 character set, IMS commands do not use this character set. IMS commands can only operate on TPNs that use uppercase alphabetic, numeric, and national characters ('@', '\$', '#'). IMS commands cannot operate on extended TPNs.

LU=

The 1- to 17-character LU option is the partner's logical unit name used with the transaction program name (TPN) to establish an LU 6.2 conversation with a partner program.

LU names can contain any character from the APPC/MVS Type A character set. LU names can contain uppercase alphabetic, numeric, and national characters ('@', '\$', '#'), and must begin with an alphabetic or national character. You can also use a 17-byte network-qualified LU name in the LU field.

When the LU and SIDE options are specified, the LU name overrides the LU name contained in the side information entry.

MODE=

The 1- to 8-character MODE option is the partner's mode name used with the logical unit name and transaction program name to establish an LU 6.2 conversation with a partner program.

MODE names can contain any character from the APPC/MVS Type A character set. MODE names can contain uppercase alphabetic, numeric, and national characters ('@', '\$', '#'), and must begin with an alphabetic or national character.

When the MODE= and SIDE options are specified, the MODE name overrides the mode name contained in the side information entry.

SIDE=

The 1- to 8-character SIDE option is the side information entry name used to establish an LU 6.2 conversation with a partner program.

SIDE names can contain any character from the 01134 character set. The 01134 character set, documented in the *CPI Communications Specification* includes uppercase alphabetic characters A through Z and numeric characters 0-9.

When the SIDE option is specified, the LU, TPN, and MODE options can also be specified to override the values in the side information entry.

SYNC=

The SYNC option allows the application to override the LU 6.2 conversation sync level default provided by IMS.

SYNC=N

Sync_Level is NONE.

SYNC=C

Sync_Level is CONFIRM.

TYPE=

The TYPE option allows the application to override the LU 6.2 conversation type default provided by IMS.

TYPE=B

Conversation type is BASIC.

TYPE=M

Conversation type is MAPPED.

If an error is found while processing the options list, error message DFS1957 DFSAPPC ERROR is sent to the terminal.

Related reading: For more information on Type A character sets, see *z/OS MVS Programming: Writing Transaction Programs for APPC/MVS*.

Asynchronous output delivery

When creating a message destined for an LU 6.2 application program, IMS establishes conversation characteristics.

These characteristics are extracted from the:

- LU 6.2 descriptor
- DL/I Change call option list
- DFSAPPC message switch options

If IMS cannot extract a particular conversation characteristic from this list, IMS uses the defaults that are shown in the following table. If a side table name is extracted, the default mode name is not used. IMS assumes that side table entries contain a mode table entry name. If an /ALLOCATE command for a particular LUNAME - TPNAME destination specifies a mode table entry name, that entry name overrides the mode table name specified for the message.

IMS uses the information in the following table to initiate a conversation with the LU 6.2 application program that is associated with the alternate PCB. Certain fields, such as LU name, are application specific. Default values are provided but can be overridden by parameters associated with the message. A default value is used by IMS only if no value is provided by any other source. The application program can modify the default conversation characteristics using an expanded interface to the DL/I CHNG call.

Table 1. APPC/IMS default conversation characteristics

Characteristics	Default value
Conversation_Type	Mapped
Deallocate_Type	Deallocate_Sync_Level
Error_Direction	Receive_Error
Fill	Fill_LL
Log_Data	Null
Log_Data_Length	0
Mode_Name	'DFSMODE'
	IMS uses the same mode name provided by the inputting LU 6.2 partner to allocate the outbound conversation. That is, whatever mode name the inputting conversation uses, IMS also uses it for outbound allocates.
Mode_Name_Length	7
Partner_LU_Name	'DFSLU'
Partner_LU_Name_Length	5
Prepare_to_receive_type	Prep_To_Receive_Sync_Level
Receive_type	Receive_and_Wait
Return_Control	When_Session_Allocated
Send_Type	Buffer_Data
Sync_Level	Confirm
TP_Name	'DFSASYNC'
TP_Name_Length	8

APPC transaction security

The security options for APPC/IMS and LU 6.2 application programs are quite extensive. The partner systems can range from a single-user terminal or workstation to a multi-user system. All systems can have their own complex security environment. Security for IMS can be simple or complex.

Every transaction program name (TPN) must pass a security check before it is executed. The user ID that initiates the transaction is identified on the LU 6.2 format header (FMH5). If no user ID exists because you specify SECURITY=NONE, you can only access resources that are not defined with UACC (NONE). Any TPNs that are accessible in all circumstances should not be defined with UACC (NONE). The TPN security definition is required.

z/OS security consists of two parts. First, z/OS authenticates the transaction user. The LU 6.2 transaction contains security information. The FMH5 contains the user ID, a “profile” name, which is used as the group name, and security options. You supply both the user ID and password. The user ID is defined to RACF, and the password must be valid for the user ID.

If Already_Verified is specified in the FMH5, the sending system verifies the user ID. This user ID must be defined to RACF on the receiving z/OS system. No password is needed in this case.

If SECURITY=NONE is specified, z/OS does no checking. Instead, z/OS builds a special security profile that corresponds to SECURITY=NONE. This allows access to z/OS and APPC/IMS resources that have UACC specified at any level other than NONE. Resources with UACC (NONE) or without a UACC specified cannot be accessed.

After the user ID is established, z/OS verifies that the user ID can execute the specific transaction. z/OS verifies that the user ID's access profile has ACCESS (EXECUTE) for the entity *dbtoken.x.tpname* in the CLASS (APPCTP). The value of *dbtoken* is the dbtoken value specified in the TP_Profile data set. Based on the APPCTPxx parameter specified for this LU, the value of *x* is either the *user ID*, *group* or *SYS1*.

If either of these security checks fails, z/OS rejects the transaction, and IMS is not informed of it. z/OS can also check:

- Session-level security (RACF resource class APPCLU)
- Port of entry (RACF resource class APPCPORT)
- Local application (RACF resource class APPL)

The IMS administrator should verify that these security checks are successful.

Related reading: For more information on coding the RACF resource classes, see *z/OS MVS Planning: APPC/MVS Management*.

The security check in IMS is based on the IMS transaction code or executed command name. If the TPN is DFSAPPC, no additional security check occurs. If RACF is used on your system, z/OS rejects the transaction if RACF is not active. The IMS command name or transaction code associated with the TPN is used in the RACF resource class associated with this IMS ('C' or 'T'). RACF checks IMS commands and transactions for all other IMS terminal types in the same way.

If the RACF check is successful, the Transaction Authorization exit routine (DFSCTRN0) is called for transactions, and DFSCCMD0 is called for command authorization. However, the following rules apply to RACF:

- For commands, default security only applies if RACF is not used.
- For remote transactions, RACF is optional.

Otherwise, the exit routines make the security decision.

The intended environment executes APPC/IMS with RACF security active. It is possible to run with RACF not active in the APPC/IMS system, but it is not possible to run with RACF not active in the z/OS system. In this sense, RACF is mandatory for LU 6.2.

The complexity of the security environment is derived partly from the many resources involved (VTAM, z/OS, and IMS) and the granularity of protection that is possible. The security definitions must be closely coordinated for successful operation of the application system.

Related concepts:

 [IMS security \(System Administration\)](#)

Part 3. Extended Terminal Option (ETO)

These topics introduce the extended terminal option (ETO) and include an overview of ETO and the information you need to administer ETO terminals in an IMS TM Network.

Chapter 5. Overview of the Extended Terminal Option

The Extended Terminal Option (ETO) of IMS allows you to add VTAM terminals and users to your IMS without predefining them during system definition.

ETO is part of the IMS Transaction Manager (TM), and provides additional features for users, such as output security, automatic logoff, and automatic signoff.

This topic provides system programmers with the conceptual information that is required to implement and administer ETO. Read the information in this topic if you are unfamiliar with ETO.

Benefits of using ETO

ETO adds essentially two major enhancements to the Transaction Manager environment. With ETO:

- Users can obtain IMS sessions with VTAM terminals that have not been defined to IMS during system definition.
- Output messages that are destined for particular users are secure, and they reach only those users.

In addition, by installing ETO, you can achieve each of the following:

- Improved system availability by reducing scheduled down time associated with adding or deleting VTAM terminals.
- Faster system availability to users, because they can establish an IMS session from any VTAM terminal in the network.
- Improved IMS security by relating output messages to users, rather than to terminals.
- Reduced number of macros required to define the terminal network. This reduces system definition time and storage requirements.
- Reduced checkpoint and restart time. For ETO terminals and user structures, resources are not allocated until they are actually required; similarly, when they are no longer required, they are deleted.
- Reduced number of skilled system programmer resources that are required for maintaining static terminal definitions.

Related tasks:

Chapter 6, “Administering the Extended Terminal Option,” on page 87

ETO terminology

Certain terms have meanings that are specific to ETO and that are therefore important for understanding and administering ETO.

Terminals

The definitions for terminal, static terminal, and dynamic terminal are described in this topic.

Definitions:

- A *terminal* is a physical VTAM logical unit (LU) that establishes a session with IMS. A physical terminal is represented using a control block.
- When terminals are not built by ETO but are defined at system definition, they are called *static terminals*. When messages are sent to a static terminal they are queued to a logical terminal (LTERM) message queue, where they await retrieval by the recipient.
- When a terminal is not defined at system definition and ETO builds a terminal, that terminal is called a *dynamic terminal*, or an *ETO terminal*. For dynamic terminals, the logical terminal (LTERM) is known as a *dynamic user message queue*, LTERM associates the messages with the user, rather than with the physical terminal. Associating messages with the users provides more security for these users, because they can access their messages only when they sign on using their unique user ID. In addition, all users in the network can access their messages from any physical terminal, instead of being restricted to using a particular physical terminal.

Dynamic users

Definition: An ETO *dynamic user* is a user who signs on to a dynamic terminal and who has a unique identification (user ID) that IMS uses for delivering messages. The user is usually associated with a person but can also be associated with another entity, such as a printer.

Terminal structures

A *terminal structure* is an IMS control block that represents a specific terminal that is known to IMS. A terminal structure is created when the individual terminal logs on to IMS. It is deleted when the terminal logs off with no remaining associated activity (such as status that must be retained for the next connection to IMS).

User structures

A *user structure* is a set of IMS control blocks, including a user block and one or more LTERM blocks. The message queues are associated with the dynamic user, as opposed to the physical terminal, and they are queued to the user ID.

The dynamic user structure connects to the physical terminal only when the user signs on. This provides a secure environment, because different users accessing the same terminal cannot receive each other's messages.

IMS creates a user structure when either of the following events take place:

- A dynamic user signs on to IMS.
- Output messages that are destined for a dynamic user are sent to the user, but the user has not signed on to IMS.

Usually, a user structure represents a person who uses IMS. The user structure name is usually the same as the user ID. A user structure can also represent a logical destination, such as a printer. In this case, the user structure name can be the same as or different from the LTERM name that your installation uses in its application programs and its exit routines. For example, you can assign the same name to a user structure for a printer that you assign to its LTERM destination node name. However, output is then queued according to the terminal, and not to the user.

The following figures show the differences between static resources and ETO dynamic resources.




Figure 12. Static resources




Figure 13. ETO dynamic resources

Descriptors

A *descriptor* provides information to IMS when IMS builds a dynamic resource for a logon or a signon.

The four types of ETO descriptors are:

- Logon descriptors
- User descriptors

MSC descriptors
MFS device descriptors

IMS stores descriptors in the following IMS.PROCLIB members:

DFSDSCMx

Contains the descriptors that are automatically generated during IMS system definition. The suffix of DFSDSCMx matches the suffix that your installation specifies on the SUFFIX= parameter of the IMSGEN system definition macro.

DFSDSCTy

Contains customized device descriptors that your installation creates. Descriptors in DFSDSCTy override duplicate descriptors in DFSDSCMx, and the last descriptor that is defined is used.

Logon descriptor

Definition: A *logon descriptor* is a skeleton that IMS uses to build an ETO dynamic terminal. It provides information regarding a terminal's physical characteristics. IMS uses logon descriptors in conjunction with exit routines to create terminal structures.

The three types of logon descriptors are: generic, group, and specific.

Definitions:

- A *generic logon descriptor* provides characteristics for all terminals of a particular type. For example, all SCS printers might share a single generic descriptor. Similarly, all 3270 terminals might share a generic descriptor.
- A *group logon descriptor* provides characteristics for a collection of terminals, each of which has compatible hardware characteristics and is defined to IMS in the same manner. The actual characteristics for these terminals are usually identical, but they can differ; IMS uses the group descriptor to derive their characteristics.

Example: You might create separate logon descriptors for different groups of terminals that differ only in the setting for the autologoff (ALOT) time value.

- A *specific logon descriptor* provides characteristics for a single terminal, and these characteristics apply only to that terminal. In this case, the descriptor name matches the name of the terminal that it describes.

Recommendation: Although you might need to use specific logon descriptors during the actual migration to ETO, use generic or group logon descriptors after you have migrated to ETO; these kinds of descriptors ease network administration.

User descriptor

Definition: A *user descriptor* is a skeleton from which a user structure is built. A user descriptor can provide user options and queue names.

MSC descriptor

Definition: An *MSC descriptor* is used to create a remote LTERM, which is an LTERM that does not exist on the local IMS. The physical terminal definition (either static or dynamic) for the remote LTERM is in the remote IMS.

Each MSC descriptor for a remote LTERM is loaded during IMS initialization and tells IMS which MSC link to use for output destined for that remote LTERM.

MFS device descriptor

Definition: *MFS device descriptors* allow you to add new device characteristics for MFS formatting without requiring an IMS system definition. The MFSDCT utility (DFSUTB00) uses MFS device descriptors to update default formats in the MFS library.

IMS also uses MFS device descriptors to update the MFS device characteristics table. IMS loads this table only during initialization; therefore, updates are not effective until the next IMS initialization.

Related concepts:

Chapter 39, “Overview of Multiple Systems Coupling,” on page 677

ETO concepts

The main purpose of ETO is to dynamically define terminals to IMS. This topic describes such things as ETO terminal and user structures, descriptors and exit routines for dynamic terminals.

When structures are created and deleted

Structures are created in the following situations:

- Log on
- Sign on
- Output is queued to your LTERM
- /ASSIGN command is used to assign an LTERM to a non-existent user
- /ASSIGN command is used to assign a non-existent LTERM to a user
- /CHANGE USER *username* AUTOLOGON command is directed to a non-existent user

In all cases, IMS searches for an existing structure (terminal or user) before creating a new one.

IMS creates and deletes user structures in the following sequence (This sequence applies only to terminal logon and logoff and to user signon and signoff. When asynchronous output is queued to a user, IMS creates the user structure, as needed.):

1. When you establish a session between IMS and an undefined terminal, IMS selects a logon descriptor.
2. Using the information in the logon descriptor, the customization defaults, and VTAM information, IMS builds a VTAM terminal control block that describes the new terminal.
3. When you sign on, if a user structure does not exist, IMS builds one, using information from a user descriptor that it selects, and then connects this user structure to the terminal structure.
4. IMS deletes terminal or user structures when they are no longer needed to maintain sessions. User structures are typically deleted when you sign off, if no special status needs to be maintained and if no messages remain queued. IMS deletes terminal structures when no terminal status exists (such as trace mode), no user is signed on, and the terminal is idle.

If you are using Resource Manager and a resource structure, IMS normally maintains status in the resource structure instead of the local control blocks. Therefore, IMS deletes the structures.

Exceptions: The following terminal structures and user structures are not deleted:

- SLU P and Finance terminal and user structures are normally only deleted during an IMS cold start if SRM=LOCAL. They can also be deleted, however, if the /CHANGE NODE COLDSESS command is used, in which case they are deleted at the first checkpoint following the command.
- ISC terminal and user structures are only deleted following a cold session termination if SRM=LOCAL.

Descriptors and exit routines

Using descriptors and exit routines, you can assign characteristics to ETO dynamic terminals and assign user structures to be associated with those terminals.

A descriptor provides the basic information for the dynamic terminal. An exit routine completes or changes this information. Two methods of using descriptors and exit routines are:

- You can use many descriptors and code little or no processing logic in exit routines.
- You can use few descriptors and code exit routines to perform much of the processing.

How descriptors are created and used

All descriptors are created during IMS initialization, prior to IMS startup. You must specify that you want ETO support and ensure that the ETO initialization exit routine (DFSINTX0) does not disable ETO support.

During IMS initialization, IMS reads and validates all ETO descriptors. IMS initialization then continues, and the descriptors remain in storage for the duration of IMS execution. Any changes you make to descriptors become effective after the next initialization of IMS.

IMS uses descriptors to create both terminal and user structures. IMS rebuilds structures during an IMS restart, if appropriate. For example, if messages are queued for a structure and IMS goes down, the structures are rebuilt when IMS restarts. IMS rebuilds these structures to be the same as they were before the IMS restart. IMS does not use the descriptors or exit routines to rebuild these structures. Therefore, any changes you make to descriptors are only reflected in new structures that are built after IMS restart, and the changes are not reflected in structures that are rebuilt during IMS restart.

Example: USERA signs on using descriptor DESCA which specifies ASOT=20. USERA starts an IMS conversation, and then IMS abnormally terminates. The system programmer changes DESCA to ASOT=10. After the IMS restart, USERB signs on using DESCA. USERA was rebuilt during the IMS restart. USERA still has ASOT=20, and USERB has ASOT=10.

Summary of ETO implementation

The following figure illustrates the ETO concepts and shows an overall view of an ETO implementation.

- A physical-to-logical relationship does not exist between the device and the LTERM.
- **8** RACF (or an equivalent product) manages security in the ETO environment.

Chapter 6. Administering the Extended Terminal Option

The IMS Extended Terminal Option (ETO) allows you to dynamically add VTAM terminals and users to your IMS without having to first define them during system definition.

ETO dynamically creates user structures for a terminal session when:

- The user signs on to IMS.
- Output messages are sent to the user and await retrieval by the user.
- The /ASSIGN command is used to assign an LTERM to a non-existent user.
- The /ASSIGN command is used to assign a non-existent LTERM to a user.
- The /CHANGE USER *username* AUTOLOGON command is directed to a non-existent user.

Some of the administrative advantages of using ETO include:

- You do not need to code the following macros for the system definition stage-1 input stream:

MSC remote NAME macros

VTAM macros: TYPE, TERMINAL, NAME, VTAMPOOL, SUBPOOL

Removing these macros reduces the complexity of network management.

- You need to perform fewer system definitions.
- You schedule fewer planned outages for new system definitions.

Using ETO, you can ensure that all terminals and users are able to establish sessions with IMS, even if these terminals and users are not defined to IMS during system definition.

You can use execution-time parameters and exit routines to authorize users to access some or all of the functions that ETO provides.

Related reading:

- For information on specific device types that IMS supports, see *IMS Version 11 Release Planning*.
- For information on using execution parameters, see *IMS Version 11 System Definition*.
- For information on using exit routines, see *IMS Version 11 Exit Routines*.

Planning for ETO

Migrating a static-terminal environment to an ETO environment requires planning.

Although you can continue to define VTAM terminals and LTERMs to IMS during system definition, if you do so:

- You cannot take advantage of the ETO features that exist for those terminals.
- You must fully define the terminal. You must supply all TERMINAL macros, NAME macros, and parameters.

ETO terminals must be VTAM terminals.

Restrictions: The following VTAM terminals cannot be ETO terminals:

- IMS master terminal (MTO)
- IMS secondary master terminal
- MSC physical and logical links
- ISC sessions that are used by XRF for surveillance
- LU 6.2 terminals (dynamically created and managed by APPC/IMS)

Identifying your requirements

The degree to which you implement ETO across your IMS installation depends on your installation requirements.

ETO is considered fully implemented when no static VTAM terminals exist in the system and when the majority of terminals and users are defined using the default logon descriptor and default user descriptor. However, because installations vary in application program dependencies, the cost of fully implementing ETO also varies.

Your installation should determine the extent to which full ETO support is required, based on the following requirements:

Full user-message security

Full implementation of ETO is required for full user-message security. In this environment, no node-name user descriptors exist. Any requirements for user structures that the default user structure does not provide must be defined by user descriptors or by the Signon exit routine (DFSSGNX0).

Dynamic terminal support only

Only partial implementation is required for dynamic terminal support. You can move network definition statements from the system definition to ETO descriptor PROCLIB members. Benefits of this implementation include:

- Fewer system definitions are required in order to maintain network definitions, because you can change descriptors between IMS warm starts.
- Shorter run times are required for system definition, because you do not need to define VTAM terminal networks.
- Improved performance exists for IMS checkpoint and restart, because dynamic terminals and user resources are allocated only when they are used.

With partial implementations, however, you do not achieve improved user-message security, because each LTERM has a fixed relationship with a physical terminal.

ETO restrictions

Before implementing ETO, ETO has a few restrictions that you should be aware of.

The restrictions for ETO include:

- Dynamic terminals are not supported for terminal-related MSDBs or for non-terminal-related MSDBs that have LTERM keys.
- Application programs that use specific LTERM names sometimes require particular ETO customization.

The DFSUSER user descriptors can help you customize ETO for application programs that have dependencies on LTERM names contained in the I/O PCBs.

Related concepts:

“Using DFSUSER user descriptors” on page 106

Defining physical terminals

When implementing ETO, to achieve your desired VTAM terminal network you need to be aware of certain requirements and aspects of how ETO and VTAM work together as you plan for and define the physical terminals in the network.

Performing the following actions can ensure that you achieve your desired VTAM terminal network:

- Assess how often IMS application programs depend on specific terminal characteristics.
- Check the accuracy of each VTAM terminal definition. For each dynamic terminal, ETO builds a terminal structure that relies on the VTAM definition for the characteristics (such as the LU type, screen size, and model) for that terminal.

Terminal characteristics that are specified in your IMS system definition might differ from (and override) those in the VTAM definitions. If these terminal characteristics in the IMS system definition are compatible with those of the actual terminal, the discrepancy is not apparent.⁴

- Either provide specific node-name logon descriptors or use the Logon exit routine (DFSLGNX0) for terminals that are not adequately defined using the default logon descriptor.
- Code one of the following two exit routines to determine the device type:
 - The Logon exit routine (DFSLGNX0) can determine the device type by examining the default logon descriptor.
 - The Signon exit routine (DFSSGNX0) can determine the device type later in the process by examining the terminal control blocks.

When receiving a request to establish a terminal session, IMS relies on the following information from VTAM session parameters:

- **UNITYPE**

The unit type of the node that is attempting to log on. IMS determines the UNITYPE by using the following fields:

- The LUTYPE field of the CINIT⁵ request provides part of the UNITYPE information. The LUTYPE value is usually found in the first byte of the PSERVIC operand of the MODEENT macro, which is used to generate the VTAM mode table entry that is used for a logon. The following figure shows the fields of the PSERVIC operand in the VTAM MODEENT macro.

4. If the actual terminal characteristics do not match those in the IMS definition or the VTAM definition, it is possible that the terminal can function with IMS.

5. *CINIT* is a network services request sent from a system services control point (SSCP) to a logical unit (LU), asking that LU to establish a session with another LU and to act as the primary end of the session.




Figure 15. VTAM MODEENT macro PSERVIC operand fields that IMS uses

The following table shows the mapping of the LUTYPE value to the IMS UNITYPE.

Table 2. Mapping for VTAM LUTYPE value to IMS UNITYPE

VTAM LUTYPE	IMS UNITYPE
LUTYPE X'06'	LUTYPE6
LUTYPE X'02'	SLUTYPE2
LUTYPE X'01'	SLUTYPE1 (default) or NTO
LUTYPE X'00'	3270, SLUTYPEP, or 3600/Finance

- If the LUTYPE field is X'00' (indicating that the terminal is 3270 non-SNA, Finance, or SLU P), IMS must check the transmission services profile specification in the TS field of the CINIT request. The TS value is usually found in the TSPROF operand of the MODEENT macro that is used to generate the VTAM mode table. This LUTYPE value must match the value in the logon descriptor that IMS selects for a logon.

The following table shows the mapping of the TSPROF specification to the IMS UNITYPE.

Table 3. Mapping for TSPROF specification to IMS UNITYPE

TSPROF specification	IMS UNITYPE
X'02' or X'03'	3270
X'04'	SLUTYPEP (default) or 3600/Finance

- **Input RU size**

The input RU (request unit) size in the BIND must be less than or equal to the RECANY buffer size for the IMS (also required for static terminals).

- **Output RU size**

The output RU size in the BIND must be greater than or equal to the OUTBUF size that IMS determines for the terminal from the selected logon descriptor. This parameter is also required for static terminals.

- **Screen size and model number**

For non-SNA 3270 and SLUTYPE2 devices, IMS retrieves both screen size (row and column) and model number from the BIND:

- For static terminals, the screen size is the value that was specified in the system definition.

- For dynamic terminals, IMS determines the screen size from the VTAM definition or from the Logon exit routine (DFSLGNX0).

Recommendation: Until the ETO feature was available, IMS ignored the screen size and model number values in the BIND, because the IMS system definition held this information. Therefore, check to ensure this definition is accurate.

If you determine that a VTAM definition is inaccurate, you can use the Logon exit routine (DFSLGNX0) to override the VTAM-provided screen size and model number. For example, use a terminal naming convention or MODETAB definition convention. The Logon exit routine can also assign USER=NODE as a name, when appropriate.

IMS uses the values in the PSERVIC operand of the MODEENT macro, which is used to generate the VTAM mode table entry that is used for a logon. MFS formats must be available for all screen sizes that IMS dynamically builds.

Restriction: IMS does not use the 3270 Read Partition Query (RPQ) command to determine the screen size from the device controller.

For more information on the Logon exit routine (DFSLGNX0), see *IMS Version 11 Exit Routines*.

Related tasks:

“Identifying VTAM device types, screen sizes, and models” on page 92

Planning for both static and dynamic terminals

Static and dynamic terminals can coexist in the same IMS. However, if users move between static and dynamic terminals, there are situations you need to plan for.

If users move between static and dynamic terminals, plan for the following situations:

- IMS maintains separate queues for static and dynamic terminals. A static terminal has one or more LTERMs associated with it, as controlled by the IMS system definition (or the /ASSIGN command to move an LTERM to a different terminal). Some users become accustomed to having their output queue follow them from ETO terminal to ETO terminal. Static terminals do not provide this feature. Users need to be able to differentiate between static and dynamic terminals, or confusion can result.
- Given one static terminal and one dynamic terminal, separate IMS conversations can exist at the same time. The dynamic terminal conversation belongs with the user structure and follows the user from terminal to terminal.⁶ The static conversation belongs to the static terminal and can only be released to a static terminal. This situation is user friendly and predictable only when the user is certain of the terminal type (static or dynamic).

Normally, the terminal operator is able to determine whether a terminal is static or dynamic by checking the security information that is provided at the end of the DFS3650 (SESSION STATUS) message:

OUTPUT SECURITY AVAILABLE

Indicates that the terminal is dynamic, and output is associated with the signon ID.

NO OUTPUT SECURITY AVAILABLE

Indicates that the terminal is either statically defined or that ETO created it by using one of two methods:

6. This assumes that the Signon exit routine (DFSSGNX0) sets the same user name each time, as is usually the case.

- Using a node user descriptor
- Using the Signon exit routine to assign the node name to the user structure

In either case, the output is associated with the terminal, rather than with the user.

Exception: Message DFS3650 might be suppressed if you use the NOTERM option.

Recommendation: To ease migration and limit possible confusion, convert to dynamic ETO terminals by using logical groupings within your organization, such as departments or floors.

Defining terminals for growth

When designing your ETO implementation, be sure to plan for growth in your network.

Recommendations:

- To increase future growth potential and system availability, minimize the use of descriptors. The Logon, Signon, and Destination Creation exit routines should provide enough customization, thereby eliminating the need for unique descriptors beyond those that are required for specific terminal types.
- To ensure user data is correctly handled, design exit routines carefully. In particular, carefully plan for user data that is specified during the logon process. Exit routines should work correctly, regardless of whether user data is specified.

Some terminal types, such as Finance and SLU P terminals, require user data that specifies signon information. If this data is missing, the equivalent information must be provided in the Logon exit routine (DFSLGNX0).

Related tasks:

“Identifying VTAM device types, screen sizes, and models”

Related reference:

 Transaction Manager exit routines (Exit Routines)

Identifying VTAM device types, screen sizes, and models

VTAM logon CINIT user data provides IMS with information to build session control blocks. This information includes logon descriptors, screen size, model numbers, and RU sizes.

The information must be accurate to ensure that the correct terminal-related control blocks is built from information in the logon descriptors. Carefully define your VTAM PSERVIC parameters to ensure that IMS selects the appropriate logon descriptors and establishes screen sizes using specific terminal characteristics.

The BIND is rejected if the input and output RU sizes in the BIND are incompatible with the IMS RECANY and descriptor OUTBUF sizes.

Defining device types

If the definitions are coded incorrectly, IMS chooses the wrong MFS format and device characteristics, possibly causing screen format errors.

IMS dynamic terminal control blocks are built from definitions in a combination of the following:

- The logon descriptor

- Information that VTAM passes to IMS during logon
- The MFS device characteristics table

VTAM passes the physical terminal characteristics to IMS. The following fields have information that IMS uses to determine device types:

- LUTYPE field from the VTAM PSERVIC parameter of the VTAM MODEENT macro
- TS profile from the TSPROF parameter of the VTAM MODEENT macro

Non-SNA 3270 printers and displays

The default descriptor names for non-SNA 3270 printers and displays are DFS327P and DFS3270.

Rules required for defining non-SNA 3270 devices are:

- The TSPROF parameter of the VTAM MODEENT macro must be 2 or 3.
- The PSERVIC parameter of the VTAM MODEENT macro must specify LU Type=0.
- Byte 12 in the VTAM PSERVIC parameter must be modified so that it distinguishes a printer from a video. The following table shows the content that specifies each device type based on the bit location in byte 12:

Table 4. Bits in byte 12 of the VTAM PSERVIC parameter

Bits in byte 12	Content								
0-1	Device type: <table> <tr> <td>00</td><td>Unspecified device type</td></tr> <tr> <td>01</td><td>Printer device</td></tr> <tr> <td>10</td><td>Display device</td></tr> <tr> <td>11</td><td>Display/printer device (3275)</td></tr> </table>	00	Unspecified device type	01	Printer device	10	Display device	11	Display/printer device (3275)
00	Unspecified device type								
01	Printer device								
10	Display device								
11	Display/printer device (3275)								
2-7	Reserved.								

If you specify nothing (B'00'), the default is B'10', which is a display device. ⁷

Related reading: For more information on defining non-SNA 3270 devices, see *z/OS Communications Server: SNA Programming*.

LU type-2 devices

The following is a list of rules required for defining SLU-2 devices:

- The TSPROF parameter of the VTAM MODEENT macro must be 3.
- The PSERVIC parameter of the VTAM MODEENT macro must specify LU Type=2.

Related reading: For more information on defining LU type-2 devices, see *z/OS Communications Server: SNA Programming*.

7. B'00' means bits 0 and 1 equal 00. B'10' means bits 0 and 1 equal 10.

3275 devices

VTAM definitions alone cannot identify a 3275 device as it is logging on. Additional information (UNIT=3275) must be specified in a logon descriptor. Define this descriptor in the CINIT user data or in the Logon exit routine. The descriptor itself must be identified as a 3275 device type.

Static 3275 devices are defined in the IMS system definition as follows:

```
TYPE          UNITYPE=3270  TERMINAL UNIT=3275,TYPE=3270-An,SIZE=(24x80),COMPT=PRT1
```

The 3275 has only one buffer. This forces the display and printer components to be the same model. The VTAM definitions for a dynamic 3275 as statically defined are:

```
PSERVIC=X'0000000000000185000007EC0'
```

NTO devices

The terminal must identify itself as an NTO device in one of the following ways:

- LU presentation services profile in the BIND image must specify LU1.
- When the LUTYPE is NTO, IMS uses the data stream compatibility byte to precisely define the specific NTO device type. If the LUTYPE specified is SLUTYPE1, the data stream compatibility byte is ignored. The terminal is assumed to be an actual SLUTYPE1 as defined on the logon descriptor.
- Logon exit routine, if available, must accept the DFSNTO (default) logon descriptor or specify an NTO logon descriptor that is defined by your installation using the node name or LOGOND user data parameter.
- The data stream compatibility byte in the CINIT specifies the device type. WTTY indicates a TTY NTO device. If nothing is specified, the device type is LUNS NTO.

LU2 and non-SNA 3270 screen size and model information

After the Logon exit routine approves the logon, BIND image data (screen-size) and feature information from the logon descriptor are used to search the MFS device characteristics table for the appropriate MFS device information.

If the model is specified in the screen size control byte the MFS device characteristics table is not searched. A DFS3646 error message is issued if no match exists on screen-size and feature. The information from the proper MFS device characteristics table entry is then used for that device. This information in the MFS device characteristics table comes from the IMS system definition or the MFSDCT utility (DFSUTB00).

If the screen-size control byte is X'7F' and both the default and alternate screen-sizes are specified, a search of the MFS device characteristics table using alternate screen size commences. If no match is found, another search begins using the default screen size. If no screen size is found, message DFS3646I is issued to the operator.

The screen size (the product of the lines and columns) must be in the range 80-16384. The lines and columns must each be in the range 1-255.

If the screen size and features of the 3270 device that is logging on maps into two or more MFS device characteristics table entries, the first entry in the table that matches the screen size and features is selected.

If the screen-size control byte (the 11th byte of the PSERVIC on the VTAM MODEENT) is X'00' and both default and alternate screen-sizes are specified, a search of the MFS device characteristics table using the default screen-size occurs. If no match is found, another search begins using the alternate screen-size.

Use the Logon exit routine, DFSLGNX0, to override the screen-size or model for the device during logon.

LU2 screen-size and model information:

Screen-size and model information applies to LU type-2 devices.

The following is an algorithm IMS uses to determine the model or screen size for ETO terminals:

1. Screen-size is established based on the model specification in the VTAM PSERVIC. IMS determines the model by checking the screen-size control byte in the VTAM PSERVIC field for an X'01', X'02', or X'03'. The model is established accordingly. Screen-size specifications are ignored when the model is specified. X'01' represents a model 1 with a 12x40 screen-size, and X'02' and X'03' represent a model 2 with a 24x80 screen-size. If the model is specified, the MFS device characteristics table is not searched for MFS device information. The features are obtained from the logon descriptor.
2. Screen-size is established based on the screen-size specifications in the VTAM PSERVIC. A value of X'7E' in the screen-size control byte causes the default screen size in the VTAM PSERVIC to be used. A X'7F' causes the alternate screen-size to be used first to search the MFS device characteristics table. If no match is found, the default screen-size is used to search; the first match is the screen-size. If no match is found, message DFS3646I is sent to the operator.
3. If the screen-size control byte is X'00', the default and alternate screen-size specifications in the VTAM PSERVIC are used to search the MFS device characteristics table for MFS device information. If a match is made on the default size, the default is used. If a match is made on the alternate size, the alternate is used. If no match is made, the logon is rejected.
4. If the screen-size control byte is X'00' and no screen-size is specified, the device is defaulted to a model 2 device, and the screen-size is established (24x80).
5. The screen-size is established in the BIND parameter in the default screen-size field. The erase write (EW) command is always used.

If you want to override the model value, use the Logon exit routine (DFSLGNX0). Valid values are X'01' and X'02', corresponding to model 1 and model 2 in the logon descriptor. Screen-size and model specifications in the VTAM PSERVIC are ignored. Model=X'01' represents a 12x40 screen-size. Model=X'02' represents a 24x80 screen-size. The MFS device characteristics table is not searched for MFS device information. The features are obtained from the logon descriptor.

If you want to override the screen-size value, use the Logon exit routine, DFSLGNX0. Be aware that if you override the model, the screen-size override is ignored. The features from the logon descriptor and the screen-size are used to search the MFS device characteristics table for the MFS device information.

Non-SNA 3270 screen-size and model information:

Model information applies to VTAM 3270 Record-Mode devices (non-SNA 3270s).

To determine the model or screen-size for ETO terminals IMS uses the following algorithm:

1. Screen-size is established based on the model specification in the VTAM PSERVIC field. IMS determines the model by checking the screen-size control byte in the VTAM PSERVIC field for X'01', X'02', or X'03'. The model is established accordingly. Screen-size specifications are ignored when the model is specified. X'01' represents a model 1 with a 12x40 screen-size, and X'02' and X'03' represent a model 2 with a 24x80 screen-size. If the model is specified, the MFS device characteristics table is not searched for MFS device information. The features are obtained from the logon descriptor. Device type and screen-size are determined from the model value.
2. Screen-size is established based on the screen-size specifications in the VTAM PSERVIC field. A value of X'7E' in the screen-size control byte causes the default screen size in the VTAM PSERVIC field to be used. A X'7F' causes the alternate screen-size to be used first in searching the MFS device characteristics table. If no match is found, the default screen-size is used to search; the first match is the screen-size. If no match is found, message DFS3646I is sent to the operator. Device type and screen-size are implied by the model value.
3. If the screen-size control byte is X'00', the default and alternate screen-size specifications in the VTAM PSERVIC field are used to search the MFS device characteristics table for MFS device information. If a match is made on the default size, the default is used. If a match is made on the alternate size, the alternate is used. If no match is found, the logon is rejected. Device type and screen-size are implied by the model value.
4. If no model information and screen-size are specified in the VTAM PSERVIC field, IMS uses the CINIT's model byte. Model type is established by the VTAM definition statement FEATUR2. The model information applies only if the screen-size control byte is X'00' and if the screen-size is also X'00'. This applies to printers and displays. X'00' is the default and corresponds to a model 1. X'01' corresponds to a model 2. Device type and screen-size are implied by the model value.
5. The screen-size is used to determine the type of write command used. If the screen-size is equal to 960 or greater than 1920, IMS uses erase write alternate (EWA). If the screen-size is less than or equal to 1920 and not equal to 960, IMS uses erase write (EW). The device type and screen size are determined from the model value.

If you want to override the model value, use the Logon exit routine (DFSLGNX0). Valid values are X'01' and X'02', corresponding to model 1 and model 2 in the logon descriptor. Screen-size and model specifications in the VTAM PSERVIC are ignored. Model=X'01' represents a 12x40 screen-size. Model=X'02' represents a 24x80 screen-size. If the model is specified, the MFS device characteristics table is not searched for MFS device information. The features are obtained from the logon descriptor. Device type and screen-size are determined from the model value.

If you want to override the screen-size value, use the Logon exit routine (DFSLGNX0). Be aware that if you override the model, the screen-size override is ignored. The features from the logon descriptor and the screen-size are used to search the MFS device characteristics table for the MFS device information. Device type and screen-size are implied by the model value.

Screen definition examples:

The following examples show the PESRVIC parameter in the VTAM mode table, what the equivalent TERMINAL macro parameters would be for a static terminal, and the corresponding MFS DEV statement TYPE parameter, which is used for both static and ETO terminals.

- LU0 (non-SNA 3270 video)
 - VTAM mode table: PESRVIC=X'00000000000000000000200'
 - TERMINAL macro: UNITYPE=3270 UNIT=3284/86 MODEL=2
 - MFS DEV statement: TYPE=(3270P,2)
- Model 2 non-SNA 3270 Printer
 - VTAM mode table: PESRVIC=X'00000000000000000000240'
 - TERMINAL macro: UNITYPE=3270 UNIT=3284/86 MODEL=2
 - MFS DEV statement: TYPE=(3270P,2)
- Non-SNA 3270 Display (model specified)
 - VTAM mode table: PESRVIC=X'00000000000000000000080'
 - TERMINAL macro: UNITYPE=3270 MODEL=1
 - MFS DEV statement: TYPE=(3270,1)

Model information can come from the FEATUR2 parameter for non-SNA 3270; if this parameter is not specified, this is a model 1 (screen-size 12x40). Assume that FEATUR2=1 (specified or default).

- Non-SNA 3270 Display (model specified)
 - VTAM mode table: PESRVIC=X'00000000000000000000080'
 - TERMINAL macro: UNITYPE=3270 MODEL=2
 - MFS DEV statement: TYPE=(3270,2)

Model information can come from the FEATUR2 parameter for non-SNA 3270; if this parameter is not specified, this is a model 2 (screen-size 24x80). Assume that FEATUR2=2.

- Non-SNA 3270 Display (screen-size specified)
 - VTAM mode table: PESRVIC=X'000000000000185000007E80'
 - TERMINAL macro: UNITYPE=3270 TYPE=3270-A2,SIZE=(24,80)
 - MFS DEV statement: TYPE=3270-A2, where A2=24x80

For a SLU-2 device, UNITYPE=SLUTYPE2 would be specified (also note the change to the PESRVIC field).

The screen-size comes from the default screen-size field (24x80). (For a SLU-2 device, the first byte of the PESRVIC field would be X'02'. The last byte would be X'00').

- Non-SNA 3270 Display (screen-size specified)
 - VTAM mode table: PESRVIC=X'00000000000020500000080'
 - TERMINAL macro: NITYPE=3270 TYPE=3270-A3,SIZE=(32,80)
 - MFS DEV statement: TYPE=3270-A3, where A2=32x80

The screen-size comes from the default screen-size field (32x80). (For a SLU-2 device, the first byte of the PESRVIC field would be X'02'. The last byte would be X'00').

- Non-SNA 3270 Display (model specified)
 - VTAM mode table: PESRVIC=X'00000000000000000000280'
 - TERMINAL macro: UNITYPE=3270 MODEL=2

- MFS DEV statement: TYPE=(3270,2)
This is a model 2 non-SNA 3270 display (24x80).
- SNA 3270 Display (model specified)
 - VTAM mode table: PSERVIC=X'020000000000000000000000280'
 - TERMINAL macro: UNITYPE=3270 MODEL=2
 - MFS DEV statement: TYPE=(3270,2)
This is a model 2 SNA 3270 display (24x80).

Planning a high-security environment with ETO

ETO enhances the security of your IMS system. You can customize the ETO security features for your installation needs.

For example, you can customize exit routines that apply to both static terminals and ETO dynamic terminals.

The ETO security features allow you to control each of the following:

- The physical connection of terminals to IMS.
- User signon to IMS.
- Output to users or nodes.
- Message queuing to users. You can customize message queuing two ways:
 - Automatically allocate an LTERM to a user that you identify at signon
 - Use the Destination Creation exit routine (DFSINSX0)
- Command and transaction security, by using RACF (or an equivalent SAF-compliant security product).

Static versus dynamic terminals

You define static terminals during system definition to associate an LTERM with a particular physical terminal. Any user at a terminal can receive output messages that are queued for that terminal, regardless of whether signon is required for that terminal.

The major benefit of ETO is that dynamic LTERMs (output message queues) are managed separately from the terminals and are assigned to a user name. The user can obtain output only after signing on. This dynamic LTERM-to-user association is maintained until the user signs off, and it remains after signoff if IMS does not delete the user structure.

Planning for MFS

After you implement ETO dynamic terminal support, the number and diversity of terminal types is likely to increase. When a terminal dynamically establishes an IMS session, MFS formats might not be available for the device type that is requesting the session.

You can solve this problem by:

- Restricting device types to only those that are used in the MFS definitions for your application programs.

Recommendation: Use the Logon exit routine (DFSLGNX0) to select the desired logon descriptor.

- Extending the MFS device output formats to include the new terminal types that connect to the IMS.

Recommendations:

- Use the MFS Generation utility (MFSGU) facilities, such as the STACK statement, to make creating these additional formats easier.
- Use the MFSDCT utility (DFSUTB00) to avoid needing to perform an IMS generation.

Often, multiple device types have similar or identical capabilities, so the distinction of device type might not be important to your environment. Although IMS application programs are usually not sensitive to device type, such a dependency can exist. For example, MFS can create different input messages from the same input data stream, due to differences in the format specification for device input.

Planning user IDs

When planning user IDs for use with ETO there are a number steps you should take.

The steps include:

- Ensure that dynamic user structures have unique names in IMS. Support for unique user IDs depends on whether user-based security is important. For user-based security, user IDs must be unique to the user. Otherwise, user IDs can be the same as the terminal LU name.
- Code the Signon exit routine (DFSSGNX0) to provide user ID suffixing, if users need to sign on to more than one terminal with the same user ID at the same time. Review and select the SGN and RCF EXEC keyword parameters.
- Analyze how application programs use the user field in the I/O PCB.

Planning user queue names

When planning user queue names for use with ETO there are a number steps you should take.

When planning user queue names:

- Ensure that dynamic user queue names are unique in the IMS. Support for unique queue names depends on whether user-based security is important. For user-based security, user queue names must be unique to the user. Otherwise, user queue names can be the same as the terminal LU name.
- Specify user descriptors or provide logic in the Signon exit routine (DFSSGNX0) if the user queue name cannot be the same as the user name.
- Analyze how application programs use the LTERM field in the I/O PCB.

Planning operations

When planning for operations there are a number of steps you should perform.

The steps you should take when planning for operations include:

- Update your MTO procedures to reflect the concept of dynamic resources. The MTO needs to become familiar with the command input fields and response formats.
- Update the help-desk procedures to reflect the IMS terminal and user resource structures, as well as the commands that are required to diagnose user's problems.

- Review and update automated-operator procedures, if necessary.
- Develop procedures for handling the dead-letter queue.
- Identify requirements for autosignoff and autologoff.
- Develop standards for the conditions under which system-wide values should be used and when they should be overridden.

Planning for MSC support with ETO

Although MSC physical and logical links are predefined during system definition, you can dynamically create MSC remote LTERMs during IMS initialization.

You can identify a remote MSC NAME to IMS using an MSC descriptor. The descriptor relates each remote resource to the link path name of a generated MSNAME macro.

Recommendation: Define remote LTERMs with MSC descriptors to maintain consistency between remote and local systems.

If you choose not to use MSC descriptors in order, the MSC Verification utility, DFSUMSV0, recognizes the remote LTERMs but not the corresponding local LTERMs in the target system. In this case, IMS issues warning message DFS2331W.

IMS processes MSC descriptors that are associated with MSC links defined within the IMS. IMS ignores other MSC descriptors. You can maintain a single network definition for remote LTERMs in IMS.PROCLIB for multiple interconnected IMS systems that use MSC. In this case, each logical link path name must be unique throughout the entire network.

Coding ETO descriptors

IMS uses ETO descriptors to dynamically build terminal structures and user structures.

The four types of ETO descriptors are:

- Logon
- User
- MFS device
- MSC

The basic format for all ETO descriptors is:

Column

Description

1	One of the following descriptor types:
L	Logon descriptor
U	User descriptor
D	MFS device descriptor
M	MSC descriptor
*	Comment line (ignored by IMS)
3-10	Name of the descriptor, using the following conventions: • Must be one to eight alphanumeric characters.

- For logon and user descriptors, characters are limited to A-Z, #, \$, and @.
- For MSC descriptors, use the link name.

12-72 One or more keywords and their parameters, separated by blanks. Use commas to separate multiple parameters for a single keyword.

73-80 Optional sequence numbers (ignored by IMS).

Related concepts:

 ETO descriptors (System Definition)

Creating descriptors using the system definition process

To ease migration, you can create a starter set of ETO descriptors (except MFS descriptors) using the ETOFEAT keyword on the IMSCTRL macro at system definition. You can add to or modify the starter set by creating an additional IMS.PROCLIB member.

For MFS descriptors, the system definition creates a device characteristics table based on the stage-1 input that can be used as one of the inputs to the MFSDCT utility (DFSUTB00). The device characteristics table contains what the system-defined device characteristics table contains, plus any additional entries from device descriptors.

Restriction: ETO descriptors are not created during system definition if the LGEN parameter is specified on the SYSTEM= keyword in the IMSCTRL system definition macro.

Creating descriptors during the system definition process saves you time and ensures that the descriptors are correct. IMS generates an ETO descriptor report, which provides information on the relationship between ETO descriptors and the IMS system definition resources that they represent.

Related concepts:

 Including IMS ETO Support in the IMS system (System Definition)

Related reference:

 IMSCTRL macro (System Definition)

Storing descriptors

Descriptors that are created during system definition are stored in the IMS.PROCLIB member, DFSDSCMx. Subsequent system definitions of the same stage-1 input deck overwrite the DFSDSCMx member.

Recommendation: To avoid losing descriptors when member DFSDSCMx is replaced, store descriptors that you create by using TSO or z/OS utilities in IMS.PROCLIB member, DFSDSCTy. If you need to update descriptors that are created at system definition in DFSDSCMx, use TSO or a z/OS utility (such as IEBUPDTE) to make the updates.

Creating logon descriptors

Logon descriptors provide IMS with information about the physical characteristics of the terminals that establish logon sessions. These characteristics must be consistent with the VTAM logon BIND characteristics.

This topic describes how to create and use logon descriptors.

Creating logon descriptors during system definition

When you specify IMS system definition options to create ETO logon descriptors, IMS dynamically creates a logon descriptor for every unique VTAM TERMINAL macro that you have specified (For each TERMINAL macro definition, a node user descriptor is also created for use as a migration step to ETO).

IMS system definition can produce up to 37 common logon descriptors for each device type. The descriptor that defines the largest number of terminals of that type becomes the default logon descriptor. This logon descriptor assumes the IMS-defined default name for that type. The other terminals of that type create their own unique logon descriptor by using a suffix on the default name.

For terminal definitions that do not match one of the 37 common descriptors, IMS creates an individual logon descriptor.

These descriptors are generated as comments (with an asterisk in column 1). For example, *L3270A. To choose a descriptor that you need, remove the asterisk.

The naming convention for the 37 common logon descriptors that are created during the system definition process is:

- The last character of the name must be unique.
- Blank for the most common, then 0-9 and A-Z.

Restrictions:

- During system definition, IMS does not create ETO logon descriptors for the primary or secondary master terminal, or for LU 6.1 terminals that are defined as XRF ISC links.
- IMS does not generate logon descriptors if the LGEN parameter is specified on the SYSTEM= keyword in the IMSCTRL system definition macro.
- The following keywords are not supported on logon descriptors: PU, SIZE, MODEL, TYPE, MSGDEL.

Criteria for selecting logon descriptors

The logon descriptor contains data that is related to the node and the VTAM CINIT. This data allows IMS to create a control block structure that supports a session.

IMS uses the following criteria (in sequence) to select a logon descriptor:

1. IMS uses existing control blocks for terminals. IMS does not look for a descriptor if it finds existing control blocks.
2. IMS determines whether the Logon exit routine is used to define logon descriptor names.

For VTAM terminals, the exit routine extracts the name from the VTAM CINIT user data, or another appropriate algorithm. The LUTYPE and TS= (transmission service level) fields in the VTAM CINIT data must agree with the selected descriptor; otherwise, the logon is rejected. IMS rejects invalid logon descriptor names.

3. In response to a request to establish a session (logon), IMS examines the LOGOND parameter. The LOGOND parameter can indicate the logon descriptor for IMS to use.

LOGOND is a parameter in the VTAM CINIT user data. LOGOND is also a keyword on the IMS /OPNDST command.

Restriction: LOGOND is not valid for ISC parallel sessions that use VTAM.

4. If IMS does not find a logon descriptor name, it looks for a logon descriptor with the same name as the VTAM CINIT LUNAME.

Recommendation: Use the node name on the logon descriptor if you expect any of the following situations:

- You do not expect to add more of the same terminal type.
- You do not have many of the terminal type.
- You want to simplify the logic in the Logon exit routine.
- For ISC, you want to specify parameters (such as OUTBUF=) other than the default ISC logon descriptor, and you have not coded an exit routine.

If these criteria do not yield a valid logon descriptor name, IMS selects a descriptor by using the default criteria.

Criteria for selecting a default logon descriptor

IMS provides a default logon descriptor for each VTAM terminal type. The logon descriptor name is based on the LU type and TS profile.

If necessary, you can add logon descriptors. If you do not provide a logon descriptor, and IMS does not locate one in IMS.PROCLIB or in the Logon exit routine (DFSLGNX0), IMS uses the algorithm shown in the following table to assign a default logon descriptor name. The VTAM LUTYPE, IMS UNITYPE, and IMS default logon descriptor names are shown for each terminal type.

The default logon descriptor is determined in one of three ways:

- IMS looks for the most common logon descriptor created during IMS system definition and uses it as the default.
- If no terminal definition exists for that terminal type, IMS creates a default logon descriptor.
- Your installation can create a default logon descriptor.

Table 5. Mapping for VTAM LUTYPE, IMS UNITYPE, and default logon descriptor names

VTAM LUTYPE	IMS UNITYPE	Default descriptor name
X'06'	LUTYPE6	DFSLU61
X'02'	SLUTYPE2	DFSSLU2
X'01'	SLUTYPE1 (Default) ¹ NTO	DFSSLU1 DFSNT0
X'00' (TS = X'04')	SLUTYPEP (Default) ² Finance 3601	DFSSLUP DFSFIN DFSFIN
X'00' (TS = X'02' or X'03')	3270 3270, UNIT=3284	DFS3270 DFS327P

Notes to table:

1. IMS does not distinguish between SLUTYPE1 and NTO terminals; SLUTYPE1 is the default. To support both SLUTYPE1 and NTO terminals, do the following:

- Override the DFSSLU1 default logon descriptor by using the Logon exit routine (DFSLGNX0) or by specifying the LOGOND parameter.
 - If your installation has no SLUTYPE1 terminals, rename the DFSNTO logon descriptor to DFSSLU1 in order to make it the default for LU type X'01' terminals.
2. IMS does not distinguish between SLU type-P, Finance, and 3601 terminals; SLUTYPEP is the default. ETO considers Finance and 3600 series terminals to be the same. To support both SLUTYPEP and 3600/Finance terminals, do the following:
- Override the DFSSLUP default logon descriptor by using the Logon exit routine (DFSLGNX0) or by specifying the LOGOND parameter.
 - If your installation has no SLUTYPEP terminals, rename the DFSFIN logon descriptor to DFSSLUP in order to make it the default for LU type X'00' (TS = X'04') terminals.

Using NTO, 3600/Finance terminals

Because of conflicting CINIT information, IMS cannot generate DFSNTO or DFSFIN as a default logon descriptor name. CINIT parameters for SLU 1 are identical to NTO, and SLU P is identical to 3600/Finance.

Recommendations: For 3600/Finance and NTO terminal types, take each of the following actions:

- Always supply the appropriate logon descriptor name in the Logon exit routine (DFSLGNX0). You can generate this name as a constant, based on LU name, LU type, or CINIT user data.
- Ensure that the logon descriptor name is provided as the LOGOND parameter from the BIND user data.
- Ensure existence of a logon descriptor that matches the node name of the terminal that is logging on. If IMS does not find a default logon descriptor, the logon attempt fails.

Recovering ETO terminals using XRF

To recover ETO terminals in an XRF environment, use the BACKUP= parameter when specifying the logon descriptor.

The BACKUP= parameter sets the priority that controls the order in which sessions are switched.

When 3600/Finance and SLU-P terminals are defined as class-2 terminals in an XRF environment, automatic re-logon and re-signon occur during a takeover.

Related reading: For general information on XRF, see *IMS Version 11 System Administration*.

Creating user descriptors

User descriptors provide information relating to user options and user structure names. IMS needs user descriptors in order to create control blocks that enable users to use ETO terminals.

The three types of user descriptors are:

- Installation-created
- Node user
- DFSUSER

IMS chooses the first valid descriptor, using the given sequence. IMS creates DFSUSER and node user descriptors using system definition options.

This topic describes how to create user descriptors.

Creating user descriptors during system definition

User descriptors are generated from each VTAM TERMINAL or VTAMPOOL SUBPOOL macro.

User descriptors that are created during system definition by using the VTAM TERMINAL macro have the same name as the terminal. User descriptors that are created by using the VTAMPOOL SUBPOOL macro have the same name as the subpool.

For user descriptors that are created by using a SUBPOOL macro, you cannot set a response option (TRANSRESP, NORESP, FORCRESP), because it is defined on the TERMINAL macro for static definitions. You need to add the response option (appropriate for your installation) to any user descriptor that is created with a SUBPOOL macro.

Criteria for selecting user descriptors

When a user signs on, if the user structure does not exist, IMS selects a user descriptor to build the user structure.

IMS selects the user descriptor according to the following criteria:

1. An installation-written exit routine can select the name of the user descriptor (user ID, node name, or DFSUSER).
2. IMS looks for USERD, provided at logon. IMS also looks for a descriptor name (user ID, node name, or DFSUSER) that is specified in the USER descriptor field of the /SIGN or /OPNDST command.
3. IMS looks for a descriptor that has the same name as the user ID (installation-created user descriptor only). If IMS finds one, and an LTERM keyword is not specified in the descriptor, IMS creates a user structure and a single LTERM, both of which have the same name as the user ID.
4. IMS looks for a descriptor that has the same name as the VTAM node (node user descriptor). If IMS finds one, and an LTERM keyword is not specified in the descriptor, IMS creates a user structure and a single LTERM, both of which have the same name as the VTAM node. However, no output security is associated with this user structure. Any user that signs on at a terminal can receive messages that are queued for that terminal.
5. IMS selects the default user descriptor, DFSUSER. IMS creates a user structure and a single LTERM, both of which have the same name as the user ID.

Using installation-created user descriptors

You can create your own installation-specific user descriptors that meet important criteria for your site.

The name of the installation-created user descriptor is the same as the user ID. You can add values to these user descriptors that are not provided with node user descriptors or DFSUSER descriptors.

Using node user descriptors

Node user descriptors help in migrating from static terminal definitions to ETO dynamic terminal definitions.

During IMS system definition, node user descriptors are created as an option. Node user descriptors can be useful when exit routines that perform descriptor selection are not yet complete.

Node user descriptors must have the same name as their associated terminals. Therefore, IMS creates unique node user descriptors that retain user options and user structure names that exist in the system definition. When IMS system definition creates ETO descriptors, a node user descriptor is created for each terminal that has a VTAM TERMINAL macro or VTAMPOOL SUBPOOL macro. Node user descriptors are created even when they match the DFSUSER options (No descriptor is created for an ISC terminal that is defined for parallel session support).

If the LTERM parameter of the node user descriptor contains the same name as a statically defined LTERM, the node user descriptor is ignored. IMS issues message DFS3673W. Consequently, all node user descriptors that are created during system definition are generated as comment statements. To use the node user descriptors, remove the asterisks.

In most cases, node user descriptors are not needed. You only require node user descriptors when the default user descriptor (DFSUSER) or the Signon exit routine (DFSSGNX0) cannot supply the user options you desire. Using a TSO or a z/OS utility, you can discard most of the node user descriptors that are created during system definition.

Be aware that when you use node user descriptors, you cannot use ETO's ability to eliminate the need for predefining terminals that connect to IMS. Using node user descriptors lose output security for each user as well.

Recommendations:

- If continuous availability of IMS is critical, eliminate node user descriptors as soon as possible. Because adding new terminals by using node user descriptors requires an IMS restart, use exit routines instead of node user descriptors.
- To avoid using node user descriptors, use the Signon exit routine (DFSSGNX0) to have IMS build a user structure using the node name as the user name. Although this is similar to using node user descriptors, it avoids predefining large numbers of these descriptors at IMS initialization.
- You can use node user descriptors to start a session when output is available, using /OPNDST, LOGON, or autologon. To use autologon, specify the SHARE option on the TERMINAL macro.
- If you select a default user descriptor using the Signon exit routine (DFSSGNX0), you can set a bit to indicate that IMS should create the user structure using the node name for the user structure name. This is the same as selecting a node user descriptor. By creating node-name structures by using the Signon exit routine, rather than by defining a node user descriptor for each terminal, large network installations benefit from decreased complexity.

Using DFSUSER user descriptors

If IMS does not find a user descriptor that has the same name as the user ID or the terminal that is signing on, and no exit routine has provided one, IMS uses DFSUSER as the default descriptor.

IMS uses DFSUSER for both signon and application program output processing. Using DFSUSER also enables IMS to dynamically create a user structure for a signon request when no other user descriptor is available.

IMS creates DFSUSER when you specify the system definition options to create ETO descriptors. DFSUSER contains the options specified most frequently in your IMS system definition, and it should satisfy most users. However, using the DFSUSER descriptor does not require you to use all of these options. You can code the Signon exit routine (DFSSGNX0) to make changes to the structures that are built by using DFSUSER.

DFSUSER builds a user structure that has the same name as the user ID that is specified on the /SIGN command. IMS allocates to this user structure a single LTERM, which also has the same name as the user ID. You can use the Signon exit routine (DFSSGNX0) in order to supply multiple queues, if necessary.

Recommendation: Fully implementing ETO involves creating most user resources using DFSUSER. After migrating to ETO, use the DFSUSER descriptor for as many users as possible. Minimizing the number of descriptors reduces the administrative workload of an IMS network.

The user structure name becomes the user ID. When you use the DFSUSER descriptor, you can modify user structure names and other options by using the Signon exit routine (DFSSGNX0).

If an application program has a dependency on an LTERM name that exists in an I/O PCB, you can customize ETO in one of two ways:

- Provide a user descriptor for each user that is to use a particular application program.
- Use the Signon exit routine (DFSSGNX0) to tailor the default action for a subset of the user population.

The Signon exit routine (DFSSGNX0) and the Destination Creation exit routine (DFSINSX0) enable you to dynamically create user structures, even when the specified user descriptors do not contain data to create LTERMs.

Related reference:

 Signon exit routine (DFSSGNX0) (Exit Routines)

 Destination Creation exit routine (DFSINSX0) (Exit Routines)

Creating MFS device descriptors

With MFS device descriptors, you can define screen size and feature combinations that are not generated during IMS system definition.

The MFSDCT utility (DFSUTB00) uses the MFS device descriptors in order to update device type, screen size, and features in the MFS device characteristics table. MFSDCT also uses MFS device descriptors to generate new MFS default formats, without requiring changes to the system definition. The MFS device descriptors are not created as part of a system definition, but are instead an optional step after IMS system definition.

MFS device descriptors define terminal characteristics of dynamic terminals that differ from static terminals. Unless you use exit routines to override screen sizes, you need to ensure that all ETO terminal screen sizes that are different from those that are defined for static terminals are defined using MFS device descriptors.

Recommendation: If you have users who log on to terminals that have different device characteristics, you need to create an expanded set of MFS formats. Input

and output messages that are delivered to non-originating terminals are formatted to the terminal according to the MFS formatting specifications that are defined for that device.

Building the device characteristics table

When someone logs onto an ETO 3270 or SLU-2 terminal that is connected to IMS, usually the screen size is provided to IMS in the VTAM CINIT PSERVIC data. Sometimes the model number is also provided.

IMS uses this screen size and the features that are specified on the logon descriptor in order to search the MFS device characteristics table, which contains one or more entries. Each entry identifies the following:

- The MFS device type
- The screen size
- The features

Searching the table, you can find the MFS device type. If the search is unsuccessful, the logon attempt is rejected, and messages DFS3646I and DFS3672I are issued.

The MFS device characteristics table is built in one of two ways:

- The IMS system definition process builds an MFS device characteristics table based on the TYPE=, SIZE=, and FEATURE= specifications of the TERMINAL macro. Each unique combination has an entry in the MFS device characteristics table.
- The MFSDCT utility (DFSUTB00) builds or modifies an MFS device characteristics table based on the input of the MFS device descriptor. These descriptors have the TYPE=, SIZE=, and FEATURE= specifications that the TERMINAL macro has. When the VTAM CINIT PSERVIC indicates a model number instead of a screen size, the MFS device characteristics table is not searched. A model number indicates a certain screen size and MFS device type.

The MFS device characteristics table is not searched if the VTAM CINIT PSERVIC field contains a model number instead of a screen size.

Related reading: For more information on using the MFSDCT utility (DFSUTB00), see *IMS Version 11 System Utilities*.

Using the MFSDCT utility (DFSUTB00)

IMS searches the device characteristics table in order to locate MFS device information during session initiation only in certain situations.

IMS searches the device characteristics table during session initiation in the following circumstances:

- The screen size is overridden.
- The model value in PSERVIC is X'00', X'7E', or X'7F'.

The MFSDCT utility (DFSUTB00) enables you to define screen sizes that are not defined during IMS system definition. The MFSDCT utility uses MFS device descriptors in PROCLIB member DFSDSCMx and DFSDSCTy, without performing an IMS system definition. The new screen size definitions are added to those that are already defined.

The MFS utility does the following:

- Reads device descriptors from IMS.PROCLIB members DFSDSCMx and DFSDSCTy.
- Builds one device descriptor table entry statement for each new device descriptor.
- Terminates if no descriptors are specified.
- Optionally loads the existing device characteristics table from IMS.SDFSRESL, and builds one entry statement for each existing table entry.
- Passes the table entry statements and the DCTBLD and MFSINIT macros as input to assembler.
- Prepares assembler output as a new or updated device characteristics table, as well as a new set of default MFS format definitions. Output is in separate files for subsequent processing.

The MFSDCT utility procedure is found in IMS.PROCLIB. You must evaluate the screen size requirements and code MFS device descriptors to meet those requirements. The MFSDCT utility must generate all possible combinations of screen sizes and features that your installation might require.

To use the MFSDCT utility, follow these steps:

1. Execute DFSUTB00 (the MFSDCT utility).
2. Assemble the new device characteristics table.
3. Link edit the new device characteristics table into IMS.SDFSRESL.
4. Execute phase 1 of the MFS Language utility in order to generate new default MFS format control blocks.
5. Execute phase 2 of the MFS Language utility; the new default MFS formats are loaded into IMS.FORMAT.

Problems with the device characteristics table are often indicated by error messages DFS3646I and DFS3672I during logon processing.

Related reference:

 MFS Device Characteristics Table utility (DFSUTB00) (System Utilities)

Creating MSC descriptors

MSC descriptors relate remote LTERMs to statically defined MSC links.

IMS creates MSC descriptors when you specify IMS system definition options to create ETO terminal descriptors. MSC descriptors relate remote NAME macros to defined MSC links.

Recommendation: Define all LTERMs in remote IMS systems by using MSC descriptors. The MSC Verification utility in the target IMS system can then associate the remote LTERMs with the corresponding local LTERMs. IMS issues message DFS2331W if the corresponding local LTERMs are not found.

IMS only processes those MSC descriptors that are associated with the MSC links that are defined within the system being initialized. IMS ignores all other MSC descriptors. If each logical link path name is unique in the network, you can maintain your network's MSC definition source in a single PROCLIB member, IMS.PROCLIB.

Exit routines

You can customize ETO with a variety of exit routines delivered with IMS.

You can use the following exit routines to customize ETO:

- Initialization exit routine (DFSINTX0)
- Logon exit routine (DFSLGNX0)
- Logoff exit routine (DFSLGFX0)
- Signon exit routine (DFSSGNX0)
- Signoff exit routine (DFSSGFX0)
- Destination Creation exit routine (DFSINSX0)
- Greetings Message exit routine (DFSGMSG0)

If ETO is enabled, these exit routines are loaded during IMS initialization. All non-MSU and non-LU 6.2 VTAM terminals (static or dynamic) can use these exit routines.

You can code these ETO exit routines with a wide range of processing logic for any of the following purposes:

- Enforcing naming conventions
- Selecting user queues
- Overriding terminal characteristics, such as screen size
- Enhancing security by limiting access to IMS terminals
- Overriding security by allowing signon without using a security product, such as RACF
- Selecting operational parameters, such as timeout values

Related reference:

 Transaction Manager exit routines (Exit Routines)

Starting ETO

To include ETO in your IMS system, specify ETO=Y in the IMS or DCC startup procedure.

The default for the ETO= keyword in the startup procedures is ETO=N. If ETO=N remains specified, IMS rejects requests to establish sessions for undefined terminals. Messages that are destined for nonexistent queues are refused, and IMS issues message DFS064 or an A1 status code.

When IMS initializes, it calls the Initialization exit routine (DFSINTX0). You can code this exit routine to disable ETO, or to create and load data that you want ETO to use. IMS maintains a pointer to this data and passes this pointer to the ETO exit routines as an input parameter. IMS passes this pointer to non-ETO exit routines through the SCDINTXP field in the system contents directory (SCD).

When you enable ETO, IMS validates each ETO descriptor:

- If a descriptor is coded incorrectly, IMS ignores it and issues message DFS3640W to the MTO.
- If IMS detects an invalid parameter within a valid descriptor, IMS substitutes the default parameter, and processing continues with message DFS3641W.

- If IMS is unable to read all the descriptors (for example, because of read errors), IMS abends with abend U0015. IMS does not abend if it finds one valid logon descriptor and one valid user descriptor.
- IMS does not start a system whose descriptor information is incomplete.

Descriptors, exit routines, and the MFS device characteristics table can be added, deleted, or updated before IMS initialization time. However, if a control block for a session exists across restart, a change in the descriptor that built the control block does not take effect until after the control block is deleted. For example, 3600/Finance, SLU-P, and ISC sessions can warm start with control blocks created from an original descriptor, even when that descriptor is changed or deleted after the control blocks were created.

Logging onto ETO terminals

You can log on to your terminal in three ways.

You can use:

- The IMS /OPNDST command, and optionally include signon data.
- The SNA commands INITSELF, INITOTHER, or USS LOGON, and optionally include user data for signon.
- ETO autologon, which includes user data based on user descriptors or exit routines. The user data is then passed to signon.

If you specify ETO=Y, you can establish sessions with ETO terminals, but only if each of the following is true:

- An available logon descriptor provides sufficient information to create the control blocks necessary to accept the session request.
- In addition to the required logon descriptor, a 3270 or SLU-2 terminal requires an entry in the MFS device characteristics table that matches its screen-size and features. However, IMS does not search the MFS device characteristics table if the 3270 or SLU-2 terminal is identified as model 1 or 2.
- If the screen-size control byte is X'7F' in the PSERVIC field of the VTAM MODEENT macro, IMS searches the MFS device characteristics table using the alternate screen-size. If no match is found, IMS searches the MFS device characteristics table using the default screen-size. If you specify both sizes, and you want to use the default screen-size, the Logon exit routine (DFSLGNX0) must specify the default screen-size as an override.
- If the user cannot specify the logon descriptor in the user data, you must use the Logon exit routine (DFSLGNX0) or let IMS choose the default logon descriptor based on the terminal type.

Limiting dynamic logon to specific terminal types

If you want to limit dynamic logon to specific terminal types, delete the default logon descriptor for those terminal types you do not want to logon dynamically.

You can also reject the default logon for specific terminal types by using the Logon exit routine (DFSLGNX0). Dynamic logon for these terminal types fails.

Related concepts:

“Criteria for selecting a default logon descriptor” on page 103

Creating and reusing LTERM control blocks

A user structure containing an LTERM can be created in a number of different ways.

A user structure containing an LTERM is created:

- When the user signs on, and the user structure does not currently exist
- When an asynchronous message is created for an LTERM, and the LTERM does not currently exist
- When you use the /ASSIGN command to assign an LTERM to a non-existent user.
- The /ASSIGN command is used to assign a non-existent LTERM to a user.
- The /CHANGE USER *user* AUTOLOGON command is directed to a non-existent user.

The user structure is allocated to a terminal after successful signon.

Using default CINIT or BIND user data formats

Each request for session initiation can include VTAM CINIT or BIND user data to provide logon descriptor or signon data. Your installation can provide a logon exit routine to process this data.

IMS can receive optional user data when you establish a session using one of the following methods:⁸

- Using an IMS /OPNDST command
- Using autologon
- The RTO provides a user logon

You can expand the user data formats to meet your own requirements. You can either supply the logon descriptor name in your logon exit routine by using user data, or you can create the logon descriptor name by using an IMS algorithm.

The user data appears in the CINIT user data field, and it is available to IMS when the VTAM Logon exit routine is scheduled. One optional parameter, the logon descriptor name, applies to the IMS logon process. The remaining parameters apply to the IMS signon process and, optionally, to RACF. During either process, IMS does minimum processing on the CINIT user data parameters before first calling the optional installation Logon exit routine (DFSLGNX0), and later calling the Signon exit routine (DFSSGNX0).

Although the Logon and Signon exit routines can translate any installation-defined user data format, IMS has defined a default user data format that:

- The installation can expand.
- IMS can process in the absence of exit routines. This format is the logon descriptor name followed by signon data in the same format as the IMS /SIGN command.

Restriction: For warm session initiation of an STSN device, the user data must be the same as in the original logon.

Related reference:

8. ETO ISC terminals do not have optional data available.

- ➞ IMS control region exit routines (Exit Routines)
- ➞ /SIGN command (Commands)
- ➞ Format for CINIT user data parameters (System Programming APIs)

Signing on and queue LTERM allocation

Signing on to an ETO terminal identifies a user to IMS, creates a user structure, and connects this user structure to the terminal structure.

Users at VTAM terminals can sign on by:

- Issuing the /SIGN command.
- Signing on at the DFS3649 message screen.
- Providing user data with the session initiation request.

Users that establish a dynamic VTAM session with IMS must enter valid signon data before LTERMs are allocated to the session. You can require that the signon data be validated by a security product, such as RACF.

Providing signon data

Users can enter signon data by using one of several methods.

The methods that users can use to enter signon data include:

- Using the /SIGN command
- Including the signon data with the logon user data
- Coding the Logon exit routine (DFSLGNX0)

Providing signon data for ISC, SLU-P, Finance, and output-only devices

For dynamic ISC, SLU-P, Finance, and output-only devices, you must provide user signon data at session initiation.

IMS validates the signon data and allocates LTERMs to the terminals, based on:

- User descriptors in IMS.PROCLIB
- Applicable LTERM and user option defaults
- Signon exit routine (DFSSGNX0) output

Restriction: A user cannot enter signon data from an output-only device during logon or when using the /SIGN command. For a dynamic output-only terminal, users must enter signon data at session initiation.

If a user forgets to include signon data during session initiation for an output-only device, IMS issues message DFS2085I (with return code 264) to the MTO. If signon data is omitted for dynamic ISC, SLU-P, or Finance terminals, IMS issues messages DFS3645I and DFS3672I.

Signing on multiple times

Users can concurrently sign on to one or more terminals. These terminals can be a combination of both static and dynamic terminals.

For both static and dynamic VTAM terminals, you must specify the EXEC parameter, SGN=M (multiple signons enabled).

For dynamic terminals, the name of the user structure that represents the user to IMS must be unique. You can use one of four methods to make the user structure unique:

- Use the Signon exit routine (DFSSGNX0) to return a 1- to 3-byte suffix to the user ID. The total length of the user ID plus the suffix cannot exceed eight characters.

Recommendation: Use a naming convention that ensures unique user IDs. For example, you can create suffixed user IDs using the Signon exit routine (DFSSGNX0), and check the uniqueness of these user IDs by using IMS Callable Services.

Recommendation: If you are operating in a sysplex environment, create a naming convention that ensures unique user IDs across the IMSplex by creating a suffix from the following:

- As the first part of the suffix, the Signon exit routine can attach the unique part of the IMS system identifier, which is specified on the IMSID parameter, to the user ID.
- As the second part of the suffix, use IMS Callable Services to choose an available value for the user on that IMS system.

Example: A user with a user ID of USER signs on to IMSA. A unique suffix of A1 is created through this two-step process:

1. The Signon exit routine appends the A from IMSA to the user ID.
2. The Signon exit routine invokes IMS Callable Services. The exit routine then appends a 1, representing the first user named USER to sign on to IMSA.

In this example, user IDs must be less than six characters in length.

- Use the Signon exit routine to specify that the user structure name is the same as the node name.
- Use the Signon exit routine to specify a name that is unrelated to the user ID or the node name.
- Specify a user descriptor name that is the same as the node name, causing the name of the user structure to be the same as the node name. You can do this by using one of the following methods:
 - Enter signon data that contains the user descriptor name.
 - Use the Signon exit routine to specify the node user descriptor.
 - Specify no user descriptor and let IMS use the node user descriptor by default.

Unless you specify SGN=M to enable multiple signons, the user ID for a dynamic user must be unique. If the user ID is used as the user-structure name, it must be unique, regardless of whether you specify SGN=M. If the user ID is not used as the user-structure name (for example, by using the Signon exit routine), you must specify SGN=M for multiple signons.

Recommendation: Assigning a single name to an IMS user is useful in determining output status. If you use multiple names for individual users (for example, when the Signon exit routine assigns them), you must provide a means to determine the user names that are created by signons.

Restriction: You cannot use the same name for a dynamic LTERM and a static LTERM that is defined during system definition.

Related reference:

-  Signon exit routine (DFSSGNX0) (Exit Routines)
-  IMS callable services (Exit Routines)

Receiving DFS3649A, the signon required message

The ETO user must sign on before a session can enter transactions or commands.

The user can enter the /RCLSDST command only before signing on. IMS issues message DFS3649A (indicating that a /SIGN command is required) in each of the following situations:

- After a signon failure
- After a signoff
- After a logon in which no user data is entered

Any terminal user can then issue a /SIGN ON command to begin the next session.

You do not receive the DFS3649A message in any of the following situations:

- Signon data is included in the VTAM CINIT or BIND.
- The session is an ISC, SLU-P, Finance, or 3270 printer session.
- The terminal is a SLU-1 terminal running in unattended mode.
- The terminal is the subject of an autologon attempt.

Related reference:

-  /SIGN command (Commands)

Related information:

-  DFS3649A (Messages and Codes)

Receiving DFS3650I, the session status message

After a user successfully signs on (or if signon is not required), IMS issues message DFS3650I, indicating the status of the session with IMS.

Exceptions: In the following situations, IMS does not send message DFS3650I:

- When SLU-1 terminals run in unattended mode.
- When you specify the NOTERM option on the user descriptor.

Message DFS3650I provides information on session status, such as:

- Whether the user is in conversation mode
- Whether user output security exists for the terminal

Related information:

-  DFS3650I (Messages and Codes)

ETO terminal-LTERM relationship

The system programmer is responsible for the relationship between a terminal that is in session with IMS and a particular user's LTERMs.

For a single-component terminal (such as a SLU 2), IMS creates a single default LTERM name that is the same as the user ID that was supplied during signon. Using an exit routine or a user descriptor, you can give the LTERM a different name.

For a multi-component terminal (such as a SLU P), you need to ensure that sufficient LTERMs are created in order to use that terminal. You can use exit routines, installation-created descriptors, or node user descriptors to ensure that a sufficient number of LTERMs is created. Otherwise, IMS creates a single LTERM that is allocated to the first component of the terminal. Regarding node user descriptors, the user needs to use the node user descriptor to establish the correct relationship between the LTERM and the node, either explicitly or in the Signon exit routine (DFSSGNX0).

It is possible for an application program to associate a specific LTERM name with a specific terminal. The system programmer must ensure that names of alternate PCBs are consistent with LTERM names defined in:

- User descriptors
- Signon exit routine
- Insert exit routine
- Recovery requirements for ISC, SLU-P, and Finance sessions

How IMS determines which queues to allocate

IMS uses the following method to determine which LTERMs to allocate.

1. If the user control blocks already exist within IMS, these control blocks are reallocated, and the Signon exit routine is called, if appropriate.
2. If the control blocks do not exist and a user descriptor is specified at signon, IMS looks for the specified descriptor.
 - If the descriptor specified is not the user ID, node name, or DFSUSER descriptor, IMS sends error message DFS3649A (with return code 148).
 - If the descriptor specified is a valid user descriptor for the user that is signing on, IMS builds a table of available descriptors and calls the Signon exit routine, if appropriate.
 - If the Signon exit routine exists and RC=0, IMS continues normal signon processing. If the return code does not equal 0, the signon is rejected.
3. If no user descriptor is specified at signon and no user descriptor is specified in the Signon exit routine, IMS chooses a valid descriptor from the following (in order):
 - a. User ID
 - b. Node name
 - c. DFSUSER

Related reference:

 Signon exit routine (DFSSGNX0) (Exit Routines)

Setting special processing modes

After user signon, you can set the following processing modes by using IMS commands.

Exclusive mode

/EXCLUSIVE command

Preset destination mode

/SET command

MFS test mode

/TEST MFS command

Test mode

/TEST command

These special processing modes, except for the test mode and the preset destination mode, are retained for the user after signing off and are reestablished with the next terminal on which the user signs on.

When you issue the following commands, IMS creates the required control blocks for a terminal or user, or it maintains a user's status:

/EXC USER

Places a user structure in exclusive mode.

/STOP NODE or /STOP USER

Stops a node or user structure.

/TEST MFS USER

Places a user structure in MFSTEST mode. When a user signs on, IMS places the terminal in MFSTEST mode, if the terminal supports MFS.

/TRACE NODE

Traces a logged-on node.

Control blocks are not immediately deleted when special status is removed. They are deleted at the next checkpoint if they meet all deletion requirements. If an outstanding status exists, they are not eligible for deletion.

To reset terminal status and make the control blocks eligible for deletion at the next simple checkpoint, use the following commands:

/END Clears exclusive and test modes.

/RESET

Removes the preset destination.

/RSTART

Starts stopped resources.

/START

Starts stopped resources. For other status reset by the /START command, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

/TRACE

Sets trace off.

Use the preceding set of commands if the control blocks exist solely for status retention.

You can use the Signoff exit routine, DFSSGFX0, to reset these states.

Related concepts:

"Improving performance by deleting ETO control blocks" on page 128

Printers with ETO

Two methods of implementing printer support exist in an ETO environment: direct printing and associated printing.

How you implement these two methods depends on how your application programs identify printer LTERM names.

Direct printing

Direct printing is a printing technique by which application programs insert messages to the VTAM LU name.

The dynamic LTERM and the user and terminal resources are all named after the VTAM LU name. Many application programs can queue data to the same printer LTERM, but this can create data interleaving problems.

Associated printing

Associated printing is a technique for directing application program printer output to a specific printer node name.

For associated printing, application programs insert messages to the queue that is associated with the screen-user queue name. User printing requests cause application programs to queue data to different printer LTERMs. This avoids data interleaving problems.

Printer signoff and signon are required to change the user queue. Overhead can be high if the printer has many users.

Associated printers are logged on automatically when their LTERMs have queued messages, usually during an IMS restart or during the creation of an LTERM.

Recommendation: Carefully plan how your printers are to be shared among application programs.

By implementing exit routines and application programs, the terminal operator can provide the destination during logon or signon.

Identifying printer node names

You can identify printer node names in one of two ways.

The ways that you can identify printer node names include:

- As logon user data, you can include one or more printer node names when a user establishes a session.
- Your installation can modify the MFS format for the DFS3649A greeting message in order to allow the user to supply the printer node names at signon time. During the signon, however, the Signon exit routine must be able to detect the printer node names as user data.

Coding the Signon exit routine for associated printing

To use associated printing, code the Signon exit routine (DFSSGNX0).

Code the Signon exit routine to do each of the following steps:

- Identify the printer LU name and user name for each screen user.
- Determine printer node names from the input user data.
- Name user-related LTERM structures to service the selected printers.
- Enable application programs to determine (using the user ID) the queues associated with a particular user so that the programs can insert to the correct message queue (alternate PCB).
- Pass the printer node names to IMS as associated print parameters.
- Determine the user name that is allocated to each printer when users supply printer node names. You can use either an algorithm or a table to make this

determination. The exit routine should pass these user names to IMS, which then creates the necessary user control blocks.

A unique name should exist for the user ID that is signed on and for each printer-related user that is required.

Example: To indicate the name of a user's printer identification, add a "P" to the end of the user ID. If the user ID is AAA, the name of this user's printer identification is AAAP.

- Specify one value for each of up to four printers (the number of simple checkpoints before the user structure is deleted).

The application program can use the same method as the Signon exit routine in order to determine the printer LTERM name from the input terminal user ID and queue output data as required through an alternate PCB. When output data is queued, IMS allocates the user structure to the correct printer and delivers the output.

After associated printer LTERMs are allocated and then emptied, the queues are deallocated from the terminal.

Related reference:

 Signon exit routine (DFSSGNX0) (Exit Routines)

Defining your printers

Printers are usually output-only devices; however, they can be implemented so that they send input.

You need to decide whether to have single-user or multiple-user structures share the same printer. If your application programs are sensitive to queue names, you might be limited to one or the other approach.

Single-user structure

Using a single-user structure that represents a printer is the simplest method for defining a printer. In this case, messages sent to the printer are printed in the sequence in which they originate. Interleaving of output can occur. Multi-segment messages always have all segments printed in contiguous sequence. Multiple messages might behave differently, depending on their method of origination.

Multiple-user structures

Multiple users are supported for printers. All messages for a single user are printed, and the next user is selected for printing only when the current user has no more output. Although interleaving of messages for the same user occurs as it does for single-user structures, the switching of users can also be equivalent to message interleaving from an application standpoint. Application and operational awareness of the way printers are shared is important.

The challenges of sharing printers are not unique to dynamic terminals, and in fact are the same as for static terminals.

IMS prints a separator page whenever the next message to be printed is from a different user. You can control the contents of this separator page by using an exit routine to change its content; however, this page is always printed, even if blank. Messages from the same user are not separated by a separator page. Using the NOTERM option can prevent the DFS3650 separator message from being used.

Sharing printers using ETO

Several users can share printers by using the same output terminal.

Two methods exist for using the same output terminal.

- Users can define a single-user descriptor that contains all the LTERM names that are used to deliver data to an output device (node).
- Users can define autologon parameters for a node.

Before deciding which of these methods to use, ask these questions:

- Are the LU-to-LTERM relationships generally unchanging?
- Is interleaving at the message level acceptable?
- Is delaying one user's output acceptable while another user's output is being printed?
- Is continuous autosignoff and autologon processing overhead too excessive for a particular terminal because of the minimal message delivery rate of each LU user?

The answers to these questions can help to determine your method of implementation:

- If the answer to all of the questions is “yes”, consider creating a single multi-LTERM user for the printer.
- If any of the answers is “no”, consider dynamic allocation of multiple-user LTERMs, using autologon.
- You can implement a combination of these two options.

Operator commands

This topic describes using the /OPNDST and /ASSIGN commands as they are used with ETO.

/OPNDST

The VTAM mode table that is used when the terminal first logs on determines the device characteristics. However, if the first reference to a terminal is through an /OPNDST command, the MODETABLE operand of that command determines the device type.

Omitting the MODETABLE operand on the command causes it to default, which might not be desirable. After the terminal is in session with IMS, the rules for block deletion apply. The device type remains set until the block is deleted; if the terminal is closed and then reopened, IMS uses the existing block, if available. If a block is deleted, that block is rebuilt during the next terminal logon. As a result, the /OPNDST command can have different results, depending on two things:

- Whether the block is still available (and has not been deleted).
- Whether the exit routine or descriptor has been changed since the previous initialization.

/ASSIGN

You can use the /ASSIGN command to move queues from one terminal to another. You cannot use it to reassign a static terminal to a dynamic (ETO) terminal, or to assign a dynamic terminal to a static terminal.

Related reference:

 /ASSIGN command (Commands)

 /OPNDST command (Commands)

System definition parameters for ETO

This topic describes the system definition parameters you can use with ETO.

Setting DEADQ status time with the DLQT parameter

User control block structures are normally created in several situations.

The situations in which user control block structures are created include:

- When a terminal is logged on and a user signs on.
- When the AO exit routine (DFS AE0U0) inserts a message to an LTERM or transaction.
- When an asynchronous transaction output message is sent, or a terminal message switch or /BROADCAST LTERM command is issued.

Messages might be sent to destinations that are unknown, no longer valid, or nonexistent. IMS sets a status of dead-letter queue (DEADQ) when an ETO user control block structure or its associated message queues have not been accessed within the time limit that is set by the DLQT execution parameter. (For DLQT, although valid values are 1-365 days, a value of 1 is not usually recommended. It can result in many premature and misleading DEADQ status settings during the first checkpoint.) The DEADQ status can be set regardless of whether messages have been queued for the user or whether the user is still allocated to a terminal. The DEADQ status is set during IMS checkpoints, and the MTO is notified with message DFS3643.

If the user has messages queued, remove the DEADQ status by signing on the user or by issuing the /DEQUEUE or /ASSIGN command.

User control block structures without queued messages can result in a DEADQ status. User control blocks are not deleted if a special status is pending. The status might have been set during the prior signon (such as response or conversation mode) or as a result of a command (such as /STOP or /EXCLUSIVE). If the control block remains unused for longer than the time specified for the DLQT execution parameter, IMS assigns the control block a DEADQ status, and the MTO is notified with message DFS3643.

If the user control block has DEADQ status, the status is removed when the user signs on or during the next checkpoint after all messages are dequeued and recoverable status conditions are removed using appropriate commands. In the latter case, the control blocks for the user and associated message queues are also deleted.

User control block structures that have DEADQ status might or might not be allocated to a terminal. In the case of LU type-6, SLU-P, and Finance terminals, the user structure can be allocated to the terminal with no active session. Logoff or other session termination can leave these terminals pending message recovery (SNA STSN). The user remains allocated to the terminal, and messages might or

might not be queued. (This combination is not possible for other VTAM terminal types, because a deallocation of the terminal and user ID is forced at logoff and signoff.)

If the user control block structure is allocated to an LU type-6, SLU-P, or Finance terminal, remove the DEADQ status by logging on and signing on the user, or by using the /DEQUEUE command. For LU type-6 (ISC) terminals use the /DEQUEUE command, or force the LU type-6 session to cold start. A forced-session cold start is a /STO NODE, /ASSIGN (USER TO VTAMPOOL), /STA NODE sequence that is valid only for LU 6. Forced cold start is not possible for SLU-P or Finance terminals. Clearing the allocation of an ETO user to a SLU-P or Finance terminal requires an IMS cold start.

In a shared-queues environment, you can use the /DISPLAY QCNT MSGAGE command to find the age of a queue.

Autosignoff (ASOT)

Autosignoff deallocates users from an idle session. If no activity occurs on a session within an allotted time, users are automatically signed off and must sign on again in order to use the session.

The system programmer sets the autosignoff time using the ASOT parameter in the DFSPByyy member. The time value on the ASOT EXEC parameter does not apply to 3600, SLU P, or ISC devices.

If more than one allotted time value exists, IMS uses the following criteria to determine which allotted time value to use:

- If the valid ASOT value is specified in the user descriptor, this allotted time value is used.
- If the user descriptor does not specify an allotted time value, the allotted time value from the logon descriptor is used.
- If the logon descriptor does not specify an allotted time value, the time value from the DFSPByyy member is used.
- If the DFSPByyy member does not specify an allotted time value, the default value of 1440 is used.
- If the value on the DFSPByyy member is not valid, the default value of 10 is used.

The Logon exit routine (DFSLGNX0) can override the ASOT and ALOT values during logon. The Signon exit routine (DFSSGNX0) can override the ASOT value during signon, even when the control block structure exists.

The values for the allotted time specified on the ASOT parameter in the DFSPByyy member are:

- ASOT = 0
 - The user is signed-off immediately when no output is available to be sent. This specification is normally used with Autologon terminals for signoff immediately when:
 - No IMS input or output message is available
 - After the last available output message completes

- This specification is not recommended for interactive terminals such as 3270 or SLU2 terminals. These terminals sessions normally return a PA key to continue following signon. Idle time results in immediate signoff, not waiting for terminal input.

The value on the DFSPByyy member is not used for these device types.

- ASOT = (10 - 1439)

The user is signed off after the allotted number of minutes has elapsed without terminal activity.

- ASOT = 1440

The user is never automatically signed off. This is equivalent to not having autosignoff. The system default value for SLU-P, 3600/Finance, and ISC terminals is 1440.

After autosignoff completes, IMS attempts to locate a user that has the same node name that is waiting for autologon. If IMS finds another user with output waiting, the user is allocated to the terminal, and the queues are drained.

Autologoff (ALOT)

Autologoff can terminate a session with IMS for a terminal that has been signed off for an allotted period. After an allotted time, the terminal is automatically logged off.

The system programmer sets this time on any of the following:

- On the ALOT (auto logoff time) parameter in the DFSPByyy member
- On the logon descriptor that is used to create the session control blocks
- On the EXEC parameter at initialization

If more than one allotted time value exists, the following criteria are used to determine which allotted time value to use:

- If the ALOT value is specified on the logon descriptor, this allotted time value is used.
- If the logon descriptor does not specify a valid allotted time value, the time value specified in the DFSPByyy member is used.
- If the DFSPByyy member does not specify an allotted time value, the default value of 1440 is used.
- If the value on the DFSPByyy member is not valid, the default value of 10 is used.

You can specify ALOT=1440 on the logon descriptor for a 3600/Finance, SLU-P, or ISC terminal in order to specify that no autologoff is desired. The system default value for these devices is 1440. Logon descriptors created for ISC terminals should have ALOT=1440 specified to show that autologoff should not occur. The Logon exit routine (DFSLGNX0) can override the ALOT value during logon, even when the control block structure exists.

The values for the allotted time specified on the ALOT parameter in the DFSPByyy member are:

ALOT = 0

The terminal is logged-off immediately when no signon is in effect. This specification is normally used in terminal sessions when the user is signed-on automatically during the logon process. During autologon, signon data can be provided in one of the following ways:

- Signon data supplied by the IMS /OPNDST command
- Signon data supplied by logon user data (BIND)
- Signon data supplied by logon exit (DFSLGNX0)

There are two modes of operation for using ALOT=0, either of which can be set using the DFSINTX0 User Initialization Exit parameter list.

In default mode, when signon errors are encountered, the session is automatically signed off and then logged off; no message is sent. If you do not supply the DFSINTX0 exit, or you supply the exit and indicate default mode for ALOT=0, then signon data must be supplied during the logon process. All of the following error conditions result in automatic logoff:

1. A non-signon, or errors detected during signon or input processing, result in immediate logoff.
2. /SIGNOFF results in immediate logoff.
3. /SIGNON signs off the current user and signs on a new user. However, errors encountered during the signon process, such as detection of an incorrect or expired password, result in immediate logoff.

Restriction: Default mode should not be used for interactive terminal sessions that require a response to the DFS3649 message; these sessions will not wait for input signon and will logoff immediately.

In alternate mode, when signon errors are encountered, the session is automatically signed off, a message is sent and the session is logged off. Signon data can be supplied but is not required. All of the following error conditions result in automatic logoff:

1. A non-signon error detected during input processing results in immediate logoff.
2. No signon data has been provided by the logon serrated (BIND) or the Logon Exit (DFSLGNX0).
3. A /SIGNOFF, or errors resulting from a /SIGNON, cause message DFS3649(A) (Signon Required) to be sent, and a fixed ten-minute timer set to wait for a new signon. If no signon occurs during that interval, then the session is logged off.

ALOT = (10 - 1439)

The session is terminated after the allotted number of minutes has elapsed without a signed-on user.

ALOT = 1440

The session is never automatically terminated. This is equivalent to not having autologoff.

Autosignoff and autologoff timer

The VTAM I/O Timeout Facility (if active) detects users or sessions that should be automatically terminated.

A timer pops at intervals of one minute if the VTAM I/O Timeout Facility is active, or five minutes if the Timeout Facility is not active. The timer starts a routine that determines which resources are due for autosignoff or autologoff. The specified timeout value is the minimum value for which a user or session is automatically terminated. Termination actually occurs the next time the timer pops after the specified timeout value.

Autologon

Instead of using the SHARE option on the TERMINAL macro to request that IMS automatically initiate a terminal session when output is available, ETO offers autologon support for ETO terminals and users. You specify autologon parameters when defining the user to IMS.

Definition: *Autologon* allows IMS to log on and sign on your terminal automatically. If you specify the autologon option for a user, the queuing of data to any of the user queues causes IMS to establish a session. You can specify autologon using:

- AUTLGN= parameter on the user descriptor
- Destination Creation exit routine (DFSINSX0)
- Signon exit routine (DFSSGNX0) for associated printers
- /CHANGE command with the AUTOLOGON keyword

Autologon includes both automatic logon and automatic signon. At restart, IMS attempts to start sessions with those terminals that are defined with autologon and that have queued data waiting. After the session is established, IMS automatically signs on the terminal. If a queue of waiting users exists and the allocated queues are drained, the autologon user is signed off, regardless of an existing ASOT value.

IMS manages a serial queue of waiting users if more than one user is contending for the same autologon terminal. After the autologon user is signed off, the next autologon user for the same terminal is automatically signed on. When all autologon users have signed off, the terminal is free to begin its ALOT cycle in order to terminate the session.

Autologon replaces the TERMINAL macro OPTIONS=SHARE for static terminals. OPTIONS=NOASR (no automatic session restart) on the logon descriptor is ignored for autologon printers. IMS always assumes OPTIONS=ASR for autologon printers.

Autologon is normally specified for output-only terminals. Autologon and occasional users generally do not share the same terminal session, but sharing terminal sessions is possible for interactive terminals. Interactive users on terminals that are subject to autologon must supply user signon data with the session initiation request (logon) in order to avoid contention with autologon output. Occasional terminal users can use autologon in order to have output delivered to default terminals when the output becomes available after signoff. If a terminal is stopped, the /START NODE command does not start a session. The /OPNDST NODE USER command can be used to restart the session.

Assigning output

The following topics discuss managing asynchronous output and output destinations.

Asynchronous output

Asynchronous output can easily be sent to an invalid destination by a simple typographical error.

This is because ETO provides IMS the flexibility to create user structures for any authorized user that is signed on. IMS automatically creates user structures for message switches and for the application insert call (ISRT) process when the LTERM cannot be found.

With ETO, all destinations are valid unless they are rejected by exit routines. Therefore, you can mistakenly create queues for users if you make typographical errors when entering any of the following:

- Alternate PCBs
- Message switches
- /BROADCAST commands
- MSC output

The LTERMs used in the previous situations are known as *dead-letter queues*. IMS provides commands for the MTO to monitor these queues and dispose of them.

The two types of asynchronous output destinations are valid and invalid.

A *valid destination* is one intended to receive output. An *invalid destination* is one that is not intended to receive output (for example, a misspelled destination).

Asynchronous output to a valid destination

You can send asynchronous output (such as inserts, broadcasts, and message switches) only after enabling ETO and defining a valid ETO descriptor and a valid destination LTERM name.

ETO is enabled by specifying ETO=Y in the IMS or DCC startup procedure.

The ETO descriptor is used for creating a user structure.

If control blocks exist for a previously created user structure and LTERM, the control blocks are reused.

Asynchronous output to an invalid destination

IMS refers to data that cannot be delivered as “dead letter”.

Data cannot be delivered in each of the following situations:

- No autologon destination is available for queued output.
- The user ID to which the data is associated is not a valid user ID.
- The user signon is always rejected by an installation Signon exit routine.
- An invalid destination is specified on the input or output message (for example, resulting from a typing error).

You can specify a DLQT value on the EXEC parameter at initialization in order to automatically notify the MTO when LTERM queues exceed this value and have not dequeued data or removed the status. You can use each of the following commands against dead-letter queues:

- Use the /DISPLAY USER DEADQ or /DISPLAY STATUS USER command to identify users whose LTERM queues are older than the dead-letter-queue time (DLQT), and who have not dequeued data or removed the status.
- Use the /ASSIGN command to reassign dead-letter queues to other dynamic users so they can review queued data.
- Use the /DEQUEUE command to purge data on the dead-letter queues.

In a shared-queues environment, you can use the `/DISPLAY QCNT MSGAGE` command to determine the messages that are considered to be dead-letter queues.

Related reference:

 [IMS commands \(Commands\)](#)

Delivering output messages to non-originating terminals

IMS sends your output to the terminal at which you are signed on.

Using ETO, you can receive your output messages at a different terminal than the one from which you entered the input. However, the input and output messages are formatted subject to the MFS specifications defined for both the terminal and messages, as follows:

- MFS formatting is mandatory for 3270R (non-SNA) and SLU-2 terminals, except when you use MFS bypass. MFS formatting is optional for all other VTAM terminals. Use MFS on a message-by-message basis, based on MID and MOD control block availability. MFS paging support is not available for SLU-1 or NTO terminals.
- You must define MID and MOD control blocks by using device statements that generate appropriate DIF and DOF control blocks. This allows you to map the message to and from a specific terminal type at the time that the message is sent to or received from the terminal. Default mapping occurs when the appropriate MFS blocks are not available. IMS issues error message DFS057 when the format blocks cannot be found.

If you plan to deliver output messages to non-originating terminals, you must develop or expand MFS formats, procedures, and restrictions. This allows users to move freely between terminals.

Inadvertent output data streams

Using ETO, dynamic terminal users can move freely between terminals; limited only by installation constraints. It is possible to erroneously send terminal-specific data to the wrong terminal type by mistyping the IMS `/ASSIGN` command.

When data is sent to the wrong terminal, the data, when delivered, has errors or is not recognizable. Be sure to create MFS definitions for all terminal types for which users can log on.

Signing off

Signing off of an ETO terminal ends the identification of a user to IMS, and in most cases disconnects the user structure from the terminal structure and deletes the user structure.

When a user signs off from an ETO VTAM terminal, IMS calls the Signoff exit routine (DFSSGFX0).

Recommendation: If you have provided a Signon exit routine (DFSSGNX0) that maintains system information, provide a Signoff exit routine (DFSSGFX0) also, to complement that processing.

Logging off

When a user logs off from a VTAM terminal when ETO is used, IMS calls the Logoff exit routine (DFSLGFX0).

Recommendation: If you have provided a Logon exit routine (DFSLGNX0) that maintains system information, provide a Logoff exit routine (DFSLGFX0) also, to complement that processing. Ensure that the Logoff exit routine handles all non-MSA and non-LU 6.2 terminals with which IMS communicates.

Using the Logoff exit routine, you might want to maintain a count of the terminals that are logged on.

Improving performance by deleting ETO control blocks

With ETO, IMS can dynamically delete control blocks. Dynamically deleting control blocks reduces storage usage and can improve performance.

If special terminal processing options (TRACE and STOPPED) are reset, IMS deletes session control blocks if one of the following occurs:

- No user is signed on, and a checkpoint occurs.
- The session is terminated normally or abnormally by either an MTO command or by an autologoff timeout.

If the node is a 3600/Finance or SLU P terminal, message resynchronization is necessary. The control blocks are not deleted following warm-session termination, but will be deleted following a /CHANGE NODE COLDSESS command. For ISC terminals, the control blocks are deleted after a cold-session termination. The control blocks remain, however, after an ISC warm-session termination.

If the user resets special processing options, such as those that exist after issuing a /SET, /TEST MFS, or /EXCLUSIVE command, and issues the /SIGN OFF command (or is automatically signed off), IMS deletes the user control blocks if:

- No messages are queued to any LTERMs related to this user.
- The user is not in conversation, Fast Path mode, or full-function response mode.

If the preceding conditions exist, dynamically created user control blocks are deleted. If the preceding conditions do not exist, the user control blocks can continue to exist until the conditions exist or until an IMS cold start occurs. After special processing options are reset and if all other criteria for deletion exist, the control blocks are deleted at the next checkpoint.

Important: User control blocks can be saved across session and IMS restarts by using the /CHANGE or the /ASSIGN commands with the SAVE keyword. These user control blocks are then retained until the commands are reentered with the NOSAVE keyword.

For terminals or users in full-function response mode, IMS does not delete user control blocks after a terminal logoff or a user signoff if both SRMDEF=LOCAL and RCVYRESP=YES, because in this case the full-function response mode is recoverable.

Note: In an IMSplex, if the status recovery mode is GLOBAL or NONE, the local control blocks are deleted immediately after logoff or signoff.

IDC0 Trace facility

You can use the IDC0 Trace facility to diagnose logon and logoff errors.

This facility provides information that the IMS message DFS3672I cannot provide. To use the facility, enter: /TRACE SET ON TABLE IDC0. The facility traces the following events:

- Errors that occur in the IMS VTAM exit routines (within module DFSCNXA0). These errors are also identified in a DFS3672I message, regardless of whether the IDC0 Trace is in effect.
- Errors that occur when attempting to log onto a nonexistent VTAM node (such as entering a /OPNDST command for a nonexistent terminal). IMS issues associated messages DFS2061I or DFS2062I.
- Synchronization anomalies that occur between the time that an IMS VTAM exit routine completes processing and the time that the request is accepted by normal IMS processing. The result is a X'6701' log record identified with a VTPO string.

Related reference:

 /TRACE TABLE command (Commands)

 Format of the 6701 log record with VTPO identifier (Diagnosis)

 IDC0 trace table entries (Diagnosis)

ETO and LU 6.1 (ISC) terminals

For LU 6.1 (ISC) terminals, IMS supports parallel sessions to the same node name. In this case, a separate structure is built for each session. However, each session and its associated structure operate independently, as a separate terminal.

ISC supports an SNA-defined user-data area within the BIND. When establishing a session for ISC, each half-session partner is identified through an appropriate session qualifier that is included as user data with the logon. These two qualifiers cannot be specified using the DFSLGNX0 exit routine. One belongs to each half session. IMS uses the session qualifier of the current half session as a user structure. This user structure is used to allocate an associated set of LTERM queues and to automatically provide a RACF signon, if required. The other half-session qualifier is saved with the IMS user structure. Both qualifiers are used for session-restart requests and SNA STSN message resynchronization after session failures.

Restriction: The SNA-predefined format for user data does not support some of the parameters and options of the non-ISC end-user format:

- The RACF password and group name are not supported. IMS supports RACF signon for ISC with PASSCHK=NO during user structure allocation as part of session initiation.
- The LOGOND and USERD are not supported. IMS uses defaults, unless specified on the Logon and Signon exit routines.
- When autologon is generated for ISC terminals, the AUTLDESC keyword in the user descriptor is ignored, and the LOGOND keyword in the user data is omitted.

Related tasks:

“Using default CINIT or BIND user data formats” on page 112

Related reference:

ETO and STSN terminals

This topic provides information on administering ETO for STSN terminals.

SNA STSN terminal considerations

ETO terminals that use the SNA STSN function (Finance, SLU-P, and ISC) must provide a user queue name during logon, because this queue is used to resolve the sequence number exchange that is part of the IMS connection process for this type of terminal.

The ways to provide the user name include:

- CINIT data sent by the terminal
- CINIT data provided using a separate terminal host product, such as VTAM unsolicited system services
- A user Logon exit routine (DFSLGNX0), except for ISC
- An IMS command (/OPNDST can specify user name)
- Autologon parameters supplied by user descriptors or by the Destination Creation exit routine

The Logon exit routine is the last opportunity to provide the user name. If no user name is provided for ETO STSN terminals, the logon is rejected with an error message. This differs slightly from other terminal types, which allow user signon after the logon has occurred. A signon is required before being able to use the terminal for IMS activity (transactions or commands), so the difference is only in requiring the signon data earlier for STSN terminals. The /SIGN command is supported for STSN terminals.

ETO and 3600/Finance and SLU P

You can sign on to static system-defined 3600/Finance and SLU P terminals in one of two ways: using the /SIGN command or using logon user data.

You can change the signon identification by using another IMS /SIGN command at any time. LTERMs are assigned to the terminal during system definition and are not affected by the signon process or by the SNA STSN message recovery process. The signon simply provides user access and input authorization.

IMS supports 3600/Finance and SLU-P terminals as dynamic terminals. They can use autologon or signon data with the logon request in order to dynamically allocate user structures. Signon data must be provided at session initiation for ETO 3600/Finance and SLU-P terminals. Signon data can be supplied in the Logon exit routine (DFSLGNX0) at the cold start of a session. The LTERMs and user IDs allocated at the cold start session are retained across sessions and IMS outages because of the VTAM STSN message resynchronization requirements. This requires that the same user signon data be used for subsequent warm-start sessions to both reverify the user and to allow message resynchronization.

When dynamic XRF Finance and SLU-P terminals are defined as XRF class-2, automatic re-signon and logon occur at takeover time.

/SIGN support for ETO STSN devices: ISC, Finance, and SLU-P

For ETO STSN devices, user data is required at the time that the session is allocated to create the user structure.

After the user structures are created and allocated to the terminal, /SIGN commands are accepted from ETO STSN terminals.

When you issue the /SIGN command from an ETO STSN terminal, IMS initiates a complete signon process to create the security profile associated with the session for the new user.

When the user signs off, the user's security profile is deleted, leaving the session without any security. RACF rejects all access to RACF-protected resources. The DFS3662 message is displayed if the failing resource is a command. The DFS2469 message is displayed if the failing resource is a transaction. When a user signs on to an STSN device, the same user structure allocated during session allocation is used for the new user. IMS updates the security profile of the session and stores the signon information.

The user ID of a user that is signing on to an ETO Finance, SLU-P, or ISC device with a /SIGN command is a different name from that of the user structure name. Such users do not require suffixing by the Signon exit routine (DFSSGNX0) in order to support multiple signons, because the user structure for these devices was already created during the session initiation.

Restriction: Because the user structure allocated to a Finance, SLU-P, or ISC device cannot be changed, most of the options available to the Signon exit routine are not supported, and the work areas are not passed to the Signon exit routine.

The DFS3650 message is displayed after a signon, and the DFS058 message is displayed after signoff. The user field in the DFS3650 message reflects the user structure's name rather than the RACF user ID. This is the same as for any other ETO terminal.

Restriction: Issuing the /SIGN command from STSN output-only devices is not allowed.

Related concepts:

 Specifying IMS execution parameters (System Definition)

 Exit routines (Exit Routines)

Conversation mode and response mode with ETO

ETO user structures are distinct from static terminal structures, because terminal status is maintained for each user rather than for each terminal.

An IMS conversation at a static terminal is distinct from an IMS conversation at a dynamic terminal, even for the same user name. The conversation at a static terminal can be held, and must be released (resumed) at the static terminal. It cannot be moved to a dynamic terminal.

With ETO, an IMS conversation can be resumed at the same terminal or another dynamic terminal. Because the conversation is an attribute of the user structure, it normally follows the user to a different terminal when the user signs on to IMS.

The following of the conversation attribute can be a cause of confusion, especially if the user is not aware of which terminals are static and which are dynamic. Also, if an exit routine selects different user structures for signons to different physical terminals, confusion can occur.

Resume full-function and Fast Path response mode at the same static terminal. For ETO, resume full-function and Fast Path response mode from the same or another dynamic terminal. This ability to resume the Fast Path response mode from any dynamic terminal is similar to the situation with conversations described in the previous paragraph, and the same considerations apply.

Conversation mode

For ETO, conversations are associated with the user—not the terminal that initiates the conversation. Conversations are also associated with the terminal, but only while the user is signed on. By signing off, the user can continue a conversation on a different terminal. This flexibility requires the installation to address output formatting problems.

Users in conversations that are not in response mode can sign off. Regardless of response mode, users in conversation can be automatically signed off through autosignoff or by using an MTO command. Any form of signoff leaves the terminal available for the next user. The conversation mode status follows the user to the next terminal or, with the Resource Manager and global status recovery mode (SRM=GLOBAL), on a different IMS in the IMSplex.

Response mode

Response mode is defined on the TERMINAL macro for static terminals, on the ETO user descriptor for dynamic (ETO) users, and on the TRANSACT macro for transactions. Also, response mode is either full function or Fast Path. You can also use response mode with conversation mode, if you are not running Fast Path. Response mode is primarily associated with the user and the transaction, rather than with the dynamic terminal.

When a user is in response mode, the keyboard or input response is locked until the output reply is available. During this time, it is not possible to enter input at the terminal. Normal signoff and logoff commands are not allowed. However, these functions can occur automatically during abnormal session termination. This can happen in one of three ways:

- VTAM can detect an error and end abnormally (abend).
- The MTO can issue the IMS /CLSDST or /STOP command.
- IMS can autosignoff after the specified autosignoff interval.

Regardless of how signoff occurs, if RCVYRESP=YES or RCVYFP=YES, the response mode is retained for the user that has been automatically signed off. The user's response mode operation is re-established with the next terminal on which the user signs onto and remains until the response-mode output reply is available.

The master terminal operator can reset the Fast Path response mode of an ETO dynamic user before a response is returned by issuing the /STOP USER and /START USER commands in sequence from the master terminal. The master terminal operator can also reset the Fast Path input response mode of a static node by issuing the /STOP NODE, /START NODE commands in sequence from the master terminal.

Related concepts:

“Delivering output messages to non-originating terminals” on page 127

Part 4. External subsystem attach facilities

IMS provides several options for accessing external subsystems from an IMS system.

Chapter 7. DB2 Attach Facility

Java message processing programs (JMPs) and Java batch programs (JBPs) in IMS can access DB2 for z/OS data using the DB2 Resource Recovery Services Attach Facility, referred to here as the DB2 Attach Facility.

Each dependent region that is set up for this support builds its own RRSF thread to access DB2 for z/OS data. This thread enables the coordination of updates that the application program makes with the resources of both IMS and DB2 for z/OS resource managers. When IMS JMPs and JBPs use the DB2 Attach Facility to access DB2 for z/OS data, IMS is not the sync point coordinator of updates and commits, as it is with ESAF. With the DB2 Attach Facility, IMS is a participant, and z/OS Resource Recovery Services is the sync point coordinator.

Preparing your system to use the DB2 Attach Facility

To prepare your system to use the DB2 Attach Facility, you must perform two tasks.

To use the DB2 Attach Facility:

1. Add the attachment facility definition to the IMS PROCLIB data set. Use of the DB2 Attach Facility requires that a subsystem member (SSM) be defined in IMS.PROCLIB. If an SSM member does not already exist in IMS.PROCLIB, you must create one. The SSM contains an entry for the DB2 for z/OS system with which IMS and the application program communicate. All Java dependent regions in IMS access a single DB2 for z/OS system.
2. Make the DB2 for z/OS RESLIB available to the IMS JMP and IMS JBP regions. After defining the attachment facility, you must provide the Java regions in IMS with access to the DB2 for z/OS RESLIB. In the JCL for the JBP and JMP regions types, add the DB2 for z/OS library definition using the DFSDB2AF DD statement, which points to the DB2 for z/OS libraries that contain modules used by RRSF. The DB2 for z/OS libraries must be APF-authorized.

Related concepts:

 Accessing external subsystem data (System Definition)

Managing how your Java dependent regions access DB2 for z/OS

Java application programs running in IMS dependent regions can access DB2 for z/OS under syncpoint control of z/OS Resource Recovery Services if a DB2 Attach Facility definition is included when the IMS control region is started.

The initialization processing of the IMS control region prepares for access to DB2 for z/OS. When Java dependent regions are subsequently started, application programs in those regions can make direct calls to both DB2 for z/OS and IMS.

Initialization of the DB2 Attach Facility does not affect the execution of other types of dependent regions. The DB2 Attach Facility definition can be retained in the IMS.PROCLIB member, even if Java dependent regions are not used. If a DB2 Attach Facility definition exists in the IMS.PROCLIB member, and the DB2 for z/OS library is defined in the Java dependent region JCL, all Java dependent regions that start will build an access thread to DB2 for z/OS. If the DB2 Attach

Facility definition exists, but the Java dependent regions do not require access to DB2 for z/OS, you can prevent access threads from being built by stopping access to the DB2 for z/OS system. Use the /STO SUBSYS command, which stays in effect until a /STA SUBSYS command is subsequently issued.

Related reference:

 [/START SUBSYS command \(Commands\)](#)

 [/STOP SUBSYS command \(Commands\)](#)

Chapter 8. External Subsystem Attach Facility (ESAF)

The External Subsystem Attach Facility enables MPP, BMP, and IFP application programs to access databases managed by other subsystems in addition to DL/I databases.

To enable access to the data resources of an external subsystem (ESS) product from IMS applications, the ESS must provide functions necessary for it to attach to the IMS subsystem and to, jointly with IMS, coordinate data access. The IMS Attach Facility presents a programming interface to the external subsystem product. Certain steps are required to install ESAF, which are described in the following information. Other IMS publications also contain ESAF information and references are included, where applicable.

Multiple external subsystems can only be attached by an online IMS system. These subsystems can be of the same, or of different, product types. The installation defines the external subsystems to IMS. A given IMS dependent region can have access to all external subsystems defined to the IMS system, to just a subset, or to none according to installation specifications. An application program executing in a dependent region can access more than one different subsystem. The installation defines a unique token for each subsystem, which IMS uses in routing application calls for external resources. Application program access to more than one subsystem of the same type is supported by IMS, but might not be supported by the external subsystem.

The facility provides for synchronization of external subsystem data resources with IMS data resources. For synchronization processing, IMS is the recovery coordinator and is responsible for directing commit or abort actions on behalf of its application programs. External subsystems are participants in the process and commit or abort data updates by IMS applications according to direction given by IMS. When resources are to be committed, IMS polls the participants as to whether or not they are ready to commit before giving final commit (or abort) direction.

For IMS batch, IMS is allowed to attach only to one external subsystem. IMS expects this external subsystem to be the Recovery Coordinator. This external subsystem has no way of coordinating with any other external subsystem that IMS attaches to, so IMS is restricted to only one external subsystem attachment in batch.

The External Subsystem Attach Facility is different from the coordinator controller (CCTL) associated with DBCTL.

Note: JMP and JBP applications use the DB2 Recoverable Resource Manager Services Attach Facility (RRSAF) instead of ESAF (external subsystem attach facility) to attach DB2 for z/OS databases.

This topic contains General-use Programming Interface information.

Related tasks:

Chapter 7, “DB2 Attach Facility,” on page 137

 Accessing DB2 for z/OS databases from JMP or JBP applications (Application Programming)

What the external subsystem must provide

The External Subsystem must provide three things: the External Subsystem Attachment Package (ESAP), the External Subsystem Module Table (ESMT), and the Resource Translation Table (RTT).

External Subsystem Attachment Package (ESAP)

The IMS Attach Facility uses an exit routine interface. That is, to accomplish external subsystem access from dependent regions, IMS activates exit routines at certain processing points. These exit routines must be supplied by the external subsystem. The exit routine functions are prescribed by IMS; the external subsystem supplies its unique implementation. The exit routines must, in fact, provide the actual linkage to the external subsystem. IMS is not sensitive to the linkage mechanism used.

IMS loads external subsystem-supplied exit routine modules in the control region and in each dependent region that can access the external subsystem. The external subsystem can supply additional modules needed for attach exit routine processing; IMS loads these modules as well. The external subsystem modules provided for attach processing in the IMS regions make up what IMS calls the External Subsystem Attachment Package (ESAP).

External subsystem module table (ESMT)

The external subsystem must specify the modules that IMS is to load in an external subsystem module table (ESMT). The external subsystem creates the module table using macros provided by IMS and makes it available to the installation. The installation specifies the name of the ESMT to IMS by including it on the definition of the external subsystem to IMS.

Resource translation table (RTT)

IMS uses a PSB (program specification block) to define the DL/I resources required by an application program. For an MPP, the PSB name is the same as the application program name; for a BMP or IFP, the PSB name can be different. The external subsystem can use a name other than the PSB name or the IMS application program name for the entity it uses to define the external subsystem resources required by the application program. If the external subsystem uses a different name, the external subsystem can provide a resource translation table (RTT) to map either PSB names or application program names to its entity names.

The external subsystem creates the RTT and makes it available to the installation. The installation specifies the name of the RTT on the definition of the external subsystem to IMS. IMS loads the RTT when it loads the ESAP.

The external subsystem is responsible for doing the actual mapping. IMS does not access the RTT; it merely loads the table and makes its address available to the ESAP. IMS does not prescribe the format of the RTT.

Related concepts:

Chapter 9, "Creating the external subsystem module table," on page 149

How external subsystems are specified to IMS

In an IMS.PROCLIB member, define all external subsystems that are to be accessed by IMS applications. The EXEC statement of the control region points to this member with the SSM parameter.

For each external subsystem defined to IMS, specify in the IMS.PROCLIB member:

- The external subsystem type
- The z/OS name of the external subsystem
- The name of the external subsystem module table (ESMT) that specifies the modules in the external subsystem attachment package (ESAP)
- The language interface token (LIT) that IMS uses to route application calls to the external subsystem
- The name of a resource translation table (RTT) supplied by the external subsystem, if needed, to identify the external resources required by IMS application programs
- The command recognition character (CRC) that IMS uses to route operator commands to the external subsystem
- The region error option (REO) code indicating the action to be taken when application calls to the external subsystem cannot be processed. When a Resource Translation Table (RTT) is used, the OPTION value specified in the RTT overrides the REO option in the SSM member.

You must also supply the external subsystem-supplied tables (ESMT and RTT) in the appropriate load module library.

You have the option to supply external subsystem definitions for MPP, BMP, and IFP dependent regions. If the SSM EXEC parameter is not specified for these types of dependent regions, the region can access all subsystems defined to the control region. If the SSM EXEC parameter is specified, the dependent region can access only those subsystems defined in the identified PROCLIB member. (The subsystems also must have been defined to the control region.) Use a dummy PROCLIB member (one having no definitions) if the dependent region is not to have access to any external subsystem.

Note: JMP and JBP regions can access control-region-defined DB2 for z/OS subsystems that use the COORD=RRS parameter in the IMS.PROCLIB member. You cannot use the SSM parameter in JMP or JBP EXEC statements. In the DFSJMP and DFSJBP procedures, you must add a DD statement (DFSDB2AF) that points to the DB2 for z/OS libraries.

The following tasks must be performed to attach an external subsystem to IMS.

To attach an external subsystem to IMS:

1. Define external subsystems to IMS:
 - a. In the IMS procedure library (IMS.PROCLIB), add a member that contains the information about each external subsystem with which IMS communicates.
 - b. On the EXEC statement of the IMS control region, MPP region, IFP region, or BMP region, specify in the SSM parameter the PROCLIB member that you created in the previous step. If you are using a DB2 for z/OS group name to access DB2 for z/OS databases, you must specify the group name

- in the EXEC statement of the dependent region. A DB2 for z/OS group name cannot be specified in the EXEC statement for the IMS control region.
2. Define a language interface module if you want to use one other than the IMS-supplied one.
 3. In the IMS OPTIONS statement, specify whether you want tracing of the external subsystem link.
 4. For the external subsystem, provide the ESMT and optionally the RRT.
 5. Ensure that the external subsystem modules and databases used by IMS are in appropriate APF-authorized libraries.

Related concepts:

 Accessing external subsystem data (System Definition)

The basics of attach processing

An external subsystem is attached to an IMS subsystem by means of a connection established from the IMS control region to the external subsystem.

A connection is also established from each dependent region that accesses the external subsystem. IMS is responsible for initiating these connections.

Subsystem connections

The connection between an IMS application program and the external subsystem is called a *thread*. Application threads are two-way communication paths between IMS application programs and external subsystem resources.

An application program can have more than one thread since it can access more than one external subsystem in one execution. However, access to multiple subsystems of the same type (multiple instances of the same subsystem type) while supported by IMS, might not be supported by the external subsystem product.

Establishing connections

IMS uses an 'identify' process to establish a connection to the external subsystem.

IMS activates an Identify exit routine contained in the ESAP to identify the control region or dependent region TCB to the external subsystem. The external subsystem can then monitor IMS TCBs in order to respond to IMS abnormal terminations. A connection is established upon successful completion of the identify process, in other words, once the region has been successfully identified to the external subsystem.

IMS provides a notify message mechanism so that if the external subsystem has not been started when IMS attempts to connect the control region, the external subsystem, once started, can notify IMS to establish the connection. If the external subsystem makes use of the notify capability, the order in which IMS and the external subsystem are started is not important.

The connection from the control region is established first before any dependent region connections are established. If the control region connection has not been established when a dependent region is started, the dependent region does not identify itself to the external subsystem. IMS uses a hierarchical relationship between control region and dependent region connections to allow the control region to act as recovery coordinator for dependent regions. If a dependent region fails, the control region takes recovery actions on its behalf.

The external subsystem can optionally provide an Initialization exit routine. IMS activates the Initialization exit routine, if provided, during control region and dependent region initialization before the region identifies itself to the external subsystem. This exit routine allows the external subsystem the chance to do any initialization processing it requires before each connection being established.

IMS can establish the control region connection automatically during control region initialization, provided the external subsystem has been started. However, the connection can be delayed to a later time. If an Initialization exit routine is not provided, or if the exit routine returns the appropriate return code, the control region identify is not done automatically. In this case, the external subsystem can activate the Subsystem Startup Service provided by IMS when it wants the connection established. Or, IMS attempts to establish the connection when a dependent region is ready to identify itself.

IMS also establishes the control region connection in response to a /START SUBSYS operator command.

User authorization processing

After a dependent region connection has been established, a signon process is performed to inform the external subsystem of the user ID associated with the IMS transaction being processed by the region. IMS activates a Signon exit routine provided by the external subsystem for this purpose. This initial signon for the region must be successful in order for a thread to be created for the application.

Signon processing can occur again during application program execution (that is, after the thread has been created). The Signon exit routine is activated for each message processed by the application program. The initial signon performed after the dependent region identify is related to the first message processed by the application (first get unique call to the message queue). Each subsequent message processed causes the Signon exit routine to be activated again to pass the new user ID. In the case of multiple mode transactions, this means that multiple signons can occur without intervening commit processing.

The external subsystem supplies a Signoff exit routine which IMS activates before terminating the dependent region connection. In the case of multiple signons for an application, signoff processing does not precede a new signon for a new message. A new signon rather replaces the previous.

Application threads

When the application program issues its first call for data resources owned by an external subsystem, a thread is created to connect the application to the external subsystem.

IMS activates a Create Thread exit routine supplied in the ESAP to identify the application program to the external subsystem. The external subsystem is expected to prepare to receive data requests from the specific application program as necessary (that is, reserve resources, create a processing structure, and so on). When the application terminates, IMS activates a Terminate Thread exit routine to terminate the thread.

Terminating connections

A Terminate Identify exit routine must be provided in the ESAP. IMS activates this exit routine when a connection is to be terminated.

Termination of the control region connection can be initiated by IMS, by the external subsystem, or by operator command (/STOP SUBSYS). IMS terminates the connection when it is shutting down. The /STOP SUBSYS command causes the connection to be terminated and also puts it in stopped status. IMS does not allow the connection to be reestablished until a /START SUBSYS command has been processed.

The external subsystem can request that the control region connection be terminated in one of two ways. One way is by posting a termination ECB. IMS provides, on the Identify exit routine invocation, the address of an ECB that is expected to be used by the external subsystem when it is terminating. When the external subsystem posts the termination ECB, IMS, after allowing dependent region connections to quiesce, terminates the connection and also puts the connection in stopped status (as it does for the /STOP SUBSYS command). The second way that the control region connection can be terminated is by activating the IMS-supplied Subsystem Termination Service from an external subsystem exit routine.

After activating the Terminate Identify exit routine in the control region, IMS activates the Subsystem Termination exit routine supplied in the ESAP. This exit routine, which can be thought of as the reverse of the Initialization exit routine, might be used by the external subsystem to reset work areas or free storage, for example.

A dependent region connection is maintained for as long as the region is active unless IMS has been requested, either by the external subsystem or by an IMS /STOP SUBSYS command, to terminate the (control region) connection. In general, it is only when IMS has been requested to terminate the connection that the external subsystem Terminate Identify exit routine is driven for dependent regions. Thus, the exit routine is not necessarily activated when a dependent region terminates. This is true also for the Signoff exit routine. Terminate Identify exit routine invocation always follows Signoff exit routine invocation.

The external subsystem is expected to monitor, through the z/OS end-of-task exit routines, the IMS TCBs identified to it and to perform the necessary signoff and terminate identify processing when an identified TCB ends.

Terminate Thread exit routine invocation always precedes normal termination of the dependent region connection if the region had a thread to the external subsystem. Thus the Terminate Thread exit routine is activated before the Signoff and Terminate Identify exit routines are activated (if they are) or before the dependent region is terminated.

Since IMS does not allow dependent region connections to exist unless the control region has a connection, the Terminate Identify exit routine is not activated for the control region until after each dependent region has either terminated or had its Terminate Identify exit routine activated.

The Subsystem Termination exit routine is not activated for dependent regions.

Inquiry parameter processing

The INQ parameter is only checked when the IMS transaction issues a Create Thread exit routine.

The INQ parameter on subsequent transactions is not checked. Therefore, if any updates are to be done in a Fast Path region between the Create Thread exit routine and the Terminate Thread exit routine, the inquiry flag in the first transaction must be INQ=NO.

For example, if the first transaction that calls DB2 for z/OS from a given Fast Path region is only going to read the DB2 for z/OS data and not update it, the transaction will set the INQ=YES flag in the Create Thread Parameter list indicating that this first transaction and all subsequent transactions in that Fast Path region are treated as inquiry only transactions. If a subsequent transaction running under the same Fast Path region calls DB2 for z/OS for update, the thread to that Fast Path region will still be set to INQ=YES, even though the transaction is correctly defined as INQ=NO. This will result in an SQLCODE817 error.

Application call processing

After a thread from the application program to the external subsystem has been created, application calls for external data resources are passed to the Normal Call exit routine supplied in the ESAP.

The language interface bound with the application provides the language interface token (LIT) for the external subsystem when it activates IMS to process calls to the external subsystem. The installation specifies a unique LIT for each external subsystem it defines to IMS. IMS matches the LIT provided by the language interface stub with the LIT specified in the definition to route the call to the external subsystem.

Resource coordination

IMS, as recovery coordinator, directs commit processing for updates to external subsystem resources initiated by IMS application programs. IMS uses a two-phase commit process to synchronize resources across external subsystems. External subsystems are participants in the process.

In the first phase of the commit process for an application, IMS polls the participants for a vote as to whether or not they are prepared to commit the updates. In the second phase, IMS directs the participants to commit or to abort. If all participants voted 'yes' on the first phase, IMS directs them to commit on the second phase; otherwise, IMS directs them to abort.

When an external subsystem determines that its resources are associated with non-update transactions (for which commit processing is not necessary), the external subsystem can perform all commit processing during the first phase, eliminating the need for the second phase. In this case, the external subsystem returns to IMS from the Commit Prepare exit routine with return code X'C' indicating that the first phase successfully completed and the second phase is not required. IMS will not initiate the second phase of commit processing for this external subsystem.

IMS uses a 16-byte recovery token to identify a unit of work across one or more subsystems. The recovery token for a unit of work is initially passed on the Signon exit routine invocation.

When application updates are to be committed, IMS activates the Commit Prepare exit routine supplied by the external subsystem. The associated recovery token is passed on the invocation. The external subsystem indicates, by the return code from the exit routine, whether or not it is prepared to perform commit processing

for the recovery token. When an application is executing in a Distributed Syncopate environment (also known as a Protected Conversation environment) and requires a subsystem SIGNON, IMS obtains the XID token and places its address in the exit parameter list before calling the subsystem's SIGNON exit.

For the second phase of the commit process, if it is required, IMS can activate either of three external subsystem exit routines: the Commit Continue exit routine, the Abort Continue exit routine, or the Terminate Thread exit routine. When the application is not terminating and all participants are prepared to commit, IMS drives the Commit Continue exit routine. At completion of the commit process, the application will continue processing the current PSB on the existing thread. When the updates are to be aborted but the application is not terminating, or being terminated, the Abort Continue exit routine is activated. In this case, the application will continue processing under the same recovery token.

The external subsystem Terminate Thread exit routine must be able to process the second phase of commit. At application termination, IMS passes the recovery token and a commit option on the Terminate Thread exit routine invocation. The commit option indicates whether to commit or abort outstanding updates.

When IMS, the external subsystem, or an application program terminates abnormally, units of work that have not been committed or aborted are left outstanding. To resolve outstanding units of work, IMS activates the external subsystem Resolve Indoubt exit routine. IMS always activates the Resolve Indoubt exit routine at least once after establishing the control region connection. IMS activates the exit routine once for each outstanding recovery token indicating whether to commit or abort the unit of work. When there are no units of work to be resolved or when IMS has exhausted the list of outstanding recovery tokens, IMS activates the exit routine to inform the external subsystem of that fact. When IMS encounters an outstanding recovery token associated with z/OS Resource Recovery Services, IMS will delay the subsystem Resolve Indoubt exit call until RRS or the IMS user has indicated (ABORT or COMMIT) which action to take. When called for RRS Resolve Indoubt, it is the subsystem's responsibility to ensure that recovery tokens are resolved in their proper order.

IMS maintains outstanding recovery tokens across normal (warm) and emergency restarts of IMS, and reconnections of the subsystems. IMS permits a connection without all recovery tokens being resolved (that is, the Resolve Indoubt exit routine return code can indicate that the recovery action was not taken). IMS destroys outstanding recovery tokens when it is cold started.

The Resolve Indoubt exit routine is also used to coordinate resources in the event of abnormal termination of an application program. Following an application programabend, the exit routine is activated from the control region if the application had a thread connection to the external subsystem.

External subsystem command support

IMS provides a command, /SSR, which allows the IMS operator to send commands to the external subsystem.

To receive commands from IMS the external subsystem must supply a Command exit routine. IMS passes the command contained in the /SSR input to this exit routine. AOI (automated operator interface) programs can also send commands to external subsystems using /SSR. The /SSR command input contains identification, a

command recognition character (CRC), of the external subsystem to which the command is directed. The CRC for a subsystem is specified as part of the definition of the subsystem to IMS.

Related tasks:

“How external subsystems are specified to IMS” on page 141

Related reference:

 /SSR command (Commands)

IMS services available to the ESAP

IMS provides exit routines that an external subsystem can activate to access certain IMS system services.

The external subsystem can:

- Request that a connection be initiated (Subsystem Startup Service).
- Request that connections be quiesced (Subsystem Termination Service).
- Have a log record written to the IMS log (Log Service).
- Have a message sent to an IMS destination (Message Service).

Related concepts:

 Exit routines (Exit Routines)

Chapter 9. Creating the external subsystem module table

The external subsystem creates the external subsystem module table (ESMT) to supply definitions of the external subsystem modules that IMS is to load.

External subsystem exit routine modules, as well as any other modules needed in the ESAP, are defined in the ESMT. The installation provides the name of the ESMT to IMS as part of the definition of the external subsystem. During initialization for attach processing in the control and dependent regions, IMS loads the ESMT and then loads the modules defined therein.

The external subsystem optionally provides in the ESMT definitions of work areas needed for its ESAP. If work area definitions are provided, IMS obtains the specified work area storage in each region after loading the defined modules.

IMS provides two macros to be used by the external subsystem to create the ESMT. The DFSEMODL macro is used to define the ESAP modules that IMS is to load. The DFSEWAL macro is used to define the work areas that IMS is to create. A series of DFSEMODL statements defining modules, optionally followed by DFSEWAL statements defining work areas, and ending with a DFSEMODL statement specifying END=LAST, generates the table.

DFSEMODL macro

The external subsystem module table (ESMT) is generated from a series of DFSEMODL statements, one for each module definition.

In addition to module definition information, information about the control block that is to contain the addresses of the modules when they are loaded is also supplied on DFSEMODL statements. IMS creates this control block before it loads the modules.

IMS provides the capability for up to three sets of modules to be loaded and anchored on separate control blocks. Accordingly, the ESMT consists of one to three subtables, each containing the specifications for a set of modules and their module address control block. The module address control block for external subsystem exit routines is the EEVT (external entry vector table).

When the module address control block is created, IMS stores its address into a source control block, which is the EEVTP (EEVT prefix).

The format of the DFSEMODL macro is:

```
(label)  DSFEMODL  DSNAME=,SOURCE=,MODNAME=,DSLABEL=,  
          SUBPOOL=,OPTION=,END=
```

where:

(label) Is optional. If coded, the macro generates ESMT subtable labels. The last label on a macro statement in the series from which a subtable is generated is used as the subtable label.

The following parameters provide control block information and need only be specified once per subtable (for example, on the first DFSEMODL statement in the series). If specified on more than one statement, the first specifications encountered are used in generating the table.

DSNAME=

(p1,p2,p3)

- p1** Name of the module address control block. The name must be specified (at least on one DFSEMODL statement) for each ESMT subtable.
- p2** Module address control block size. The size must be specified. IMS obtains storage of the specified size to create the module address control block.
- p3** Subpool number for the module address control block storage request. This parameter is optional. If 251 is not specified, IMS obtains the storage from subpool 230.

SOURCE=

(p1,p2)

- p1** Name of the source control block. This parameter is required. DFSEEVTP must be specified. (See the following discussion.)
- p2** Label in the source control block of the location to store module address control block address. This parameter is required.

The following parameters provide module definition information.

MODNAME=

Name of the module IMS is to load. MODNAME must be specified on all DFSEMODL statements that do not specify the END parameter. (END can be specified with or without MODNAME.)

DSLABEL=

Label in the module address control block of the location to store the module address after it is loaded. DSLABEL must be specified (when MODNAME is specified).

SUBPOOL=

For resources that reside on a PDS data set:

The subpool into which IMS is to load the module. SUBPOOL must be specified when MODNAME is specified. For the control region, IMS loads the module into the subpool specified. For dependent regions, IMS loads the module into subpool 251 if SUBPOOL=251 is specified. Otherwise, the module is loaded into subpool 230. Valid specifications are 0, 229, 230, 231, 241, 251, 252.

For modules that reside on a PDSE data set:

The SUBPOOL parameter is not used. Modules residing on a PDSE are loaded in one of the following methods:

- Modules that are linked as reentrant (RENT) are loaded into subpool 252, key 0. These modules are not fetch-protected.
- Modules that are linked as not reentrant (NORENT) are loaded into subpool 251, TCBKEY, and are fetch-protected. You must ensure that the correct protect key is in use before accessing these modules.

OPTION=

(p,p)

This parameter is optional. Two options, NOCTL and NODEP, are supported. (Position of an option in the subparameter list is not important.)

NOCTL

The module is not to be loaded in the control region.

NODEP

The module is not to be loaded in dependent regions.

The END parameter controls ESMT generation.

END=

- YES** Must be specified to indicate the end of a subtable in the ESMT being generated. END=YES is used only when the ESMT is to contain more than one subtable. It is specified to end each subtable except the last (or only). DFSEMODL statements for the next subtable in the ESMT are to follow the END=YES specification.
- LAST** Must be specified on the last DFSEMODL definition statement for the ESMT being generated. The next DFSEMODL or DFSEWAL statement (if any) causes a new ESMT generation to be started.

You must bind the ESMT module into a program library (SDFSRESL) using a binder ENTRY statement that specifies MAINEP as the entry point. A table definition header is generated at the end of the ESMT module. The ENTRY statement allows IMS to correctly reference the header for subsequent processing.

DFSEMODL supports an execute form (MF=E) for internal use only. It cannot be used for ESMT generation. The list form (MF=L) described is the default.

Mapping DSECTs for the module address and source control blocks must be included in the ESMT generation source, otherwise the assembly will fail.

The following restrictions apply to external subsystems:

- Specifying the source control block.
DFSEEVTP must be specified as the source control block name (SOURCE(p1)) for all IMS-defined subtables. Otherwise, although DFSEMODL accepts other specifications, the module load process will fail, prohibiting a connection for the region experiencing the failure. The user must use the EEVTP mapping layout, as this is the layout that IMS expects.
- Defining subsystem exit routine modules.
DFSEEVT must be specified as the module address control block name (DSNAME(p1)) for the subtable that contains the subsystem exit routine module definitions.
The module address control block size (DSNAME(p2)) must be specified according to the size indicated by the EEVT mapping (EEVTLGTH) as shown in "Control block mapping" in *IMS Version 11 Exit Routines*.
EEVPPEA must be specified as the label in the source control block (SOURCE(p2)) to anchor the module address control block (EEVT). IMS does not check for this (nor does the macro) but uses the offset generated from the label specified to store the address. If the offset is incorrect, IMS will not be able to activate the exit routines.

The label (DSLABEL) specified for a particular subsystem exit routine module is used to generate the offset that IMS uses to store the exit routine address in the module address control block (for example, in the EEVT). Thus these labels must be specified according to the EEVT mapping.

- Generating additional subtables.

The ESMT must always have one subtable containing definitions for exit routine modules. The external subsystem could choose to have other modules needed in its ESAP anchored on a separate control block, which means that another subtable would be generated.

Although the DFSEMODL macro does not restrict the number of subtables that can be generated, problems can occur during processing if more than three (3) are generated. For each subtable, IMS creates a module address control block and stores its address in the EEVTP. There are only three fields in the EEVTP that could be used as anchors for these control blocks, one of them being the anchor for the EEVT.

The EEVTLDIR and EEVPEWA fields are not used by IMS and thus are available for this purpose. The discussion on defining external subsystem work areas in “DFSEWAL macro” suggests how EEVPEWA might be used to anchor a work area address control block.

- Defining external subsystem-unique modules.

If the ESAP needs non-IMS exit routine modules (for example, modules that the external subsystem activates without any knowledge of or support from IMS), the external subsystem can define these modules in an additional subtable as previously discussed. The external subsystem must supply the mapping DSECT for the module address control block for these modules.

Recommendation: Do not define other modules in the subtable containing exit routine modules or extend the size of the EEVT to include their addresses. The EEVT is an IMS control block which IMS reserves the right to extend at any time, which could require the external subsystem to regenerate the ESMT and re-compile modules.

DFSEWAL macro

The work areas that IMS is to create for the external subsystem must be defined by including DFSEWAL macro statements along with the DFSEMODL statements provided for ESMT generation.

The DFSEWAL statements, one for each work area defined, follow the DFSEMODL statements, except that the last statement in the series must be a DFSEMODL statement specifying END=LAST. The DFSEWAL statements cause a table of work area definitions to be built in the generated ESMT.

Work areas can be defined in each subtable generated in the ESMT. At least one module must be defined in each subtable. If a subtable is generated containing only work area definitions, an error occurs during IMS processing of the ESMT.

IMS creates the work areas defined in a subtable after loading the modules defined in the subtable. IMS stores the addresses of the created work areas in a work area list control block. This control block is also defined by the DFSEWAL macro and can either be contained in the module address control block for the subtable or be created as a separate control block. For this discussion, EWAL is used to refer to the external subsystem work area list control block.

Recommendation: Contain the EWAL in the module address control block rather than creating it as a separate control block. When IMS creates the EWAL, its address is not (explicitly) provided to the external subsystem. If, instead, the EWAL is contained in the module address control block, which IMS anchors in the EEVTP, the external subsystem specifies its location (with DFSEWAL) and thus knows how to access it. (When IMS creates the EWAL, it stores its address in the in-storage ESMT for internal use. The format of the ESMT is not included in the documented attach interface.) The following figure shows a representation of the relationship between the EWAL, EEVTP, and EEVT.

Figure 16. EWAL, EEVTP, and EEVT relationship

The second larger block of example code below illustrates how the external subsystem might specify work area definitions.

```
DFSEWAL DSNAME, SOURCE=, WALSP=, NAME=, DSLABEL=,
      SUBPOOL=, LV=, OPTION=
```

The following parameters provide control block information and need only be specified once per subtable (for example, on the first DFSEWAL statement in the series). If specified on more than one statement, the first specifications encountered are used in generating the table.

DSNAME=

(p1,p2)

- p1** Name of the work area list control block mapping DSECT. The DSECT name must be specified. If IMS creates the work area list, this name is given to the job pack entry for the storage acquired.
- p2** Work area list size. If the size is specified, IMS obtains storage of the specified size to create the work area list. If the size is not specified IMS does not create the work areas unless the source control block for the work area list (DFSEWAL SOURCE(p1) specification) is the module address control block specified for the modules defined in the subtable (DFSEMODL DSNAME(p1)). (See the following discussion.)

SOURCE=

(p1,p2)

- p1** DSECT name for the control block in which the work area list is to be anchored. This parameter must be specified. (See the following discussion).
- p2** Label in the source control block DSECT of the location that is to contain the work area list. This parameter is required. IMS does not store the work list area address into this control block. (See the following discussion.)

WALSP=

Subpool number for the work area list storage request. This parameter is optional. If WALSP=251 is not specified, IMS obtains the storage from subpool 230.

The following parameters provide work area definition information and must be specified on each DFSEWAL statement.

NAME=

The name given to the job pack directory entry created for the work area storage acquired. This parameter is required.

DSLABEL=

Label in the work area list control block DSECT of the location into which IMS is to store the work area address. DSLABEL must be specified.

SUBPOOL=

Subpool from which IMS is to obtain storage for the work area. The subpool must be specified. IMS acquires subpool 251 storage if SUBPOOL=251 is specified; otherwise, the work area is created in subpool 230. The macro allows 0, 229, 230, 231, 241, 251, or 252 to be specified.

LV= Work area size. The size must be specified.

OPTION=

(p,p)

This parameter is optional. Two options, NOCTL and NODEP, are supported. (Position of an option in the subparameter list is not important.)

NOCTL

Work area is not to be created in the control region.

NODEP

Work area is not to be created in dependent regions.

Mapping DSECTs for all referenced control blocks must be included in the ESMT generation source, otherwise the assembly will fail.

The source control block DSECT name and label must be specified. However, IMS does not store the EWAL address into this control block.

To indicate that the EWAL is to be contained in the module address control block:

- The size for the EWAL (DFSEWAL DSNAME(p2)) must not be specified.
- The module address control block DSECT name (DFSEMODL DSNAME(p1)) must be specified as the EWAL source control block DSECT name (DFSEWAL SOURCE(p1)).
- The source control block label (SOURCE(p2)) must specify the location of the work area list in the module address control block.

If the size for the EWAL is specified, IMS obtains storage for the EWAL without checking if the module address control block was specified as the EWAL source. If the EWAL size is not specified and the module address and EWAL source control block DSECT names do not match, IMS does not create the work areas. (IMS does not know the address of the source control block. IMS does not indicate that the work areas were not created.)

IMS reserves the EEVPEWA field in the EEVTP control block for the address of an EWAL. The following code illustrates how definitions can be specified by an external subsystem to anchor a work area list in this field. What really happens is that a module address control block is created, anchored at EEVPEWA, and used as the EWAL.

```
DFSEMODL DSNAME=(DFSEEVTP,68,230),SOURCE=(DFSEEVTP,EEVPEEA),
          MODNAME=INITEXIT,DSLABEL=EEVTINIT,SUBPOOL=230
DFSEMODL MODNAME=IDEXIT,DSLABEL=EEVTID,SUBPOOL=230
DFSEMODL MODNAME=RIDEXIT,DSLABEL=EEVTRID,SUBPOOL=230,
          OPTION=NODEP
          .      .      .
          .      .      .
          .      .      .
DFSEMODL MODNAME=CMDEXIT,DSLABEL=EEVTCMD,SUBPOOL=230,
          OPTION=NODEP
DFSEMODL END=YES
DFSEMODL DSNAME=(ESSEWAL,40,230),SOURCE=(DFSEEVTP,EEVPEWA),
          MODNAME=MODX,LABEL=MODXADDR,SUBPOOL=230
DFSEWAL  DSNAME=(ESSEWAL),SOURCE=(ESSEWAL,ESSEWAL),WALSP=230,
          NAME=WKA1,DSLABEL=WKA1ADDR,SUBPOOL=230,LV=200
DFSEWAL  NAME=WKA2,DSLABEL=WKA2ADDR,SUBPOOL=230,LV=100
DFSEWAL  NAME=WKA3,DSLABEL=WKA3ADDR,SUBPOOL=230,LV=100
DFSEMODL END=LAST
```

Notes to example:

- The existence of a DSECT named ESSEWAL created by the external subsystem to map the EWAL is assumed.
- Two subtables are defined for completeness:
 - The first subtable contains exit routine module definitions.
 - The second subtable contains work area definitions:

- A module is defined in this subtable with EEVPEWA specified as the anchor field for the module address control block. (If the external subsystem does not really want a module loaded for this subtable, both the NOCTL and NODEP options can be specified.)
- The module address control block DSECT, ESSEWAL, is specified as the EWAL source control block DSECT and the EWAL size is not specified, indicating that the EWAL is to be contained in the module address control block.
- ESSEWAL is also specified as the label in the source block for the EWAL, indicating that the EWAL starts at offset zero in the module address control block. Thus, the module address control block itself is the EWAL, anchored at EEVPEWA in the EEVTP.

Chapter 10. IMS External Subsystem Attach Facility processing

The IMS External Subsystem Attach Facility performs processing during control region initialization.

Loading the External Subsystem Attachment Package

During control region initialization, IMS loads external subsystem-supplied tables using the table names specified by the installation (in the external subsystem definition member in IMS.PROCLIB).

IMS loads the external subsystem module table (ESMT) and then loads the external subsystem modules defined in the table. The resource translation table (RTT) is also loaded, if provided. If an error occurs during this process, IMS puts the subsystem in 'stopped' status and does not establish a connection. However, IMS will reaccess the definition (PROCLIB) and reattempt this process if a /START SUBSYS command is received.

Possible errors are:

- Unable to process the PROCLIB member
- Unable to open the external subsystem load library
- Unable to load the ESMT (incorrect name specified, not in library)

IMS stores the addresses of the ESMT and the RTT in the EEVTP control block fields EEVPESMT and EEVPRTA, respectively.

Creating the EEVT control block

IMS creates and initializes the EEVT control block based on information contained in the external subsystem module table (ESMT) that is generated from DFSEMODL macro statements.

The size and subpool for the EEVT storage request are obtained from the ESMT. The EEVT is created in subpool 230 unless subpool 251 is specified in the ESMT. The external subsystem must ensure that the size specified for the EEVT is at least as large as the size indicated in the IMS EEVT mapping.

IMS stores the EEVT address into the EEVTP control block based on an offset specified in the ESMT. The external subsystem must ensure, therefore, that the offset generated in the ESMT (using the DFSEMODL macro) points to the EEVPEEA field in the IMS EEVTP mapping.

IMS does not check whether the EEVT pointer field, EEVPEEA, in the EEVTP is initialized by this process. In fact, the offset in the ESMT could cause IMS to store the address into some other field in the EEVTP designated for some other use, possibly causing a problem. Thus, **the external subsystem must ensure that the correct offset is generated into the ESMT.**

This process allows the external subsystem to specify another set of modules for IMS to load (IMS would not activate these modules). Both lists of module addresses, one being the EEVT, would be anchored in the EEVTP.

IMS does not use the EEVTLDIR field. Actually more than two sets of modules could be defined in the ESMT (subtables) and loaded by IMS except that there are not enough fields in the EEVTP to anchor the address lists.

Related concepts:

Chapter 9, “Creating the external subsystem module table,” on page 149

Loading external subsystem modules

As IMS loads external modules defined in the ESMT, the module addresses are stored in the EEVT.

The module definitions provide the offsets to the locations in the EEVT for the addresses. IMS does not check whether or not required exit routine addresses have been set by the module loading process. If the external subsystem chooses, the ESAP can set exit routine addresses in the EEVT once IMS has passed control to it. For example, the external subsystem can provide multiple exit routines in one load module and have the ESAP set the individual exit routine module addresses.

Whether through module definitions in the ESMT or through ESAP processing, the external subsystem must ensure that the address of an exit routine is present in the EEVT when IMS needs to activate the exit routine. Exit routine addresses must be placed in the EEVT according to the IMS EEVT mapping.

Some of the exit routines prescribed by IMS are activated only in the control region; some are activated only in the dependent regions. The external subsystem can indicate, in the module definition, if a module is not to be loaded in the control region or in dependent regions. Exit routine module definitions should specify loading according to the following:

- Exit routines activated in the control and dependent regions:
 - Identify
 - Initialization
 - Terminate Identify
- Exit routines activated only in the control region:
 - Command
 - Echo
 - Resolve Indoubt
 - Subsystem Termination
- Exit routines activated only in dependent regions:
 - Abort Continue
 - Commit Continue
 - Commit Prepare
 - Commit Verify
 - Create Thread
 - Normal Call
 - Signoff
 - Signon
 - Subsystem Not Operational
 - Terminate Thread

In the control region, IMS loads external subsystem modules into the subpools specified in the module definitions. If subpool 251 is specified for a module in dependent regions, IMS loads the module in subpool 251; otherwise, it is loaded in subpool 230.

An external subsystem uses only those exit routines that it needs to communicate with IMS, although some exit routines are required and others are optional. When a required exit routine does not exist, IMS generates an error message when it tries to call the exit routine and terminates the connection with the external subsystem.

If your external subsystem does not need the function that an exit routine is designed to perform, you can write the exit routine so that one exists when IMS calls it but that so no operations are performed. (An exit routine can contain common code, such as SR 15,15 and BR 14 logic, which ESS branches to when the exit routine is called and which does not perform any specific operation.) During processing of the Initialization exit routine, the external subsystem can update the addresses in the DFSEEV T DSECT (from both the control region and dependent regions, if necessary) and point to these exit routines. This action allows IMS to function normally yet not issue error messages and terminate an external connection if an exit routine does not exist.

Creating work areas for the ESAP

After loading the ESAP modules, IMS obtains work area storage for the ESAP if work area definitions are contained in the ESMT.

The IMS DFSEWAL macro is used to generate work area definitions in the ESMT.

The process IMS uses to create the work areas is similar to the process used to load ESAP modules except that IMS can either:

- Create the control block for the work area addresses (as it creates the EEVT for the ESAP module addresses), or
- Store the work area addresses into the same control block that has the module addresses.

The intended use of the EEVPEWA field in the EEVTP is to hold the address of a control block referred to as the external subsystem work area list (EWAL) that contains the addresses of the work areas created for the ESAP. However, the external subsystem must have provided the appropriate specifications in the ESMT to cause IMS to store the address of the EWAL in this field.

IMS creates each work area either in the control region or in dependent regions or both, according to the definition. Storage is obtained in subpool 251, if specified; otherwise it is obtained in subpool 230.

Related concepts:

Chapter 9, "Creating the external subsystem module table," on page 149

Initiating the external subsystem connection

IMS automatically connects to the external subsystem during control region initialization processing (for example, when IMS is started) unless the external subsystem chooses to defer the control region identify to a later time.

The external subsystem defers the connection by returning from the control region Initialization exit routine with return code 4 (do not identify), or by not providing an Initialization exit routine for the control region.

If the external subsystem uses the notify message mechanism provided by IMS (and if the external subsystem is not up when IMS activates the Identify exit routine) the connection is automatically established when the external subsystem is started. Return Code 4 from the Identify exit routine causes IMS to wait for the external subsystem to send the notify message passed to the exit routine, and to reactivate the exit routine when the message is received.

If the notify mechanism is not used, the Identify exit routine should return with return code 12, in which case the connection is put in stopped status. Stopped status must be removed by a /START SUBSYS command before IMS will establish a connection.

Once the connection has been established, IMS performs Resolve Indoubt processing to resolve any outstanding recovery tokens with the external subsystem. If outstanding recovery tokens exist and a Resolve Indoubt exit routine was not supplied, IMS terminates and stops the connection; otherwise dependent regions are allowed to connect to the external subsystem.

Deferring the control region identify

The external subsystem can defer the control region identify if it prefers to have the connection established at some later time.

IMS schedules and gives control to application programs whether or not a connection exists to the external subsystem. The external subsystem thus could choose to wait until an application program call has to be serviced (first call for its resources) before connecting to IMS. Of course, to process calls, the control and dependent region connections and the application thread must exist.

When the control region identify is deferred, the identify is done when:

- The external subsystem, with an exit routine, activates the IMS Subsystem Startup Service.
- An MPP or IFP dependent region that can access the external subsystem is started.
- A /START SUBSYS command naming the external subsystem is processed.

Using the IMS Subsystem Startup Service

When the control region identify is deferred, the external subsystem can activate the IMS Subsystem Startup Service when it wants the control region Identify exit routine to be driven.

To be more specific, if a connection does not exist when the first application call for external subsystem services is processed by a dependent region, IMS does not automatically attempt to identify. The external subsystem must activate the Startup Service to establish the connection (if it **wants** to process application calls).

The Startup Service is also used to establish dependent region connections. When an external subsystem call from an application is processed before the control region or dependent region has been identified to the external subsystem, the

dependent region activates the Subsystem Not Operational exit routine. The external subsystem is expected to call the Subsystem Startup Service from this exit routine to establish the connection.

When activated, the Startup Service establishes the control and dependent region connections, if the control region identify has not been done. If the control region identify has been done, it establishes only the dependent region connection. If IMS is waiting for the external subsystem to send the notify message, which it accepted on a previous Identify exit routine invocation, the Startup Service returns an error return code and does not establish the connection. For details on using the Subsystem Startup Service exit routine, see *IMS Version 11 Exit Routines*.

Related concepts:

“Establishing dependent region connections”

Establishing dependent region connections

Connections can be established to MPP, IFP, or BMP regions.

MPP and IFP regions

Identify processing for an MPP or an IFP dependent region is similar to identify processing for the control region in that the dependent region automatically activates its Identify exit routine to establish a connection during dependent region initialization, unless the external subsystem defers the identify.

If the control region connection has not been established when the dependent region would automatically identify, IMS attempts to identify the control region. If successful, the dependent region identify is performed. Thus, if the control region identify is deferred but dependent regions are allowed to connect automatically, the control region Identify exit routine might be activated (automatically) when a dependent region is started.

When the identify for a dependent region is deferred, the connection to the external subsystem is established when the first application program call to the external subsystem is issued in the region. In this case, connection processing is the same as for a BMP dependent region.

BMP regions

The connection from a BMP dependent region is not established until the first application call to the external subsystem is processed by the region. The connection is automatically established.

Return code 4 from the Initialization exit routine (deferred identify) for a BMP region has no effect.

Notify message

IMS passes the address of a notify message on the Identify exit routine invocation for the control region.

If the external subsystem is not active (has not been started), the Identify exit routine can indicate (return code 4) to IMS that the notify message has been accepted and will be sent to IMS when the external subsystem is active. The external subsystem, once started, sends the message to IMS using an internal z/OS

and padded with blanks on the right. The Identify exit routine can obtain the job name from the TIOT pointed to by the current TCB.

COMMA

Is a 1-byte field containing a comma (C',').

MESSAGE_TEXT

Is the notify message text passed to the Identify exit routine.

Application program request support

Application calls are passed to IMS from the language interface module that you bind with the application program. The language interface branches to the appropriate IMS program request handler passing the application program call parameter list.

For calls directed to external subsystems, the language interface must also pass an external subsystem parameter list which it constructs. The purpose of this parameter list is to pass the LIT (language interface token) for the external subsystem to which the call is directed. IMS routes the application call to the external subsystem whose LIT value matches the LIT value passed on the call.

IMS passes both the call parameter list and the external subsystem parameter list to the Normal Call exit routine for the intended external subsystem. The first word of the external subsystem parameter list contains the address of a 4-byte field containing the LIT value in character format, left justified and padded on the right with blanks. IMS prescribes only the first word in the external subsystem parameter list (address of the LIT). The parameter list can be extended to provide external subsystem-dependent information to the Normal Call exit routine.

Language interface definition

IMS provides a language interface module, DFSLI000, which supports the value of SYS1. The installation can use this module or it can define its own language interface if it wants to use a LIT value other than SYS1.

When two or more external subsystems are accessed by the IMS system, the installation must define its own language interface modules because each subsystem has a unique LIT.

IMS provides the DFSLI macro to assist the installation in generating a language interface module. The code necessary to perform the language interface function is generated in the DFSLI macro expansion. The IMS macro library must be supplied when the macro statements are compiled to generate the module.

Related concepts:

 Defining the language interface module (System Definition)

Language interface entry points unique to external subsystems

The IMS language interface module provides three entry points that application calls directed to an external subsystem can exclusively use.

Two of the entry points are associated with an implied LIT value specified with the DFSLI macro. (The language interface module, generated by the DFSLI macro, contains the specified LIT value as a hard-coded constant.) The third entry point is not associated with an implied LIT value; it allows the application program to

specify the LIT value when it makes the application call. For all entry points, register 1 contains the address of the parameter list which IMS passes to the external subsystem. The following are language interface entry points to external subsystems:

DSNHLI

Entry point associated with an implied LIT value.

The application program does not need to know which subsystem provides access to the external resources it uses. (If the external subsystem is DB2 for z/OS, this entry point is used for SQL calls.)

DSNWLI

Entry point associated with an implied LIT value.

The application program does not need to know which subsystem provides access to the external resources it uses. (If the external subsystem is DB2 for z/OS, this entry point is used for Instrument Facility calls.)

DFSESS

Entry point allowing an application program to specify an LIT value.

The application program must know which subsystem provides access to the external resources it uses. The application program must specify the address of the LIT value as the first parameter in the application call list. Before it passes control to the external subsystem, IMS increments the address of the application call list by four to skip over the LIT value parameter.

Restriction: Do not use the DFSESS entry point to communicate with a DB2 for z/OS subsystem.

Accessing multiple external subsystems

An application program can access DL/I and an external subsystem in the same execution. Whether or not an application program can access more than one external subsystem in the same execution can be restricted by the language interface.

Where the data (call) interface provided to application programs by one external subsystem (product) is distinct from the interface provided by another external subsystem (for example, DL/I calls as distinct from SQL calls), an application can access both subsystems because the language interface paths can be different. Where the data interface is the same, as in the case of two external subsystems of the same type (two instances of the external subsystem) or two external subsystem products that use the same call interface (for example, SQL), an application cannot access both in the same execution unless the application is written to be dependent on data location. (The dependency is intrinsic in the case of different call interfaces.)

Resource recovery token

A 16-byte recovery token is used to uniquely identify a unit of work across all subsystems to which the application has thread connections. IMS passes the token to the Signon exit routine before the thread is created.

For commit and resolve indoubt processing, IMS passes the recovery token to identify the unit of work for which the requested action is to be taken.

The recovery token is constructed as shown in the following figure.




Figure 19. Recovery token format

where:

IMS-id

Is the IMS system ID (1 to 4 characters), left justified and padded with blanks on the right to eight bytes.

OASN

Is a 4-byte binary origin application sequence number assigned to the application when it is scheduled. The OASN is assigned based on the scheduling order within the IMS system since the last cold start. It is also referred to as the application schedule number.

commit_number

Is a 4-byte binary commit number. The commit number is initialized to binary zeroes when the application is scheduled and then incremented after each commit is processed for the application.

The external subsystem should check the recovery token passed at signon for uniqueness. Cold starts of IMS can cause a recovery token to be generated that is a duplicate of a recovery token that is indoubt in the external subsystem. The Signon exit routine can indicate to IMS that the recovery token passed was found to be a duplicate, in which case IMS terminates the application program with an abend. The Commit Prepare exit routine can also indicate that the token is a duplicate that supports external subsystems that choose not to associate the recovery token with the unit of work until commit is processed.

The installation uses the /DISPLAY SUBSYS command with the OASN keyword to determine what units of work are in indoubt status in IMS. The installation can use the /CHANGE command (when necessary) to manually delete indoubt units of work in IMS. The /CHANGE command only affects unit of work status in IMS. There is no communication with the external subsystem. These commands use only the OASN and not the full recovery token; /DISPLAY lists only the OASN portion of the recovery token (in decimal format) and /CHANGE accepts just the OASN (again in decimal format). (Within IMS, the OASN is unique across all known units of work.)

Terminating the external subsystem connection

IMS terminates the external subsystem connection (control region connection) in an orderly manner in a number of different circumstances.

IMS terminates the external subsystem connection (control region connection) in an orderly manner when one of the following occurs:

- IMS processes a /STOP SUBSYS command.
- The external subsystem (ESAP) activates the IMS subsystem termination service exit routine.
- The external subsystem posts the termination ECB provided on the Identify exit routine invocation.

- Certain attach processing errors are encountered.
- IMS is shutting down.

IMS allows existing threads to the external subsystem to complete processing. When all threads have terminated, IMS terminates the connection by activating the Terminate Identify exit routine from the control region.

When the connection is terminated due to a /STOP SUBSYS command, the termination ECB being posted, or processing errors, IMS puts the external subsystem connection in stopped status. Once in stopped status, IMS does not allow a connection to be reestablished. A /START SUBSYS command is required to remove the stopped status.

Termination requested by the external subsystem

The external subsystem can cause the connection to be terminated either by posting the termination ECB or by activating the Subsystem Termination Service exit routine from the ESAP. The connection is not put in stopped status when the service is used.

The termination service might be used in conjunction with the external subsystem command exit function. For example, when the command exit routine is activated with an external subsystem termination command supplied on an IMS /SSR command, the exit routine could activate the Subsystem Termination Service exit routine to cause the connection to the external subsystem to be terminated.

On the initial identify performed in the control region, IMS provides the external subsystem with the address of a termination notification ECB. When the subsystem needs to terminate the connection, it posts the ECB. The ECB is located in CSA. Depending on the post code, IMS terminates the connection in the following manner:

Deactivate all active threads, prohibit the initiation of any new threads, and then terminate the connection. Upon completion of the terminate function, the connection is set in a stopped state.

Supported Post Codes:

X'40000000'

Reserved for IMS.

X'40000008'

External subsystem is terminating in an orderly fashion.

X'4000xxxx'

All other post codes are interpreted as a quick or catastrophic shutdown of the external subsystem.

Related reference:

 Subsystem Termination Service exit routine (Exit Routines)

Dependent region connections

At dependent region termination, the Signoff and Terminate Identify exit routines are not activated unless the control region connection is to be terminated.

That is, the dependent region Signoff and Terminate Identify exit routines are activated only when the IMS system continues to run without a connection to the external subsystem, such as when either the external subsystem has posted the

termination ECB provided on the identify, the ESAP has activated the IMS Subsystem Termination Service, or IMS has processed a /STOP SUBSYS command. Otherwise, IMS does not communicate dependent region termination to the external subsystem.

IMS expects the external subsystem to monitor the dependent TCB with a z/OS End Of Task (EOT) exit routine. The subsystem should do any signoff and terminate identify processing it requires when the EOT exit routine is notified of the region termination.

The external subsystem must also monitor EOT exit routine for dependent regions for which a thread was created. When an application program terminates abnormally, the Terminate Thread exit routine is not called.

When the control region connection is to be broken, a signoff followed by a terminate identify is done for the dependent region at region termination or after thread termination if a thread was active when the request to break the system connection was received. The Signoff exit routine is called only once even though more than one signon might have been done for the region. IMS continues its signoff and terminate identify processing and does not reactivate these exit routines if they encounter errors.

Explanation of stopped status

The installation should be aware of the conditions that cause IMS to stop the external subsystem connection.

After an external subsystem connection is stopped, the /START SUBSYS command must be used to reestablish the connection. Stopped status is carried across restarts. The following list includes the conditions that cause the stopped status to be set:

- When the external subsystem posts the subsystem termination ECB provided during identify processing. Regardless of the post code type (for example, orderly or catastrophic), the connection is stopped upon completion of the termination processing.
- On abnormal termination of the IMS task (TCB) in the control region under which the ESAP is activated, the external subsystem connection is marked stopped.

If the control region abends, it is unlikely the stopped state will be set. After a successful IMS restart, the connection is in the state it was prior to the abend.

- The obvious case is after processing of the /STOP command. Even if IMS abends while processing the /STOP command, the stopped state is set.
- When IMS restarts, if outstanding recovery tokens exist for an external subsystem that is no longer defined to IMS (for example, the installation deleted its definition from the external subsystem definition member in IMS.PROCLIB), stopped status is set for that external subsystem.

Part 5. IMS Connect and TCP/IP communications

The IMS Connect function of IMS provides access to both IMS DB and IMS TM from TCP/IP-enabled environments.

Chapter 11. Overview of IMS Connect

IMS Connect provides high performance TCP/IP and local z/OS communications between one or more IMS Connect clients and one or more IMS systems. IMS Connect supports both IMS DB and IMS TM systems.

IMS Connect enables:

- Distributed clients to exchange messages with IMS DB by using TCP/IP connections and the Open Database Manager (ODBM) component of the IMS Common Service Layer (CSL).
- Distributed clients to exchange messages with IMS TM by using TCP/IP connections and OTMA.
- Local z/OS clients that use the local option with IMS TM Resource Adapter to exchange messages with IMS TM by using the z/OS Program Call (PC) function (also referred to as the *local option*) and OTMA.
- IMS Operators that use IBM Management Console for IMS and DB2 for z/OS to issue commands to an IMSplex and receive command replies by using TCP/IP and the IMS Operations Manager (OM).

IMS Connect provides the following features:

- Commands to manage the communication environment.
- Assistance with workload balancing.
- Reduced design and coding efforts for client applications.
- Support for IMSplexes, which enables communications with other IMSplex members, such as:
 - ODBM
 - OM
- TCP/IP access to IMS application programs and operations on demand with advanced security and transactional integrity.
- TCP/IP access for distributed applications to IMS databases on demand through clients such as the IMS Universal drivers and ODBM.
- XML conversion support for certain IMS Connect clients, such as IMS Enterprise Suite SOAP Gateway and IMS Web 2.0 Solution for IBM Mashup Center. The IMS Connect XML conversion support converts input messages into the data structures expected by IMS application programs written in select programming languages, thereby eliminating the need to create or modify IMS application programs to process XML.

IMS Connect connects to IMS DB through ODBM for direct access to databases that are managed by IMS DB. IMS Connect connects to IMS TM for transaction processing support through Open Transaction Manager Access (OTMA).

Communications between IMS Connect and other IMSplex members, such as ODBM and OM, requires the use of the IMS Structured Call Interface (SCI).

IMS Connect performs router functions between its clients and IMS and IMSplex resources. Request messages that are received either from distributed clients through TCP/IP connections or from local clients through the z/OS program call interface are passed to an IMS system, referred to as a data store, through z/OS

cross-system coupling facility (XCF) sessions. IMS Connect receives response messages from the data store and then returns them to the originating TCP/IP or local option clients.

If the data store terminates, the status of the data store is sent to IMS Connect from OTMA through XCF. If IMS Connect was connected to the data store when the data store terminated, when the data store is restarted, IMS Connect is notified and automatically reconnects to the data store. You do not need to manually reconnect to the data store.

Generally, the term *data store* refers to an IMS system. More precisely, however, the term data store represents the OTMA target member (tmember) connection to an IMS system. For example, an instance of IMS Connect can have multiple data store definitions for the same IMS system, in which case each data store represents a different OTMA tmember connection to that IMS system.

IMS Connect also supports callout requests from IMS application programs running in IMS dependent regions. IMS application programs issue callout requests to request data or services from a provider that is external to the IMS installation. During a callout request, IMS acts as a client and the external provider is the server.

Related concepts:

➞ IMS Connect definition and tailoring (System Definition)

Related tasks:

Chapter 15, “IMS Connect support for callout requests,” on page 223

Related reference:

➞ HWSCFGxx member of the IMS PROCLIB data set (System Definition)

➞ Integrated IMS Connect commands (Commands)

➞ IMS Connect exit routines (Exit Routines)

➞ HWS messages (IMS Connect) (Messages and Codes)

➞ IMS Connect return and reason codes (Messages and Codes)

IMS Connect client support

As a TCP/IP server and a message router for IMS, IMS Connect provides access to IMS TM, IMS DB, and the CSL Operations Manager (OM). The client support provided by IMS Connect differs, depending on which type of access the IMS Connect client needs.

IMS Connect supports TCP/IP clients that communicate with socket calls, as well as TCP/IP clients that communicate with different input data stream formats.

IMS DB client support

For access to IMS DB, IMS Connect, with the CSL Open Database Manager (ODBM), supports the following types of clients:

- Application programs that use the IMS Universal Database resource adapter for the Java EE platform
- Application programs that use the IMS Universal JDBC driver
- Application programs that use the IMS Universal DL/I driver

- User-written client application programs that use the open standard DRDA[®] communications architecture

The following figure illustrates an IMS Connect system configuration that supports IMS DB client communications:




Figure 20. Overview of IMS Connect support for IMS DB systems

IMS TM client support

For access to IMS TM, IMS Connect supports the following types of clients:

- User-written TCP/IP client application programs that use the IMS request message (IRM) header to communicate protocol options to IMS Connect
- IMS TM Resource Adapter (previously known as IMS Connector for Java)

- IMS Enterprise Suite SOAP Gateway, which includes XML message conversion support
- IMS Web 2.0 Solution for IBM Mashup Center, which includes XML message conversion support
- Local z/OS application programs that use the local option for client communications and IMS TM Resource Adapter

The following figure illustrates an IMS Connect system configuration that supports IMS TM client communications:




Figure 21. Overview of IMS Connect support for IMS TM systems

For IMS TM IMS Connect clients, you can write user-message exits that execute in the IMS Connect address space to convert the format of client input messages to the OTMA message format before IMS Connect sends the message to IMS. The user-written message exits also convert the OTMA message format to the customer message format before returning the message to IMS Connect. IMS Connect then sends output to the client.

IMS Connect supports TCP/IP communications between IMS TM and distributed Java applications through the IMS TM Resource Adapter running under either WebSphere Application Server for distributed platforms or WebSphere Application Server for z/OS.

For IMS TM, IMS Connect also supports the SOAP Gateway, which is a web services solution that enables IMS applications to inter-operate outside of the IMS environment through SOAP to provide and request services that are independent of platform, environment, application language, or programming model.

IMS Connect supports the IMS Web 2.0 Solution for IBM Mashup Center, which enables the integration of existing IMS assets into Web 2.0 mashup and application solutions, providing users access to IMS transactions through RSS, ATOM, or XML feed.

For SOAP Gateway and IMS Web 2.0 Solution clients, IMS Connect also provides XML conversion support, which converts incoming XML messages into the data structures of some of the common programming languages supported by IMS application programs.

In addition to TCP/IP client communications, IMS Connect also supports local communication through the *local option*. This option provides a non-socket (non-TCP/IP) communication protocol for use between WebSphere Application Server for z/OS and IMS Connect in the z/OS environment. Servlets that run in WebSphere Application Server for z/OS and the use the IMS TM Resource Adapter can communicate with IMS Connect through the local option.

To use the local option for client communications, IMS Connect and an instance of WebSphere Application Server for z/OS running IMS TM Resource Adapter must reside in the same z/OS image.

The local option for client communications is supported only through IMS TM Resource Adapter and the IMS Connect HWSJAVA0 user message exit.

IMSplex operations support

To issue IMS type-2 commands to the CSL OM and receive command responses through a TCP/IP connection, IMS Connect supports clients such as IBM Management Console for IMS and DB2 for z/OS. A single IMS Connect can support communication between the TCP/IP client and any IMS within an IMSplex.

The following figure illustrates an IMS Connect system configuration that supports IMS TM client communications:




Figure 22. Overview of IMS Connect support for TCP/IP clients

IMS Connect support for access to IMS DB

For IMS Connect clients, such as the IMS Universal drivers, that access databases that are managed by IMS DB in DBCTL and DB/TM environment, IMS Connect manages TCP/IP connections and routes incoming access requests among the instances of the CSL Open Database Manager (ODBM) and the IMS DB systems in an IMSplex.

IMS Connect is the TCP/IP server and front-end IMSplex message router for the IMS Universal drivers, which include:

- IMS Universal Database resource adapter for the Java EE platform
- IMS Universal JDBC driver
- IMS Universal DL/I driver

The IMS Universal drivers provide direct, non-transactional access to IMS databases through TCP/IP connections for distributed application programs or local z/OS application programs.

The IMS Universal drivers simplify the development of IMS Connect client application programs that access IMS DB through TCP/IP connections. The IMS Universal drivers do not use the IMS Connect IMS Request Message (IRM) communications protocol. The IMS Universal drivers also shield application developers from the underlying Distributed Relational Database Architecture™ (DRDA) that is used internally for communications between the IMS Universal drivers and IMS Connect.

Because IMS Connect supports the DRDA protocol and, with ODBM, is a complete DRDA target server, you can write application programs to the DRDA protocol directly; however, the IMS Universal Database resource adapter for the Java EE platform is the recommended API for accessing IMS databases through TCP/IP from a distributed environment.

IMS Connect support for the IMS Universal drivers includes support for global two-phase commit transactions.

IMS Connect supports communication with the IMS Universal drivers only on dedicated DRDA ports and only through shareable persistent sockets.

IMS Connect security support includes the IMS Connect DB Security user exit routine (HWSAUTH0), which can be used for greater control over the authentication of user IDs on connections that access IMS DB. RACF is also supported.

IMS Connect support for the IMS Universal drivers is defined by the ODACCESS configuration statement in the IMS Connect configuration PROCLIB member and requires at least one instance of ODBM running in the same IMSplex as IMS Connect.

Connection routing for IMS Connect clients that connect to IMS DB

IMS Connect routes incoming connections from client applications that are connecting to IMS DB based on an *alias name* and the instance of ODBM to which the alias name belongs. The client application programs can specify an alias name on their incoming request messages.

IMS Connect keeps track of which alias names belong to which instance of ODBM in a internal tracking table that is populated during registration with each instance of ODBM. Information about the ODBM instances and alias names is also stored in an ODBM list in an exit interface block mapped by the HWSXIBOD macro. The address of the ODBM list is stored in a table mapped by the HWSXIB macro.

IMS Connect also provides the IMS Connect DB Routing user exit routine (HWSROUT0) that you can use to further control the routing of incoming message from clients that access IMS DB.

If IMS Connect receives a blank alias name from the client application or the HWSROUT0 exit routine, IMS Connect routes incoming connections in round-robin fashion among all of the instances of ODBM known to IMS Connect. If IMS Connect receives from the client application or the HWSROUT0 exit routine an alias name that is shared by multiple instances of ODBM, IMS Connect routes the incoming connections that specify that alias name in round-robin fashion among all of the instances of ODBM that share the alias name.

IMS Connect support for the IMS TM Resource Adapter

To support Java application programs connecting to IMS TM, IMS Connect supports the IMS TM Resource Adapter.

The IMS TM Resource Adapter runs under WebSphere Application Server in either z/OS or distributed environments and is used by Java applications, Java Platform, Enterprise Edition (Java EE) applications or web services to access IMS transactions

that are running on host IMS systems. It is also used by IMS applications that run in IMS dependent regions to make asynchronous callout requests to external Java EE applications.

When running in a z/OS environment, the IMS TM Resource Adapter can connect to IMS Connect through either the TCP/IP communications protocol or through z/OS program calls. The use of z/OS program calls is referred to as the *local option*. When running in a distributed environment, the IMS TM Resource Adapter connects to IMS Connect through only the TCP/IP communications protocol.

When a Java application submits a transaction request to IMS through the IMS TM Resource Adapter, IMS Connect sends the transaction request to IMS Open Transaction Manager Access (OTMA) by using the z/OS cross-system coupling facility, and the transaction runs in IMS. The response is returned to the Java application using the same path.

When an IMS application invokes an external enterprise JavaBeans™ (EJB) component or web service through a callout request, IMS Connect retrieves the callout request from a hold queue and passes it to the IMS TM Resource Adapter, which in turn passes the request to an EJB application in WebSphere Application Server that is set up to receive callout requests. The EJB starts a connection to IMS Connect via the IMS TM Resource Adapter. The IMS TM Resource Adapter polls IMS Connect to retrieve the callout requests from the hold queue. The EJB processes the request, and returns any response data to IMS by issuing a normal IMS transaction request.

IMS Connect identifies both TCP/IP connections and local option connections by a client ID that is submitted by the IMS TM Resource Adapter.

For connections on persistent TCP/IP sockets, IMS Connect can generate a client ID for the IMS TM Resource Adapter in any of the following cases:

- The IMS TM Resource Adapter requests that IMS Connect generate a unique client ID in the event that a client ID submitted by the IMS TM Resource Adapter is a duplicate of an existing IMS TM Resource Adapter client ID.
- The IMS TM Resource Adapter submits a blank client ID.

For connections that use the local option, IMS Connect also ensures the uniqueness of client IDs.

Related concepts:

 [IMS TM Resource Adapter overview](#)

Overview of IMS Connect XML Conversion Support

For certain IMS Connect clients, IMS Connect can convert the XML data contained in an input message into the data structures used by IMS application programs written in either COBOL or PL/I. The data in the corresponding output message is also converted from programming language of the IMS application program back to the XML data that IMS Connect client expects.

The IMS Connect XML conversion support enables IMS to accept messages in an XML format without having to create or modify IMS application programs to support XML.

The IMS Connect XML conversion support uses the HWSSOAP1 user message exit to identify the appropriate XML adapter and XML converter. The HWSSOAP1 user message exit is used by both IMS Enterprise Suite SOAP Gateway and IMS Web 2.0 Solution for IBM Mashup Center.

IMS Connect calls the XML adapter, which serves as the interface to the XML converter. The XML converter is generated based on either the COBOL copybook or the PL/I source of the IMS application program, depending on which language the IMS application program is written in. After the XML data has been converted into the data structures of the programming language of the IMS application program, the input message is passed to OTMA.

Related concepts:

Chapter 16, “IMS Connect XML message conversion,” on page 231

Related tasks:

 Configuring XML conversion support for IMS Connect clients (System Definition)

IMS Connect support for z/OS Sysplex Distributor

IMS Connect includes a variety of features that facilitate its execution in a z/OS Sysplex Distributor environment.

In a z/OS Sysplex Distributor environment, incoming messages are typically distributed among multiple instances of IMS Connect to balance the workload and increase availability. In such an environment, client applications have no control over which instance of IMS Connect receives their input messages and which IMS Connect receives subsequent requests for asynchronous output. IMS Connect, with OTMA, provides several features to support operating in such an environment, such as rerouting asynchronous output to an alternate tpipe, sharing asynchronous output by using an OTMA super member tpipe, retrieving output from an alternate tpipe queue associated with another client, and purging undeliverable output.

Related concepts:

“Rerouting commit-then-send output” on page 301

Related tasks:

“Sharing asynchronous commit-then-send output: the OTMA super member function” on page 843

“Purging undeliverable commit-then-send output” on page 299

“Retrieving output from alternate OTMA tpipe hold queues” on page 339

Overview of IMS Connect security

IMS Connect provides different security options depending on whether a client is accessing IMS DB or IMS TM.

IMS DB clients can implement security by using the IMS Connect DB Security user exit routine (HWSAUTH0), a security product such as RACF, or both. For Secure Sockets Layer (SSL) support, IMS DB clients can use IBM z/OS Communications Server Application Transparent Transport Layer Security feature (AT-TLS). IMS Connect does not provide SSL support for IMS DB clients.

IMS TM clients can implement security using any combination of the IMS Connect user message exit routines, a user security exit routine, and a security product such as RACF. For IMS TM clients, IMS Connect provides direct support for SSL and support for RACF PassTickets.

In an IMS Connect configuration, security can be implemented by using various combinations of the following components:

- On the client side:
 - The client application
 - The server of the client application
- IMS Connect
- A security product, such as RACF
- For IMS TM connections:
 - The IMS Connect user message exit routines
 - OTMA, including the OTMA Resume Tpipe Security exit routine (DFSYRTUX)
- For IMS DB connections, the IMS Connect DB Security user exit routine
- IMS
- An IMS exit routine
- An IMS application program

When you configure IMS Connect, you can enter your security specifications in the following places:

- HWSCFGxx configuration member
- The RACF FACILITY class

Related concepts:

Chapter 14, “IMS Connect security support,” on page 199

Overview of defining and invoking IMS Connect

The steps below provide a high-level overview of defining and invoking IMS connect.

Configuring IMS Connect includes the following high-level steps.

1. Authorize SDFSRESL to the Authorized Program Facility (APF) by online command or by running a JCL job.
2. Ensure that the z/OS Program Properties Table (PPT) is updated to allow IMS Connect to run in authorized supervisor state and in key 7. The specification in the z/OS PPT must match the program specification in the EXEC statement of the IMS Connect startup JCL. You can specify either BPEINI00 or HWSHWS00.
3. Create an IMS Connect configuration member, and define the IMS Connect configuration statements that IMS Connect uses during initialization.
4. Create the BPE configuration member for IMS Connect.
5. Create the HWSUINIT initialization exit routine, which is not shipped as a load module.
6. If you are connecting to IMS TM, create the user message exit routines that you need, such as HWSSMPL0, HWSSMPL1, or HWSJAVA0. These user message exit routines are also not shipped as load modules.
7. Define IMS Connect security and, optionally, OTMA client security.
8. Optionally, enable the IMS Connect XML message conversion support.

Invoke IMS Connect using either a z/OS procedure or a z/OS job. If you start multiple instances of IMS Connect with the same configuration, a connection outage can occur.

Related concepts:

 [IMS Connect definition and tailoring \(System Definition\)](#)

Chapter 12. Overview of IMS Connect exit routines

IMS provides a variety of exit routines to support IMS Connect.

IMS Connect exit routines fall into two general categories:

- User message exit routines that manage the messages to and from the various types of IMS Connect TCP/IP clients
- Exit routines that provide general functionality, such as security and routing

The IMS Connect user message exit routines, which are used only with clients connecting to IMS TM, include:

- HWSSMPL0 and HWSSMPL1 user message exits, for user-written IMS Connect client applications
- HWSJAVA0 user message exit, for the IMS TM Resource Adapter
- HWSSOAP1 user message exit, for the IMS Enterprise Suite SOAP Gateway
- HWSCSLO0 and HWSCSLO1 user message exits, for the OM command clients

All of the IMS Connect client user message exits allow you to call the z/OS TCP/IP IMS Listener security exit routine (IMSLSECX), issue the RACF function in these user message exit routines, or use the IMS Connect user RACF function.

Attention: Do not issue any z/OS calls in the user message exit that result in an MVS WAIT. If you modify the user message exit and add code that results in an MVS WAIT, all work on the TCP/IP PORT will halt until the WAIT has been posted. The user message exits cannot be modified to free any storage passed to the exit, and IMS Connect will not free any storage obtained by the user message exit when the exit returns to IMS Connect. All storage obtained by IMS Connect must be released by IMS Connect and cannot be freed by the User Message Exit without causing failures.

The exit routines that provide general functionality to IMS Connect include:

- The sample IMS Connect Destination Resolution exit routine (HWSYDRU0), which is a modified version of the OTMA Destination Resolution exit routine (DFSYDRU0)
- IMS Connect User Initialization exit routine (HWSUINIT)
- z/OS TCP/IP IMS Listener security exit routine (IMSLSECX)
- IMS Connect Password Change exit routine (HWSPWCH0)
- IMS Connect Event Recorder exit routine (HWSTECL0)
- IMS Connect Port Message Edit exit routine for TCP/IP clients
- IMS Connect DB Routing user exit routine (HWSROUT0)
- IMS Connect DB Security user exit routine (HWSAUTH0)

This topic contains Product-sensitive Programming Interface information.

Related reference:

 [IMS Connect exit routines \(Exit Routines\)](#)

Overview of user message exit routines

For most types of IMS Connect clients, IMS Connect requires the use of a user message exit routine to manage the messages that are received from and sent to the client.

The user message exit routines can perform a number of tasks related to the management of messages, including:

- Translating input messages into the protocol or format required by IMS and the IMS Open Transaction Manager Access (OTMA) component
- Rerouting messages
- Checking security for input messages
- Returning user-defined messages in response to certain user-defined criteria

The following IMS Connect clients are listed with the IMS Connect user message exit that they require:

User-provided clients that access IMS TM

The HWSSMPL0 or HWSSMPL1 user message exit routine, or a user-written user message exit routine.

Optionally, you can also use an IMS Connect Port Message Edit exit routine, which receives control between IMS Connect and the z/OS TCP/IP stack, to modify the format of input and output messages.

IMS TM Resource Adapter

The HWSJAVA0 user message exit routine.

Optionally, you can also use an IMS Connect Port Message Edit exit routine, which receives control between IMS Connect and the z/OS TCP/IP stack, to modify the format of input and output messages.

IMS Enterprise Suite SOAP Gateway

The HWSSOAP1 user message exit routine, which is provided as object code only.

Optionally, you can also use an IMS Connect Port Message Edit exit routine, which receives control between IMS Connect and the z/OS TCP/IP stack, to modify the format of input and output messages.

Clients that access IMS DB, such as the IMS Universal drivers.

IMS Connect does not support user message exit routines for clients that access IMS DB, such as the IMS Universal drivers and user-provided clients that use the Distributed Relational Database Architecture (DRDA) interface. Instead, you can use the following IMS Connect exit routines for message routing, security, and message editing:

- IMS Connect DB Routing user exit routine (HWSROUT0)
- IMS Connect DB Security user exit routine (HWSAUTH0)
- IMS Connect Port Message Edit exit routine

Clients that submit commands to the Operations Manager (OM)

The HWSCSLO0 or HWSCSLO1 user message exit routines, which are provided as object code only.

Related reference:

 [IMS Connect user message exit routines \(Exit Routines\)](#)

Security and the IMS Connect user message exit routines

IMS Connect user message exit routines can perform security checking. If you configure your exit routines to check security, you must provide a security exit or use the z/OS TCP/IP IMS Listener security exit (IMSLSECX).

IMS does not provide a sample security exit due to the many options available for security and the fact that most installations have their own specific security method. The call to RACF is performed by IMS Connect if RACF parameters are provided in the OTMA header when the message exit returns the message.

IMSLSECX is the name of the security exit called by the following IMS Connect user message exit routines:

- HWSSMPL0
- HWSSMPL1
- HWSSOAP1
- HWSCSLO0

If you use HWSSMPL0 or HWSSMPL1, you can change the name of the security exit that they call by changing EXTRN IMSLSECX to a name of your choice. If you change the name of the security exit, you must define the security exit in the HWSSMPL0 or HWSSMPL1 user message exit.

You can also provide the name of the security exit called by HWSJAVA0 and define it in the HWSJAVA0 message exit.

For more information about the IMSLSECX exit routine, see:

- *IMS Version 11 Exit Routines*
- *z/OS Communications Server: IP IMS Sockets Guide*

Related concepts:

“IMS Connect security exit routine” on page 201

“OTMA Resume TPIPE Security exit routine (DFSYRTUX)” on page 202

Related reference:

“HWSSMPL0 and HWSSMPL1 security actions” on page 202

User-defined messages

The IMS Connect HWSJAVA0, HWSSMPL0, and HWSSMPL1 user message exits can return user-defined messages to the IMS Connect client when criteria that you define are met.

Any client application programs that might receive a user-defined message must be able to recognize and process user-defined messages and any associated return and reason codes. When a user-defined message is returned to the client application program, the original input message is not sent to IMS.

When a user message exit returns a user-defined message, the original input message is not passed on

Upon returning one of these user-defined messages, you can also have the user message exit request that IMS Connect either keep the socket connection open or close it, depending on your needs.

User-defined messages can be 1 to 128 characters in length. Any message longer than 128 characters is truncated.

If the message you define is longer than the client input message that is received by the user message exit routine from the IMS Connect, increase the buffer size used by the exit routine by specifying the needed extra bytes in the EXPINI_BUFINC field returned by the exit.

To request that IMS Connect keep a socket connection open after returning a user defined message, the exit routine must set EXPREA_RETCODE to 20 (X'14'). To terminate a socket connection, the exit routine must set EXPREA_RETCODE to 4 (X'04'). Because other factors can cause IMS Connect to terminate a connection, specifying 20 in EXPREA_RETCODE does not guarantee that the socket connection will remain open.

For HWSJAVA0, you can return user-defined message text or only a return code and reason code. To return user-defined message text, the HWSJAVA0 user message exit must set OMUSER_RETCODE to 48 (X'30') and OMUSR_RSNCODE to ICONSUCC. Any other combination of values for OMUSER_RETCODE and OMUSR_RSNCODE returns a return and reason code without user-defined message text.

For HWSSMPL0 and HWSSMPL1, you return user-defined message text by placing the message text in the output message buffer.

IMS Connect does not support user-defined messages for clients that use the Local Option.

Related reference:

➞ User message exit routines HWSSMPL0 and HWSSMPL1 (Exit Routines)

➞ IMS TM Resource Adapter user message exit routine (HWSJAVA0) (Exit Routines)

Overview of function-specific exit routines

IMS provides several IMS Connect exit routines that perform specific functions for IMS Connect for increased flexibility.

You can use the following function-specific exit routines with IMS Connect:

- IMS Connect User Initialization exit routine (HWSUINIT)
- IMS Connect DB Routing user exit routine (HWSROUT0)
- IMS Connect DB Security user exit routine (HWSAUTH0)
- IMS Connect sample OTMA Destination Resolution exit routine (HWSYDRU0)
- z/OS TCP/IP IMS Listener security exit (IMSLSECX)
- IMS Connect Event Recorder exit routine (HWSTECL0)
- IMS Connect Password Change exit routine (HWSPWCH0)

Related reference:

➞ IMS Connect function-specific exit routines (Exit Routines)

Macros that support IMS Connect exit routines

IMS provides macros that support the IMS Connect exit routines.

Macros used for IMS Connect Exit Routines

The macros include:

HWSAUTPM

Maps the parameter list for the IMS Connect DB Security user exit routine (HWSAUTH0). A copy of this macro is in SDFSMACT.

HWSEXPIO

Maps the parameter list for the IMS Connect Port Message Edit exit routine (HWSPIOX0). A copy of this macro is in SDFSMACT.

HWSEXPRM

Maps the parameter list that is passed to the user exit routine on each subroutine call. A copy of this macro is in SDFSMACT. To see the structure, assemble the macro.

HWSOMPFX

Maps the OTMA message prefix format to the output buffer that the user exit routine returns on each READ subroutine call and the input buffer that is passed to the user exit on each XMIT subroutine call. A copy of this macro is in SDFSMACT. To see the structure, assemble the macro.

HWSIMSCB

Maps the IMS request message (IRM) header and BPE header formats used by the HWSSMPL0 and HWSSMPL1 user message exit routines. A copy of this macro is in SDFSMACT. To see the structure, assemble the macro.

HWSIMSEA

Maps the storage area used by the HWSSMPL0 and HWSSMPL1 user message exit routines. A copy of this macro is in SDFSMACT. To see the structure, assemble the macro.

HWSROUPM

Maps the parameter list that is passed to the IMS Connect DB Routing user exit routine (HWSROUT0) on each subroutine call. A copy of this macro is in SDFSMACT. To see the structure, assemble the macro.

HWSXIB

Maps the exit interface block used by IMS Connect user message exit routines and the HWSUINIT exit routine. Contains the addresses of the data store list (HWSXIBDS) and the HWSXIB1 control block used by the IMS Connect DB Routing user exit routine. A copy of this macro is in SDFSMACT. To see the structure, assemble the macro.

HWSXIB1

Maps the exit interface block used by the HWSROUT0 user exit routine. HWSXIB1 contains the address of the ODBM list and optional user data. The HWSXIB1 exit interface block is pointed to by HWSXIB. A copy of this macro is in SDFSMACT. To see the structure, assemble the macro.

HWSXIBDS

Maps the entry in the exit interface block data store list used by the IMS Connect user message exit routines and the HWSUINIT exit routine. The list contains the data store name, the data store availability and status information, and a user field. A copy of this macro is in SDFSMACT. To see the structure, assemble the macro.

HWSXIBOD

Maps the ODBM list that contains the name and status of each ODBM instance known to IMS Connect, as well as a user field and the names and

statuses of the IMS aliases associated with each ODBM instances. The address of HWSXIBOD is stored in the HWSXIB1 exit interface block. A copy of this macro is in SDFS MAC. To see the structure, assemble the macro or refer to the macro prologue.

Exit interface blocks

IMS Connect provides exit interface blocks to support the processing of IMS Connect exit routines that support connections to either IMS DB or IMS TM systems.

XIB exit interface block for connections to IMS TM

IMS Connect provides the XIB exit interface block to support the processing of IMS Connect user message exit routines that are used when connecting to IMS TM. You can use the XIB exit interface block and the user area it includes to store information that is used by your exit routines.

IMS Connect stores the following types of information in the exit interface block:

- Information about the XIBDS.
- Information about the IRL architecture level used.
- User data.

You can also use the XIB_USERAREA field of an exit interface block for any purpose. You can code the IMS Connect User Initialization exit routine (HWSUINIT) or your IMS Connect user message exit routine to use the XIB_USERAREA field.

The exit interface block is mapped by the HWSXIB macro.

Format of the XIB exit interface block

The XIB exit interface block is mapped by the HWSXIB macro.

The XIB exit interface block is provided to support the processing of IMS Connect exit routines. For example, the XIB exit interface block can be used to store the address of a table that an IMS Connect user message exit routine uses during processing.

The following table describes the fields and field offsets of the exit interface block.

Field	Dec Offset	Hex Offset	Length	Value
XIB_HEADER	0	X'00'	0	
XIB_EYE	0	X'00'	4	Eye catcher.
XIB_DATASTORES	4	X'04'	4	The address of the XIBDS data store list.
XIB_UFLD_CNT	8	X'08'	4	The number, in hexadecimal format, of fullword user fields in the XIB_USERAREA.
XIB_XIBDS_LEN	12	X'0C'	2	Length of the XIBDS entries.

Field	Dec Offset	Hex Offset	Length	Value
XIB_ARCHLVL	14	X'0E'	1	Architecture level of the XIB exit interface block. X'01' XIB_ARCH1 - Architecture level 1 X'02' XIB_ARCH2 - Architecture level 2 X'03' XIB_ARCH3 - Architecture level 3 XIB_ARCH3 Highest architecture level
	15	X'0F'	1	Reserved.
XIB_VERSION	16	X'10'	4	The IMS version of the IMS Connect instance. A value of X'0B01000' indicates an IMS Version 11 instance of IMS Connect.
XIB_XIB1	20	X'14'	4	Address of the XIB1 exit interface block for connections to IMS DB.
	24	X'18'	12	Reserved.
XIB_USERAREA	36	X'24'	0F	Start of the user area of the XIB.

XIBDS exit interface block for IMS TM data store information

For connections to IMS TM, IMS Connect keeps track of the status of IMS data stores in entries in the XIBDS exit interface block for data store information. IMS Connect user message exit routines can reference the XIBDS exit interface block to make routing decisions for incoming messages based on the status of the data store.

IMS Connect stores the following types of information about IMS data stores in the exit interface block data store entries:

- Availability of the IMS data stores.
- The state of the IMS data stores; that is, how well the data store is processing messages and, if processing is degraded or completely unavailable, what conditions in the IMS data store might be causing the degraded or unavailable state.
- A time stamp that records the time of the last status change or heartbeat message from OTMA.

IMS Connect updates the state information for an IMS data store when a data store connection is first established and when OTMA notifies IMS Connect of changes in the state of the data store.

OTMA issues a heartbeat message every 60 seconds to indicate that the data store is still communicating. If the time stamp in a data store entry is older than 60 seconds, OTMA could be experiencing problems.

You can also use the XIBDS_USER field of an exit interface block data store entry to for any purpose. You can code the IMS Connect User Initialization exit routine (HWSUINIT) to set the XIBDS_USER field during IMS Connect startup.

The XIBDS exit interface block data store entries are mapped by the HWSXIBDS macro as shown in the following table.

Format of XIBDS exit interface block

The XIBDS exit interface block data store entries are mapped by the HWSXIBDS macro.

The following table describes the fields and field offsets of an entry in the exit interface block.

Field	Dec Offset	Hex Offset	Length	Value
XIBDS_NAME	0	X'00'	8	Name of the data store.
XIBDS_STATUS	8	X'08'	1	Availability of the data store. The possible values are: X'00' The data store is not active. X'01' The data store is active. X'02' The data store is disconnected.
XIBDS_FLAG	9	X'09'	1	Data store entry flags. X'80' Identifies the last entry in the exit interface block.
	10	X'0A'	2	Reserved
XIBDS_USER	12	X'0C'	4	User field
XIBDS_ST_STATUS	16	X'10'	2	The overall state of the IMS data store. The values in this field indicate how well the data store is processing input messages. The possible values are: 3 Normal state: IMS is available and is processing input messages normally. 2 Degraded state: IMS is processing OTMA messages slowly. One or more conditions indicate that IMS is not processing input messages as quickly as it should. 1 Unavailable state: IMS can no longer accept input messages for processing. One or more severe conditions prevent IMS from processing OTMA messages.
XIBDS_ST_SVRSTT			0	The beginning of the fields used to identify the conditions that are contributing to the unavailable state of data store processing.
XIBDS_ST_SVRFLG1	18	X'12'	1	Reserved
XIBDS_ST_SVRFLG2	19	X'13'	1	Reserved
XIBDS_ST_SVRFLG3	20	X'14'	1	Reserved
XIBDS_ST_SVRFLG4	21	X'15'	1	X'01 The data store is flooded with messages from this IMS Connect instance and is no longer accepting input from this instance.
XIBDS_ST_WRNSTT			0	The beginning of the fields used to identify the conditions that are contributing to the degraded state of data store processing.
XIBDS_ST_WRNFLG1	22	X'16'	1	X'80' The global number of messages that are waiting to be processed by the data store has passed 80 percent of the maximum allowable number of waiting messages that is defined for all OTMA clients in this z/OS cross-system coupling facility (XCF) group. If the number of waiting messages reaches 100 percent of maximum allowable number, OTMA sets a flood condition and rejects all incoming messages from all OTMA clients in the XCF group.
XIBDS_ST_WRNFLG2	23	X'17'	1	Reserved

Field	Dec Offset	Hex Offset	Length	Value
XIBDS_ST_WRNFLG3	24	X'18'	1	Reserved
XIBDS_ST_WRNFLG4	25	X'19'	1	X'01' The number of messages that are waiting to be processed by the data store has passed 80 percent of the maximum allowable number for waiting messages that is defined in OTMA for this instance of IMS Connect. If the number of waiting messages reaches 100 percent of the maximum allowable number, OTMA rejects all incoming messages from this instance of IMS Connect.
XIBDS_ST_UTC	26	X'1A'	12	The UTC time at which this status was issued by OTMA

XIB1 exit interface block for connections to IMS DB

The XIB1 exit interface block supports IMS Connect exit routines on connections to IMS DB through the Open Database Manager (ODBM).

The XIB1 exit interface block contains the address of the XIBOD exit interface block, which stores information about ODBM and IMS data stores. You can also use XIB1 exit interface to store user data for use by exit routines, such as the IMS DB Routing user exit routine (HWSROUT0).

You can use the XIB1_USERAREA field of an exit interface block for any purpose. Both the IMS Connect User Initialization exit routine (HWSUINIT) and the IMS Connect DB Routing user exit routine can be coded to use the XIB1_USERAREA field.

The exit interface block is mapped by the HWSXIB1 macro.

Format of the XIB1 exit interface block

The XIB1 exit interface block is mapped by the HWSXIB1 macro.

The following table describes the fields and field offsets of the XIB1 exit interface block.

Field	Dec Offset	Hex Offset	Length	Value
XIB1_HEADER	0	X'00'	0	Align on doubleword.
XIB1_EYE	0	X'00'	4	Eye catcher.
XIB1_ODBMS	4	X'04'	4	Address of the XIBOD exit interface block that contains ODBM and data store information.
XIB1_UFLD_CNT	8	X'08'	4	User field count.
XIB1_ARCHLVL	12	X'0C'	1	Architecture level of the XIB1 exit interface block. X'01' XIB1_ARCH1 - Architecture level 1 XIB1_ARCH1 Highest architecture level
Reserved	13	X'0D'	3	Reserved for IMS Connect.
XIB1_XIB	16	X'10'	4	Address of the XIB.
Reserved	20	X'24'	16	Reserved.
XIB1_USERAREA	36	X'24'	0F	Start of the user area of the XIB1.

XIBOD exit interface block for ODBM and IMS DB data store information

For connections to IMS DB, IMS Connect keeps track of the status of Open Database Manager (ODBM) instances and IMS data stores in entries in the *XIBOD exit interface block for ODBM and data store information*. The IMS Connect DB Routing user exit routines can reference the XIBOD to make routing decisions for incoming messages based on whether the ODBM instances and data stores are active.

For the ODBM instances known to IMS Connect, the XIBOD keeps track of the following ODBM states:

- ODBM is running and connected to IMS Connect.
- ODBM is running but not connected to IMS Connect. You can issue the IMS Connect command STARTOD to establish a connection to ODBM.
- ODBM is not running. ODBM must be restarted before a connection can be made between IMS Connect and ODBM.

The XIBOD block also stores ODBM version information.

Data stores are known to IMS Connect by the alias names that are assigned to the data store in the ODBM configuration member CSLDCxxx during ODBM system definition. For IMS Connect to route an incoming request to a specific data store, the alias name of the data store must be active in both the ODBM in which it is defined and in IMS Connect.

For each alias name defined to the ODBM instances that are known to IMS Connect, the XIBOD keeps track of the following states of the connection to the data store represented by the alias name:

- The connection to the data store is complete. The alias name is active in both the ODBM instance and in IMS Connect (XIBOD_ICACTIVE EQU X'20').
- The connection to the data store is not complete. The alias name is active in ODBM, but has never been activated in IMS Connect (XIBOD_IACTIVE EQU X'80').
- The connection to the data store is not complete. The alias name was active in both ODBM and IMS Connect, but is no longer active in IMS Connect. The alias name is still active in ODBM (XIBOD_ICINACTIVE EQU X'10').
- The connection to the data store is not complete. The alias name was active in both ODBM and IMS Connect, but is no longer active in ODBM. Consequently, the alias name is no longer active in IMS Connect (XIBOD_IINACTIVE EQU X'40').

You can also use the XIBOD_USER field of an exit interface block data store entry for any purpose. You can code the IMS Connect User Initialization exit routine (HWSUINIT) to set the XIBOD_USER field during IMS Connect startup.

The exit interface block data store entries are mapped by the HWSXIBOD macro.

Format of XIBOD exit interface block

The XIBOD exit interface block for ODBM and data store information is mapped by the HWSXIBOD macro.

Map of the ODBM entries in the XIBOD exit interface block

The ODBM entries in the XIBOD exit interface block are mapped by the HWSXIBOD DSECT in the HWSXIBOD macro, as shown in the following table.

Field	Dec Offset	Hex Offset	Length	Value
XIBOD_HDR	0	0	0	
XIBOD_EYE	0	0	4	Character data.
				XIBOD XIBOD_EYEID: The eye catcher.
XIBOD_NAME	4	X'04'	8	Name of the Open Database Manager (ODBM) instance.
XIBOD_OSTATUS	12	X'0C'	1	The status of ODBM. The possible values are:
				X'80' XIBOD_OACTIVE: ODBM is running and connected to IMS Connect.
				X'40' XIBOD_OINACTIVE: ODBM is running but not connected to IMS Connect. You can issue the IMS Connect command STARTOD to establish a connection to ODBM.
				X'20' XIBOD_ODISC: ODBM is not running or is no longer a member of the IMSplex. ODBM must be restarted in the IMSplex before a connection can be made between IMS Connect and ODBM.
XIBOD_ODBMRRS	13	X'0D'	1	The character Y or N.
				Y indicates that ODBM is using z/OS Resource Recovery Services (RRS).
				N indicates that ODBM is not using RRS.
Reserved	14	X'0D'	2	Reserved
XIBOD_ODBMVER	16	X'10'	4	ODBM version number
XIBOD_USER	20	X'14'	4	User field
XIBOD_NEXTODBM	24	X'18'	4	Address of the next ODBM
XIBOD_NEXTATBL	28	X'1C'	4	Address of the table of alias names defined to the ODBM at address XIBOD_NEXTODBM
Reserved	32	X'20'	16	Reserved
XIBOD_IMSatable	48	X'30'	512	Character data. The alias name table has 32 alias name entries. Each alias name entry is 16 bytes long. See the following table for a map of an alias name entry.

Map of the alias name entries in the XIBOD exit interface block

The alias name entries are mapped by the XIBOD_IMSAENT DSECT in the HWSXIBOD macro, as shown in the following table. The alias name table in the XIBOD exit interface block contains 32 alias name entries.

Field	Dec Offset	Hex Offset	Length	Value
XIBOD_IMSA	0	0	4	Character data. The alias name.

Field	Dec Offset	Hex Offset	Length	Value
XIBOD_ISTATUS	4	X'04'	1	The status of the connection to the data store named by the alias.
			X'80'	XIBOD_IACTIVE: The connection to the data store is not complete. The alias name is active in ODBM, but has never been activated in IMS Connect.
			X'40'	XIBOD_IINACTIVE: The connection to the data store is not complete. The alias name was active in both ODBM and IMS Connect, but is no longer active in ODBM. Consequently, the alias name is no longer active in IMS Connect.
			X'20'	XIBOD_ICACTIVE: The connection to the data store is complete. The alias name is active in both the ODBM instance and in IMS Connect.
			X'10'	XIBOD_ICINACTIVE: The connection to the data store is not complete. The alias name was active in both ODBM and IMS Connect, but is no longer active in IMS Connect. The alias name is still active in ODBM.
			X'08'	XIBOD_DELETED: The IMS alias name has been deleted from the ODBM CSLDCxxx configuration member and can no longer be used.
Reserved	5	X'05'	3	Reserved for IMS Connect.
Reserved	8	X'08'	8	Reserved for IMS Connect.

Related reference:

[➡](#) IMS Connect UPDATE ODBM command (Commands)

[➡](#) STARTOD command (Commands)

Chapter 13. IMS Connect support for IMSplex and shared queues

IMS Connect can send and receive IMS Operations Manager (OM) commands and response string messages between a TCP/IP client and the OM in an IMSplex by using the IMS Structured Call Interface (SCI).

IMS Connect can send and receive IMS Operations Manager (OM) commands and response string messages between a TCP/IP client and the OM in an IMSplex by using the IMS Structured Call Interface (SCI). IBM Management Console for IMS and DB2 for z/OS is an IBM-provided client that uses this interface.

The following subsections provide detailed information about the environment requirements and how to set up IMS Connect to support OM.

IMS Connect support for IMSplex

The IMS Connect support for IMSplex enables access to an OM in an IMSplex from IBM Management Console for IMS and DB2 for z/OS (Management Console) or other OM command client through TCP/IP.

The IMSplex support accesses OM through the IMS SCI. The IMSplex statement in the IMS Connect configuration file (HWSCFGxx) defines IMS Connect for IMSplex support. If the IMSplex statement is omitted, then IMSplex support is not available.

IMSplex support sends IMS command string messages directly for a client (for example, the Management Console) to a selected OM within an IMSplex. One or more IMSplexes can be defined to IMS Connect to receive Management Console command messages. SCI is used to communicate between IMS Connect and the IMSplex. To gain access to the selected OM, include the IMSPLEX configuration statement in the IMS Connect configuration member in the IMS.PROCLIB data set.

The same security method (authentication of the userid, groupid, and password) used for accessing data stores also applies to IMSplex support. An IMS Connect command message exit performs similar functions as an IMS Connect user message exit.

The OM command client exit routines (HWSCSLO0 and HWSCSLO1) are designed specifically for IMSplex support. The HWSCSLO0 and HWSCSLO1 exit routines are similar to the other user message exit routines provided by IMS Connect for client access to the data store. The HWSCSLO0 and HWSCSLO1 message exits are designed to be used only by OM command clients and cannot be used by any client that sends messages to a data store. The HWSCSLO0 and HWSCSLO1 exit routines process only command string messages from OM command clients and are delivered as object code only.

IMSplex support environment

IMS Connect requires that IMS, the CSL Operations Manager (OM), and the CSL Structured Call Interface (SCI) be running to communicate with OM.

IMS must be running in the same z/OS image as IMS Connect or a different z/OS image on the same IMSplex. The OM must be in the same z/OS image or a different z/OS image in the same IMSplex. The SCI must be in the same z/OS image as IMS Connect.

IMS Connect can be brought up before or after IMS, SCI, OM, and RM. During IMS Connect initialization, connection to SCI is made. IMS Connect attempts to connect to SCI for 30 minutes. If SCI connection is not made, then an OPENIP command will need to be issued to connect to the SCI after the SCI has been initialized. If the SCI terminates normally or abnormally, IMS Connect will automatically reconnect to the SCI when the SCI is restarted.

Related concepts:

 IMS system administration considerations and tasks (System Administration)

Installing IMS Connect support for IMSplex

The IMS Connect support for an IMSplex requires the CSL Operations Manager (OM), the CSL Structured Call Interface (SCI), and TCP/IP.

IMS must be installed to use the IMS Connect IMSplex support. Ensure that the IMS.SDFSRESL data set has been added to the STEPLIB to enable access to SCI for IMS Connect IMSplex support.

OM requires the appropriate IMSplex component (IMS, IMS Connect, RM, or ODBM) to be running to provide the OM command capability. IMS Connect IMSplex support does not require these components to be running. However, if these components are not running, none of the commands can be processed.

Install or modify the following components of the IMS Connect support for IMSplex in the following order:

1. The IMS Connect configuration file:
 - a. Add the IMSplex statement.
 - b. Add the HWSCSLO0 exit to the TCPIP statement EXIT= parameter.
2. The BPE configuration file:
 - a. Add OMDR and HWSO statements, if the IMS Connect trace entries are listed separately.
3. The OM command client:
 - a. Identify the IMS Connect HWS ID= value.
 - b. Identify the IMS Connect IMSPLEX tmember= value.

When installation is complete, start the following features in order:

1. Start SCI.
2. Start OM.
3. If you are using RM, start RM.
4. Start IMS.

You can start IMS Connect before, during, or after the steps listed above. IMS Connect attempts to connect to SCI for 30 minutes. If SCI is not brought up or the connection attempt fails, you must issue an OPENIP *imsplex_name* command.

The following IMS Connect commands are enabled by IMSplex support:

- STOPIP

- OPENIP
- VIEWIP

Retrieving ALTPCB output in a shared queues environment

In a shared queues environment, to retrieve messages inserted to an ALTPCB output queue by IMS systems to which IMS Connect is not directly connected, use the OTMA super member function.

To use the OTMA super member in a configuration that includes multiple IMS systems, you must have shared queues enabled for all those IMS systems. If a back-end IMS system creates ALT-PCB output in the super member-enabled environment, the output can be retrieved from any front-end IMS with an OTMA client in the super member set.

To activate the super member function:

Specify a 1- to 4-character super member name in the SMEMBER parameter on either the HWS statement or the DATASTORE statement in the IMS Connect configuration member (HWSCFGxx). Super member names must be unique and cannot be the same as existing OTMA member names.

Related information:

Chapter 14. IMS Connect security support

IMS Connect includes a variety of options for implementing and modifying the security checking performed on messages as they arrive in IMS Connect and, for IMS TM connections, as they arrive at the data store.

For connections to IMS TM, IMS Connect provides two options for checking security within IMS Connect: you can configure IMS Connect to call RACF directly or you can have the IMS Connect user message exit routines call a Security user exit routine.

For connections to IMS DB, the IMS Connect DB Security user exit routine (HWSAUTH0) authenticates the user and you can use RACF as well. For connections to IMS DB, IMS Connect does not check the authority of the user to perform any action, but can pass a RACO token to the CSL Open Database Manager for the purposes of authorization.

Additional security features provided by IMS Connect include:

- Password management support
- Trusted-user classification for messages arriving at the data store
- OTMA accessor environment element (ACEE) timeout specification support
- For connections from clients connecting to IMS TM, Secure Sockets Layer (SSL) support. To use SSL to secure connections from clients connecting to IMS DB, you can use IBM z/OS Communications Server Application Transparent Transport Layer Security feature (AT-TLS).

IMS Connect support for RACF

IMS Connect can be configured to call RACF directly and to support RACF PassTickets.

By default, IMS Connect does not call RACF. When IMS Connect is configured to call RACF, IMS Connect can validate the user IDs and passwords on incoming messages with RACF directly.

When configured for direct RACF support, IMS Connect also supports RACF PassTickets.

If RACF is configured to support mixed-case passwords, you can also configure IMS Connect to support mixed-case password support in IMS Connect.

IMS Connect calls RACF by issuing the RACF command RACROUTE REQUEST=VERIFY to verify the user IDs and passwords received from clients in the IRM of incoming messages. You can also define a default RACF ID for IMS

If a RACF security failure occurs, IMS Connect includes the return code from the RACROUTE REQUEST=VERIFY command in the request status message (RSM) for diagnostic purposes.

If you configure IMS Connect to call RACF, evaluate the impact of the RACF calls on IMS Connect performance.

Related tasks:

“IMS Connect password management” on page 208

Related reference:

“HWSSMPL0 and HWSSMPL1 security actions” on page 202

Enabling RACF security checking in IMS Connect

You can enable RACF security checking in IMS Connect either by specifying RACF=Y in the HWS configuration statement or by issuing an online IMS Connect command.

To enable RACF in the HWS configuration statement, add the RACF=Y parameter. For example:

```
HWS ID=HWS01 RACF=Y
```

To enable RACF online:

- Issue any one of the following IMS Connect commands:
 - The WTOR format command, SETRACF ON
 - The z/OS MODIFY command format, UPDATE MEMBER TYPE(IMSCON) SET(RACF(ON))
- Verify that RACF security checking is enabled by issuing any of the following IMS Connect commands:
 - The WTOR format command, VIEWHWS
 - The z/OS MODIFY command format, QUERY MEMBER TYPE(IMSCON) SHOW(ALL)

The output from both the WTOR and z/OS MODIFY commands includes a line similar to the following: HWSC0001 HWSID=HW01 RACF=Y.

Related reference:

-  [IMS Connect UPDATE MEMBER command \(Commands\)](#)
-  [SETRACF command \(Commands\)](#)
-  [HWS statement \(System Definition\)](#)
-  [IMS Connect QUERY MEMBER command \(Commands\)](#)
-  [VIEWHWS command \(Commands\)](#)

IMS Connect default RACF user ID

You can set a default RACF user ID for IMS Connect to use when the IRM of an input message either does not contain the IRM_RACF_USERID field or the IRM_RACF_USERID field is blank.

When the default RACF user ID is used, IMS Connect passes it in the OMSECUID field of the input message to OTMA. When OTMA security checking is enabled, OTMA uses the RACF user ID for authorizing commands, transactions, and RESUME TPIPE calls with RACF.

When both a default RACF user ID is defined and the IRM_RACF_USERID field of an incoming message is not blank, IMS Connect uses the user ID defined in the IRM_RACF_USERID field.

Define the default RACF user ID in the RACFID parameter of the IMS Connect TCPIP configuration statement.

IMS Connect security for clients of IMS DB

For clients that connect to IMS DB through ODBM, such as the IMS Universal drivers and clients using the Distributed Relational Database Architecture (DRDA), IMS Connect authenticates the user, but does not check the authority of the user to perform any actions.

To authenticate a user ID for an IMS DB client, IMS Connect can use the IMS Connect DB Security user exit routine (HWSAUTH0), a security product such as RACF, or both.

IMS Connect always calls the HWSAUTH0 user exit, regardless of whether RACF or another security product is enabled. If RACF support is included in your IMS Connect configuration, IMS Connect calls the HWSAUTH0 user exit before invoking RACF.

The HWSAUTH0 user exit routine can override the input user ID with a different user ID and can provide a RACF group ID to be authenticated further by IMS Connect.

The HWSAUTH0 user exit routine is a BPE type-1 user exit routine and is refreshable.

IMS Connect does not support Secure Sockets Layer (SSL) directly for clients that connect to IMS DB. To secure connections to IMS DB with SSL, use IBM z/OS Communications Server Application Transparent Transport Layer Security feature (AT-TLS). The use of AT-TLS is transparent to IMS Connect.

Related concepts:

 [IMS Connect definition and tailoring \(System Definition\)](#)

Related tasks:

[“Enabling RACF security checking in IMS Connect” on page 200](#)

IMS Connect security exit routine

If any IMS Connect user message exit routine performs security checking, you must provide a security exit routine or use the z/OS TCP/IP IMS Listener security exit routine (IMSLSECX).

IMS does not provide a sample security exit routine due to the many options available for security and the fact that most installations have their own specific security method. The call to RACF is performed by IMS Connect if RACF parameters are provided in the OTMA header when the message exit routine returns the message.

The name of the security exit routine called by HWSSMPL0, HWSSMPL1, HWSSOAP1, or HWSCSLO0 is IMSLSECX.

If you use HWSSMPL0 or HWSSMPL1, you can change the name of the security exit routine that they call by changing EXTRN IMSLSECX to a name of your choice. If you change the name of the security exit routine, you must define the security exit routine in the HWSSMPL0 or HWSSMPL1 user message exit.

You can also provide the name of the security exit routine called by HWSJAVA0 and define it in the HWSJAVA0 message exit routine.

OTMA Resume TPIPE Security exit routine (DFSYRTUX)

The OTMA Resume TPIPE Security exit routine (DFSYRTUX) is not an IMS Connect exit routine, but it is one of two possible methods that you can use to secure messages queued on the OTMA asynchronous hold queue. The other method is to use an external security product, such as RACF. You can use the DFSYRTUX exit routine and an external security product each by itself or in combination.

The DFSYRTUX exit routine runs in the IMS control region and not in the IMS Connect address space.

Related tasks:

“Securing messages on the asynchronous hold queue” on page 804

Related reference:

 OTMA Resume TPIPE Security exit routine (DFSYRTUX) (Exit Routines)

HWSSMPL0 and HWSSMPL1 security actions

The sample user message exits HWSSMPL0 and HWSSMPL1 always perform certain security actions and perform other security actions only when the IMSLSECX security exit is or is not called.

Security actions when HWSSMPL0 and HWSSMPL1 do not call IMSLSECX

The following tables define the action that HWSSMPL0 and HWSSMPL1 take when they do not call the security exit (IMSLSECX).

Table 6. USERID results if security exit not called

	USERID field present in IRM	IRM USERID field blank/null	RACF parms results passed in OTMA security header
USERID	Yes	Yes	Default RACFID
USERID	Yes	No	IRM USERID
USERID	No	N/A	Default RACFID

Table 7. GROUPID results if security exit not called

	GROUPID field present in IRM	IRM USERID field blank/null	RACF parms results passed in OTMA security header
GROUPID	Yes	Yes	Blanks/nulls
GROUPID	Yes	No	IRM GROUPID
GROUPID	No	N/A	Blanks/nulls

Table 8. Password results if security exit not called

	Password field present in IRM	IRM PASSWORD field blank/null	RACF parms results passed in OTMA security header
PASSWORD	Yes	Yes	Blanks/nulls

Table 8. Password results if security exit not called (continued)

	Password field present in IRM	IRM PASSWORD field blank/null	RACF parms results passed in OTMA security header
PASSWORD	Yes	No	IRM PASSWORD
PASSWORD	No	N/A	Blanks/nulls

Security actions when HWSSMPL0 and HWSSMPL1 call IMSLSECX

The following tables define the action that HWSSMPL0 and HWSSMPL1 take when they call the security exit (IMSLSECX).

Table 9. USERID results if security exit called; returns blank or non-blank USERID

	USERID field present in IRM	IRM USERID field blank/null	Security exit return USERID	RACF parms results passed in OTMA security header
USERID	Yes	Yes	No	Default RACF USERID
USERID	Yes	Yes	Yes	Security exit returned USERID
USERID	Yes	No	No	USERID passed in IRM
USERID	Yes	No	Yes	Security exit returned USERID
USERID	No	N/A	No	Default RACF USERID
USERID	No	N/A	Yes	Security exit returned USERID

Table 10. GROUPID results if security exit called; returns non-blank USERID

	GROUPID field present in IRM	IRM GROUPID field blank/null	Security exit return GROUPID	RACF parms results passed in OTMA security header
GROUPID	Yes	Yes	No	Blank GROUPID
GROUPID	Yes	Yes	Yes	Security exit returned GROUPID
GROUPID	Yes	No	No	Blank GROUPID
GROUPID	Yes	No	Yes	Security exit returned GROUPID
GROUPID	No	N/A	No	Blank GROUPID
GROUPID	No	N/A	Yes	Security exit returned GROUPID

Important: If the security exit returns a blank USERID, as shown in the following table, then the GROUPID that is returned by the exit is not used.

Table 11. GROUPID results if security exit called; returns blank USERID

	GROUPID field present in IRM	IRM GROUPID field blank/null	Security exit return GROUPID	RACF parms results passed in OTMA security header
GROUPID	Yes	Yes	No	Blank GROUPID
GROUPID	Yes	Yes	Yes	Blank GROUPID
GROUPID	Yes	No	No	IRM GROUPID
GROUPID	Yes	No	Yes	IRM GROUPID
GROUPID	No	N/A	No	Blanks
GROUPID	No	N/A	Yes	Blanks

Security actions that HWSSMPL0 and HWSSMPL1 always perform

The following table defines the action that HWSSMPL0 and HWSSMPL1 take regardless of whether the security exit (IMSLSECX) is called. The password is based on the IRM, not on the security exit.

Table 12. Password results regardless of whether security exit called

	PASSWORD field present in IRM	PASSWORD field blank/null	Security exit return PASSWORD	RACF parms results passed in OTMA security header
PASSWORD	Yes	Yes	N/A	Blanks/nulls
PASSWORD	Yes	No	N/A	IRM PASSWORD
PASSWORD	No	N/A	N/A	Blanks/nulls

IMS Connect responses to errors on RACROUTE calls from the sample exits

If an error occurs on a RACROUTE call from one of the sample user message exits, HWSSMPL0 or HWSSMPL1, IMS Connect decides what error actions to take depending on the RACF parameter setting in the IMS Connect configuration file, as well as the specific circumstances that cause the error.

The different actions IMS Connect takes are described in the following series of tables.

IMS Connect response when RACF=Y and RACROUTE call parameters are in error

The following table describes the error actions that IMS Connect takes if RACF=Y, based on required RACROUTE call parameters.

Table 13. IMS Connect error actions taken based on RACROUTE call parameters (RACF=Y)

USERID	PASSWORD	GROUPID	Action taken
Non-blanks	Non-blanks	Non-blanks	RACROUTE call issued

Table 13. IMS Connect error actions taken based on RACROUTE call parameters (RACF=Y) (continued)

USERID	PASSWORD	GROUPLD	Action taken
Non-blanks	Blanks	Blanks	• Error message HWSP1503 issued • Input rejected • RACROUTE call not issued • Set OMUSR_RETCODE='X'04' • Set OMUSR_RESCODE='SECFNOPW' • Password cleared in OTMA header • * Security failed, no password *
Non-blanks	Blanks	Non-blanks	• Error message HWSP1503 issued • Input rejected • RACROUTE call not issued • Set OMUSR_RETCODE='X'04' • Set OMUSR_RESCODE='SECFNOPW' • Password cleared in OTMA header • * Security failed, no password *
Blanks	Non-blanks	Blanks	• Error message HWSP1503 issued • Input rejected • RACROUTE call not issued • Set OMUSR_RETCODE='X'04' • Set OMUSR_RESCODE='SECFNUID' • Password cleared in OTMA header • * Security failed, no password *
Blanks	Non-blanks	Non-blanks	• Error message HWSP1503 issued • Input rejected • RACROUTE call not issued • Set OMUSR_RETCODE='X'04' • Set OMUSR_RESCODE='SECFNUID' • Password cleared in OTMA header • * Security failed, no password *
Blanks	Blanks	Non-blanks	• Error message HWSP1503 issued • Input rejected • RACROUTE call not issued • Set OMUSR_RETCODE='X'04' • Set OMUSR_RESCODE='SECFNPUT' • Password cleared in OTMA header • * Security failed, no password *
Blanks	Blanks	Blanks	• Error message HWSP1503 issued • Input rejected • RACROUTE call not issued • Set OMUSR_RETCODE='X'04' • Set OMUSR_RESCODE='SECFNPUT' • Password cleared in OTMA header • * Security failed, no password *

Table 13. IMS Connect error actions taken based on RACROUTE call parameters (RACF=Y) (continued)

USERID	PASSWORD	GROUPIP	Action taken
Non-blanks	Blanks	Non-blanks	• Error message HWSP1503 issued • Input rejected • RACROUTE call not issued • Set OMUSR_RETCODE=X'04' • Set OMUSR_RESCODE='SECFNOPW' • Password cleared in OTMA header • * Security failed, no password *

IMS Connect responses when RACF=Y and either RACROUTE call fails or OTMA header data is in error

The following table describes the error actions that IMS Connect takes if RACF=Y, either based on OTMA header data, or should the RACROUTE call fail.

Table 14. IMS Connect actions taken for RACROUTE call failure or OTMA header data error and RACF=Y

RACROUTE call failure or OTMA header data	Action taken
No security header	• Error message HWSP1503 issued • Input rejected • Set OMUSR_RETCODE=X'04' • Set OMUSR_RESCODE='NOSECHDR' • Password cleared in OTMA header • No RACF call made
Security header < X'6A'	• Error message HWSP1503 issued • Input rejected • Set OMUSR_RETCODE=X'04' • Set OMUSR_RESCODE='INVSECHL' • Password cleared in OTMA header • No RACF call made
Conversation continued	• No error message issued • Input accepted • Password cleared in OTMA header • No RACF call made
Response message	• No error message issued • Input accepted • Password cleared in OTMA header • No RACF call made
UTOKEN present	• No error message issued • Input accepted • Password cleared in OTMA header • No RACF call made

Table 14. IMS Connect actions taken for RACROUTE call failure or OTMA header data error and RACF=Y (continued)

RACROUTE call failure or OTMA header data	Action taken
RACROUTE call failed	• Error message HWSP1500 issued • Input rejected • Set OMUSR_RETCODE=X'04' • Set OMUSR_RESCODE='SECFAIL' • Password cleared in OTMA header • RACF return/reason codes in HWSP1500 message
All others	See Table 13 on page 204

IMS Connect response when RACF=N and RACROUTE call parameters are in error

The following table describes the error actions that IMS Connect takes if RACF=N, based on required RACROUTE call parameters.

Table 15. IMS Connect error actions taken based on RACROUTE call parameters (RACF=N)

USERID	PASSWORD	GROUPID	Action taken
Non-blanks	Non-blanks	Non-blanks	• Password cleared • Bypass RACROUTE call • Pass these parameters to OTMA
Non-blanks	Blanks	Blanks	• Password cleared • Bypass RACROUTE call • Pass these parameters to OTMA
Non-blanks	Blanks	Non-blanks	• Password cleared • Bypass RACROUTE call • Pass these parameters to OTMA
Blanks	Non-blanks	Blanks	• Password cleared • Bypass RACROUTE call • Pass these parameters to OTMA
Blanks	Non-blanks	Non-blanks	• Password cleared • Bypass RACROUTE call • Pass these parameters to OTMA
Blanks	Blanks	Non-blanks	• Password cleared • Bypass RACROUTE call • Pass these parameters to OTMA
Blanks	Blanks	Blanks	• Password cleared • Bypass RACROUTE call • Pass these parameters to OTMA
Non-blanks	Blanks	Non-blanks	• Password cleared • Bypass RACROUTE call • Pass these parameters to OTMA

IMS Connect response when RACF=N and RACROUTE call fails or OTMA header data is in error

The following table describes the error actions that IMS Connect takes if RACF=N, either based on OTMA header data, or should the RACROUTE call fail.

Table 16. IMS Connect error actions taken for RACROUTE call failure or OTMA header data (RACF=N)

RACROUTE call failure or OTMA header data	Action taken
No security header	• Password cleared • Bypass RACROUTE call • Pass OTMA headers and data to IMS OTMA
Security header < X'6A'	• Password cleared • Bypass RACROUTE call • Pass OTMA headers and data to IMS OTMA
Conversation continued	• Password cleared • Bypass RACROUTE call • Pass OTMA headers and data to IMS OTMA
Response message	• No error message issued • Input accepted • Password cleared in OTMA header • No RACF call made
UTOKEN present	• Password cleared • Bypass RACROUTE call • Pass OTMA headers and data to IMS OTMA
RACROUTE call failed	• Password cleared • Bypass RACROUTE call • Pass OTMA headers and data to IMS OTMA
All others	• Password cleared • Bypass RACROUTE call • Pass OTMA headers and data to IMS OTMA

IMS Connect password management

IMS Connect provides several features to help you manage RACF passwords. Some of these features only apply when IMS Connect is configured to call RACF directly.

Changing RACF passwords by using client messages

When IMS Connect is configured to call RACF directly, users of the user message exit routines HWSSMPL0, HWSSMPL1, and HWSJAVA0 can change RACF passwords by submitting a client message that includes a password change request keyword.

To enable this feature, you must bind the HWSPWCH0 object code with the user message exit routine you are using. The HWSPWCH0 object code is stored in the IMS.ADFSLOAD member of the distribution library.

The password change request keyword must appear at the beginning of the application data section of the message and be followed by a blank, the old password, the new password, and the new password again.

For example, for the HWSSMPL0 and HWSSMPL1 user message exit routines, a password change request message has the following format:

```
1111IRM11zzHWSPWCH old-password/new_password/new_password|11zz
```

For the HWSJAVA0 user message exit routine, a password change request message has the following format:

```
1111IRMOTMALLzzHWSPWCH old_password/new_password/new_password|11zz
```

The password change request keyword that is defined in the sample user message exit routines is “HWSPWCH.” You can change this keyword by modifying the user message exit routine. Any password change request keyword that you define must be followed by a blank as a delimiter.

IMS Connect returns a response message to the client application in one of the following formats:

- For the user message exit routine HWSSMPL0, the format is:
11zzmessage_textCSM11zz
- For the user message exit routine HWSSMPL1, the format is:
1111message_textCSM11zz
- For the user message exit routine HWSJAVA0, the format is:
11110TMAheadermessage_text11zz

The communication sequence for a password change request is:

1. Connect.
2. Send the password change request.
3. Receive message HWSC00 xy , where xx is the final two digits of the message number and y is the message type identifier. If you are using HWSSMPL0 or HWSSMPL1, the user-written client application receives the message HWSC00 xy appended with *CSMOKY*. If you are using HWSJAVA0, the IMS TM Resource Adapter client application receives the HWSC00 xy message in the application data portion of the OTMA header and the appropriate return and reason codes in the user data portion of the OTMA header.
4. Resolve socket connection. IMS Connect maintains the socket connection, except in the following circumstances:
 - If the socket is not persistent, IMS Connect disconnects the socket.
 - If IMS Connect expects an ACK or NAK response from the client application, IMS Connect disconnects the socket.
 - If the client application is in IMS conversational mode, IMS Connect disconnects the socket.

Related reference:

 [IMS TM Resource Adapter user message exit routine \(HWSJAVA0\) \(Exit Routines\)](#)

 [User message exit routines HWSSMPL0 and HWSSMPL1 \(Exit Routines\)](#)

Enabling mixed-case password support

IMS Connect supports mixed-case passwords. To use mixed-case passwords, RACF (or your similar security product) must also support mixed-case passwords.

By default, IMS Connect support for mixed-case passwords is determined by the configuration of the security product in use. For example, if RACF supports mixed-case passwords, IMS Connect automatically supports mixed-case passwords also.

If IMS Connect cannot determine if the security product supports mixed-case passwords, IMS Connect translates all passwords to uppercase, unless instructed to do otherwise during system definition or by an online command while IMS Connect is running.

During IMS Connect system definition, you can enable, disable, or accept the RACF specification for mixed-case password support by using the PSWDMC parameter in the HWS configuration statement.

In the online system, you can enable, disable, or accept the RACF specification for mixed-case password support by using any of the following commands:

- The IMS Connect WTOR command SETPWMC ON | OFF | RCF
- The IMS Connect z/OS command UPDATE MEMBER TYPE(IMSCON) SET(PSWDMC(ON | OFF | RCF))

You can view the current setting for mixed-case password support by issuing any of the following commands:

- The IMS Connect WTOR command VIEWHWS
- The IMS Connect z/OS command QUERY MEMBER

Related reference:

-  [IMS Connect UPDATE MEMBER command \(Commands\)](#)
-  [IMS Connect QUERY MEMBER command \(Commands\)](#)
-  [SETPWMC command \(Commands\)](#)
-  [VIEWHWS command \(Commands\)](#)
-  [HWS statement \(System Definition\)](#)

IMS Connect support for RACF PassTicket

An alternative to the RACF password is a PassTicket. PassTicket allows you to communicate with a host without using a RACF password. When IMS Connect is configured to call RACF directly, you can use PassTicket to authenticate user IDs and log on to computer systems that contain RACF.

You can select PassTicket support through an IMS Connect client and send a PassTicket in the IRM in place of a RACF password. IMS Connect issues a RACF call using PassTicket and blanks out the PassTicket field in the OTMA user data header before sending the message to IMS. Because PassTicket occupies the same field as the RACF password and PassTicket cannot be translated to uppercase, the RACF password is also not translated to uppercase. You can use a user message exit to provide uppercase translation.

The IMS Connect PassTicket support parallels IMS PassTicket support.

- You can use existing APPLname definitions for newly connecting IMS Connect clients.
- Each DATASTORE statement has a parameter APPL=APPLname, where:
 - Each APPL= can be a unique RACF APPLname for each data store.

- Each APPL= can be the same name for each data store, as required for VGR support, or can be unique per data store.
- The default APPL=APPLname value is blank.
- The IMS Connect client can pass an APPLname in the IRM to the user message exit which sets the APPLname in the OTMA user data header or the user message exit can pass and set the appropriate APPLname in the OTMA user data header.

For PassTicket support, you are responsible for all definitions to RACF. You need to establish the RACF encoding and decoding routines and to supply the encoding routine to the distributed platform.

RACF PassTicket is only supported for customer-written IMS Connect client applications. The IMS TM Resource Adapter does not currently support RACF PassTicket.

This support might require changes to the customer-written user message exits and customer-written client application code. The following list describes options you can select for PassTicket support:

- **Support for passing an APPLname in the IRM to IMS Connect**

This support has been added to the IRM definition. A new 8 byte field, IRM_APPL_NM, has been added to the end of the IRM structure. If you want to implement the PassTicket function, then the client code must pass the APPLname to IMS Connect in this field.

Note: This **will** change the length of the IRM by 8 bytes and the total length of the message by 8 bytes.

The supplied user message exits (HWSSMPL1 and HWSSMPL0) have been modified so that a client can send an APPLname to IMS Connect in the IRM_APPL_NM field.

If you choose this option, you need only to pass the APPLname in the IRM. HWSIMSCB and IMS Connect have been modified to support this function.

- **No support for passing an APPLname in the IRM to IMS Connect**

This support has been added to the IRM definition. A new 8 byte field IRM_APPL_NM has been added to the end of the IRM structure. If you do not want to implement the PassTicket function, you have two options:

- **Option 1: Blank APPLname**

You can choose to pass a blank APPLname to IMS Connect in the IRM_APPL_NM field to IMS Connect.

Note: This **will** change the length of the IRM by 8 bytes and the total length of the message by 8 bytes.

The supplied user message exits (HWSSMPL1 and HWSSMPL0) have been modified so that a client can send a blank APPLname in the IRM_APPL_NM field to IMS Connect.

If you choose this option, you need only to pass a blank APPLname in the IRM. HWSIMSCB and IMS Connect have been modified to support this blank APPLname function.

- **Option 2: No APPLname**

The customer can choose to pass no APPLname to IMS Connect in the IRM_APPL_NM field to IMS Connect.

Note: This **will not** change the length of the IRM or the total length of the message.

The supplied user message exits (HWSSMPL1 and HWSSMPL0) have been modified so that a client does not have to send an APPLname in the IRM_APPL_NM field to IMS Connect.

If you choose this option, you do not need to perform any action. HWSIMSCB and IMS Connect have been modified to support this function of not passing an APPLname.

The APPLname is always passed to RACF. This is true even if PassTickets are not used. As a result, the APPL= keyword on the DATASTORE statement can be used to verify a user's authority to access an IMS Connect data store, even if PassTickets are not used.

PassTicket replay protection considerations

You can bypass PassTicket replay protection, which you might do if you have multiple end-users sharing the same user ID.

If you have multiple users with the same user IDs, it is possible for them to request access to an application during the same time interval. In this situation, the same PassTicket is generated for different users. As a result, if PassTicket replay protection is not bypassed, the users will be using the same PassTicket and be denied access to the application. Bypassing the PassTicket replay protection allows the same PassTicket to be used by multiple users.

Similarly, if you are stress testing your system where there is no think time driving requests to IMS Connect and have numerous requests to the same application occurring in the same time interval, you may want to consider bypassing PassTicket replay protection. This option allows the same PassTicket to be used within a ten minute period.

You can specify NO REPLAY PROTECTION in the APPLDATA field of the PTKTDATA profile for one or more of the selected applications to allow the same PassTicket to be generated within a ten minute period.

For additional information about no replay options, see “Protecting General Resources” in the *z/OS Security Server RACF Security Administrator's Guide*.

IMS Connect SSL connections

To use Secure Sockets Layer (SSL) with IMS Connect, you can either configure the IMS Connect support for SSL by using the IMS Connect SSL interface or you can use the IBM z/OS Communications Server Application Transparent Transport Layer Security feature (AT-TLS).

If you use AT-TLS, the use of SSL is transparent to IMS Connect and you do not need to configure the IMS Connect SSL interface. The configuration and administration of SSL is performed by your z/OS TCP/IP administrator.

In addition to simplifying IMS Connect security implementation, AT-TLS provides greater flexibility with respect to the use of ports. IMS Connect support for SSL is restricted to only a single port. If more than one SSL port is defined on the IMS Connect TCPIP configuration statement, IMS Connect abends during startup.

However, if AT-TLS is used, connections that are secured by SSL are unencrypted by AT-TLS before they reach IMS Connect and can therefore use any IMS Connect port.

Recommendation: Use the z/OS Communication Server AT-TLS feature to enable SSL on TCP/IP connections to IMS instead of configuring the IMS Connect support for SSL.

The IMS Connect SSL interface includes support for:

- SSL Version 2.0
- SSL Version 3.0
- Transport Layer Security (TLS) Version 1.0
- Transport Layer Security (TLS) Version 1.1
- Transport Layer Security (TLS) Version 1.2
- FIPS mode

Throughout this topic, the term SSL is used to describe both the SSL and TLS protocols. This topic applies only to the configuration of IMS Connect support for SSL.

SSL and TLS protect the privacy and integrity of data that is transferred through a network. SSL rests on top of TCP/IP to provide a mechanism for secure sockets. SSL uses a combination of public and private keys for initial contact and authentication between the client and servers and symmetric key encryption for the subsequent communication flows.

During the initial handshake protocol, the client and server agree on how to encrypt and decrypt information and define the format used to transmit the encrypted data. After the SSL connection is established, the encryption and decryption of communication data flows by using symmetric key encryption is transparent to the client and server.

Because the SSL encryption is transparent to the client and server programs, the message formats they send and receive are the same regardless of whether the connection is an SSL socket or a regular unencrypted socket.

When you are setting up IMS Connect SSL configuration, you must consider any limitations or restrictions of the client's SSL support. TCP/IP consistently and reliably transfers information across the internet domain, but it does not secure the information that is transferred.

SSL protects information from:

- Eavesdropping
- Data theft
- Traffic analysis
- Data modification
- Trojan horse browser / server

X.509 certificates are used by both the client and server when securing communications. The client must verify the server's certificate based on the certificate of the Certificate Authority (CA) that signed the certificate or based on a self-signed certificate from the server. The server must verify the client's certificate

(if requested) using the certificate of the CA that signed the client's certificate. The client and the server then use the negotiated session keys and begin encrypted communications.

Related concepts:

➡ IMS Connect definition and tailoring (System Definition)

➡ z/OS: Application Transparent Transport Layer Security data protection

Related reference:

➡ z/OS: Cryptographic Services System SSL Programming

z/OS key management

SSL connections use public and private key mechanisms for authenticating each side of the SSL session (the server side and the client side) and agree on bulk encryption keys to be used for the SSL session.

Note: Authenticating both the client and server during SSL handshaking is not always necessary. In many cases, you can choose to generate a public and private key pair only on the server side to facilitate server authentication during SSL handshaking. The fact that the client sends a user ID and password over the secured connection after the SSL handshaking completes can be considered authentication of the client.

To use public and private key mechanisms (PKIs), public and private key pairs must be generated. In addition, X.509 certificates (which contain public keys) might either need to be created, or certificates must be requested, received, and managed.

SSL for z/OS supports the following two methods for managing PKI private keys and certificates:

- A z/OS shell-based program named *gskkyman*. *gskkyman* creates, fills in, and manages a z/OS HFS file that contains PKI private keys, certificate requests, and certificates. This z/OS HFS file is called a key database and, by convention, has a file extension of *.kdb*.
- The z/OS Security Server (RACF) *RACDCERT* command. The *RACDCERT* command installs and maintains PKI private keys and certificates in RACF. RACF supports the management of multiple PKI private keys and certificates as a group. These groups are called key rings. RACF key rings are the preferred method for managing PKI private keys and certificates for SSL.

For more information about z/OS and SSL, see *z/OS Cryptographic Services System SSL Programming*.

For details about the *RACDCERT* command, see *z/OS Security Server RACF Command Language Reference*.

SSL initialization

The TCP/IP variable *SSLENVAR* points to an SSL interface configuration file that contains variable assignment statements that provide the information necessary for SSL initialization.

When you are configuring the IMS Connect SSL interface, you must also consider the ability of the IMS Connect client to support SSL connections.

The following example shows an SSL interface configuration file with the default SSL configuration:

```
#####  
#   This is my SSL interface configuration file   #  
#####  
GSK_PROTOCOL_SSLV2=GSK_PROTOCOL_SSLV2_ON  
GSK_PROTOCOL_SSLV3=GSK_PROTOCOL_SSLV3_ON  
GSK_PROTOCOL_TLSV1=GSK_PROTOCOL_TLSV1_ON  
GSK_KEYRING_FILE=IMSCONNECT  
GSK_KEYRING_LABEL=IMS CONNECT  
GSK_KEYRING_PW=  
GSK_KEYRING_STASH_FILE=  
GSK_CLIENT_AUTH_TYPE=GSK_CLIENT_AUTH_FULL_TYPE  
GSK_SESSION_TYPE=GSK_SERVER_SESSION  
GSK_V2_CIPHER_SPECS=642  
GSK_V3_CIPHER_SPECS=0906030201
```

The SSL interface configuration file can contain both variable assignments and comments.

Each comment in the SSL interface configuration file begins with a number sign (#) when the default host codepage 037 is used. For example, #This is a comment.

Note: In the host codepage 037, the number sign is represented by X'7B'. If you are using a codepage other than host codepage 037, you must use the character in the codepage that is represented by X'7B'. For example, if you are using host code page 1144 for Italy, X'7B' represents the pound sign (£). Consequently, you must begin each comment line with “£”, instead of “#”.

Each variable assignment statement consists of a variable name and the value that you assign to the variable. Variable assignment statements have the following format: GSK_VARIABLE_NAME=value.

Variable assignment statements can continue across multiple lines. To continue a variable assignment statement on the next line, add a blank space and a dash (“ -”) at the end of the line. The line immediately following the blank space and dash character sequence is then read as a continuation of the line. Any characters after the blank space-dash sequence (“ -”) on the same line are ignored.

It is not possible to embed the blank space-dash (“ -”) character sequence within a variable assignment.

For example, the following excerpt from an SSL configuration file has two variable values that contain dashes, one without a blank space, “IMS-Ring”, and one with a blank space, “IMSCONNECT - Test”:

```
// GSK_KEYRING_FILE=IMS-Ring  
// GSK_KEYRING_LABEL=IMSCONNECT - Test  
// GSK_KEYRING_PW=  
// GSK_KEYRING_STASH_FILE=
```

When the above configuration file is processed, GSK_KEYRING_FILE=IMS-Ring processes successfully, but GSK_KEYRING_LABEL=IMSCONNECT Certificate - Test produces an error because the blank space and dash is read as an end of the line. Consequently, the characters after the dash, “ Test”, are discarded and the following line, GSK_KEYRING_PW=, is read as a continuation of the value assigned to GSK_KEYRING_LABEL. To illustrate, here is the job log produced by the above configuration file:


```
// HWSSL00I GSK_KEYRING_FILE = "IMS-Ring"
// HWSSL00I GSK_KEYRING_FILE...SUCCESS (rc = 0)
// HWSSL00I GSK_KEYRING_LABEL = "IMSCONNECTGSK_KEYRING_PW="
// HWSSL00I GSK_KEYRING_LABEL...SUCCESS (rc = 0)
// HWSSL00I GSK_KEYRING_STASH_FILE = NULL
// HWSSL00I GSK_KEYRING_STASH_FILE...SUCCESS (rc = 0)
```

The possible variables and the values associated with the variables are as follows:

GSK_KEYRING_FILE

Name of the key database or RACF keyring. If a RACF key ring is specified, it must be an existing key ring and the current user ID must have READ access to the IRR.DIGTCERT.LISTRING and the IRR.DIGTCERT.LIST resources in the FACILITY class.

GSK_KEYRING_LABEL

Label name in the key database file of RACF key ring used. If this is not set, or set to NULL, the default key database or key ring entry is used.

GSK_KEYRING_PW

Password of the key database. This must be NULL when a RACF key ring is used or when a stash file is specified.

GSK_KEYRING_STASH_FILE

Name of the file that contains the password for the keyring. This value must be NULL when a RACF key ring is used.

GSK_V2_CIPHER_SPECS

A null-terminated character string which specifies the ciphers to enable for SSL V2.0. If this is not specified, the default cipher spec list is used. The default list is 713642 if the Security Level 3 update is installed and 624 if it is not installed.

- 1 - RC4 US
- 2 - RC4 Export
- 3 - RC2 US
- 4 - RC2 Export
- 6 - DES 56-bit Export
- 7 - Triple DES US

Usage example: GSK_V2_CIPHER_SPECS=6321

GSK_V3_CIPHER_SPECS

A null-terminated character string that specifies the ciphers to enable for the SSL V3.0 protocol and the TLS 1.0 and higher protocols. If no value is specified, the default cipher spec list is used. The default list is 05040A0306090201 if the Security Level 3 update is installed and 0306090201 if it is not installed.

For the full list of supported cipher specifications, see *z/OS Cryptographic Services System SSL Programming*.

- 01 - NULL MD5
- 02 - NULL SHA
- 03 - RC4 MD5 Export
- 04 - RC4 MD5 US
- 05 - RC4 SHA US
- 06 - RC2 MD5 Export
- 09 - DES SHA Export

0A - Triple DES SHA US

Usage example: GSK_V3_CIPHER_SPECS=0306090201

GSK_PROTOCOL_SSLV2

Used to enable or disable SSL V2.0. Possible values are GSK_PROTOCOL_SSL_V2_ON and GSK_PROTOCOL_SSLV2_OFF.

Note: All SSL V2.0 non-US encryption schemes have been decrypted. Therefore, SSL V2.0 should not be enabled unless the client does not support SSL V3.0 or TLS V1.0 and higher protocols.

GSK_PROTOCOL_SSLV3

Used to enable or disable SSL V3.0. Possible values are GSK_PROTOCOL_SSLV3_ON and GSK_PROTOCOL_SSLV3_OFF.

GSK_PROTOCOL_TLSV1

Used to enable or disable TLS V1.0. Possible values are GSK_PROTOCOL_TLSV1_ON and GSK_PROTOCOL_TLSV1_OFF.

GSK_PROTOCOL_TLSV1_1

Used to enable or disable TLS V1.1. Possible values are GSK_PROTOCOL_TLSV1_1_ON and GSK_PROTOCOL_TLSV1_1_OFF.

GSK_PROTOCOL_TLSV1_2

Used to enable or disable TLS V1.0. Possible values are GSK_PROTOCOL_TLSV1_2_ON and GSK_PROTOCOL_TLSV1_2_OFF.

GSK_CLIENT_AUTH_TYPE

Indicates the type of client authentication to take place. Two options are available: GSK_CLIENT_AUTH_PASSTHRU_TYPE and GSK_CLIENT_AUTH_FULL_TYPE. The GSK_CLIENT_AUTH_PASSTHRU_TYPE specifies to not authenticate if the client sends a certificate. GSK_CLIENT_AUTH_FULL_TYPE validates all received certificates. If the certificate cannot be validated, the connection is terminated. If no certificate is sent by the client, the connection is unsuccessful.

GSK_SESSION_TYPE

Indicates whether or not to require client authentication. A value of GSK_SERVER_SESSION does not require authentication. GSK_SERVER_SESSION_WITH_CL_AUTH does require client authentication.

GSK_V2_SESSION_TIMEOUT

The number of seconds before the SSL V2.0 session identifier expires. The valid range is from 0 to 100 seconds. If the session timeout value has not expired, the client and server, as well as peer clients (multiple client connections from same client computer) do not need to perform a handshake when starting a new connection.

GSK_V3_SESSION_TIMEOUT

The number of seconds before the SSL V3.0 session identifier expires. The valid range is from 0 to 100 seconds. If the session timeout value has not expired, the client and server as well as peer clients (multiple client connections from same client computer) do not need to perform a handshake when starting a new connection.

GSK_V2_SIDCACHE_SIZE

The maximum number of session ID elements that can be stored in the SSL V3.0 cache. The range is 0 to 32000 entries.

FIPS_MODE

Used to enable or disable FIPS mode. Possible values are GSK_FIPS_STATE_ON and GSK_FIPS_STATE_OFF.

DEBUG_SSL

Indicates whether or not to turn on SSL debugging information. If the debug information is requested, it can be found in the job output after the IMS Connect job has completed. Possible assignment values are ON and OFF.

Usage example: DEBUG_SSL=ON

Related reference:

 z/OS: Cryptographic Services System SSL Programming

SSL default setup

If the SSL initialization file does not exist, a default setup of SSL occurs.

The default setup variables and their values are:

```
GSK_PROTOCOL_SSLV2=GSK_PROTOCOL_SSLV2_ON
GSK_PROTOCOL_SSLV3=GSK_PROTOCOL_SSLV3_ON
GSK_PROTOCOL_TLSV1=GSK_PROTOCOL_TLSV1_ON
GSK_KEYRING_FILE=IMSCONNECT
GSK_KEYRING_LABEL=IMS CONNECT
GSK_KEYRING_PW=
GSK_KEYRING_STASH_FILE=
GSK_CLIENT_AUTH_TYPE=GSK_CLIENT_AUTH_FULL_TYPE
GSK_SESSION_TYPE=GSK_SERVER_SESSION
GSK_V2_CIPHER_SPECS=642
GSK_V3_CIPHER_SPECS=0906030201
```

Note: When creating a new PROCLIB file, ensure that sequence numbers are not automatically inserted in the SSL configuration file. The sequence numbers will cause parsing errors of the SSL options.

Sample JCL for RACF-Managed SSL

RACF as a security manager can create certificates and keyrings, authorize certificates, and store the certificates in a keyring. RACF can also be configured to be a certificate authority.

The following sample JCL illustrates how to set up keyrings and certificates in RACF. In the sample, SSLRING is the sample keyring name and CONNECT is the sample certificate name. RACF stores the certificates in the OMVSADM data set. You must provide a copy of OMVSADM.CERTAUTH.CERT to the client's keyring so RACF can authorize the client.

```
//SSLRACF JOB MSGLEVEL=(1,1),USER=OMVSADM,PASSWORD=,
//          CLASS=A,MSGCLASS=A
//*
/*ROUTE PRINT THISCPU/CSDM09
// EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSABEND DD SYSOUT=*
//SYSTSIN DD *
* Remove labels from mykey and delete the key.
RACDCERT REMOVE(LABEL('CLIENT') RING(mykey))
RACDCERT DELETE(LABEL('CLIENT'))
RACDCERT DELRING(mykey)
RACDCERT LIST(LABEL('CLIENT'))
RACDCERT LISTRING(mykey)
* Remove labels from SSLRING and delete the key.
```

```

RACDCERT REMOVE(LABEL('CONNECT') RING(SSLRING))
RACDCERT DELETE(LABEL('CONNECT'))
RACDCERT REMOVE(LABEL('CERTAUTH') RING(SSLRING))
RACDCERT DELETE(LABEL('CERTAUTH'))
RACDCERT DELRING(SSLRING)
RACDCERT LIST(LABEL('CONNECT'))
RACDCERT LIST(LABEL('CERTAUTH'))
RACDCERT LISTRING(SSLRING)
* Create CLIENT certificate, export to data set, and connect to mykey.
RACDCERT GENCERT-
  SUBJECTSDN(CN('CLIENT') OU('IMS') O('IBM') C('US'))-
  SIZE(512) WITHLABEL('CLIENT')
RACDCERT EXPORT(LABEL('CLIENT')) DSN(CLIENT.CERT) FORMAT(CERTB64)
RACDCERT ADD(CLIENT.CERT) TRUST WITHLABEL('CLIENT')
RACDCERT ADDRING(mykey)
RACDCERT CONNECT(LABEL('CLIENT') DEFAULT RING(mykey))
RACDCERT LISTRING(mykey)
RACDCERT LIST(LABEL('CLIENT'))
* Create SSLRING.
RACDCERT ADDRING(SSLRING)
* Create CERTAUTH certificate, export to data set, and connect to SSLRING
RACDCERT CERTAUTH GENCERT-
  SUBJECTSDN(CN('CERTAUTH') OU('IMS') O('IBM') C('US'))-
  KEYUSAGE(CERTSIGN) WITHLABEL('CERTAUTH')
RACDCERT CERTAUTH EXPORT(LABEL('CERTAUTH')) DSN(CERTAUTH.CERT)
RACDCERT CONNECT(CERTAUTH LABEL('CERTAUTH') RING(SSLRING))
* Create CONNECT certificate, export to data set, and connect to SSLRING.
RACDCERT GENCERT-
  SUBJECTSDN(CN('CONNECT') OU('IMS') O('IBM') C('US'))-
  WITHLABEL('CONNECT')-
  SIGNWITH(CERTAUTH LABEL('CERTAUTH'))
RACDCERT EXPORT(LABEL('CONNECT')) DSN(CONNECT.CERT)
RACDCERT CONNECT(LABEL('CONNECT') DEFAULT RING(SSLRING))
* ?????
RACDCERT ADD(CLIENT.CERT) TRUST WITHLABEL('IMS CONNECT CLIENT')
* ?????
RACDCERT LIST(LABEL('CONNECT'))
RACDCERT CERTAUTH LIST(LABEL('CERTAUTH'))
SETROPTS RACLIST(DIGTCERT) REFRESH
RDEFINE FACILITY IRR.DIGTCERT.LISTRING UACC(NONE)
RDEFINE FACILITY IRR.DIGTCERT.LIST UACC(NONE)
PERMIT IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(OMVSADM) ACCESS(ALTER)
PERMIT IRR.DIGTCERT.LIST CLASS(FACILITY) ID(OMVSADM) ACCESS(ALTER)
SETROPTS CLASSACT(FACILITY)
//

```

For information about the RACF and SSL, see *z/OS Security Server RACF Security Administrator's Guide*.

For details about the RACDCERT command, see *z/OS Security Server RACF Command Language Reference*.

Trusted-user support for IMS Connect messages

When IMS Connect is configured to call RACF directly, you can modify your user message exit to treat specific messages as *trusted users*. When a message is classified as a trusted user, IMS Connect does not call RACF to check security for that message, but instead passes the specified user ID to OTMA without authentication.

After bypassing IMS Connect security, messages that are classified to IMS Connect as a trusted user are still subject to any security checking that might be performed by IMS OTMA. OTMA and IMS do not recognize the IMS Connect trusted user classification.

The following IMS Connect user message exits support the trusted user function:

- HWSSMPL0
- HWSSMPL1
- HWSJAVA0
- User-written user message exits

To enable trusted user support:

- Select one or more IRM fields or an HWSJAVA0 OTMA header fields to contain the flags that identify an input message as a trusted user.
- For IRM fields or customer-written prefix fields, define the bytes and byte settings that represent the trusted user flag so that the definitions are unique to your system. IMS Connect does not define which flag bytes to set or what settings to use.
- Code the user message exit to read the field for the flag and, when the trusted user flag is found, to set the X'80' bit for the OMUSR_TRSTUSR EQU flag in the OMUSR_FLAG2 field in the OTMA user data section of the message that the exit passes back to IMS Connect.

The IRM header fields you can use to identify trusted user messages can include one or more of the following fields: PORTID, CLIENTID, USERID, TRANSACTION CODE fields, and user data.

For example, you might decide to add three one-byte fields in the IRM and to set different values in each field. The client application sets the flag. When the message is passed to the user message exit, the exit interrogates the three fields. If the fields identify the message as a trusted user, the user message exit passes a request to IMS Connect in the OTMA header to bypass the call to RACF.

Sample logic, which is commented out, is provided in both HWSSMPL0 and HWSSMPL1 and can be found by looking for the following comment lines:

```
*****
*****TRUSTED USER SUPPORT*****
*****
```

If you are using the HWSJAVA0 user message exit, the exit identifies trusted user messages by the existing data in OTMA headers fields such as OMUSR_DESTID (DataStore), OMUSR_ORIGIN (ClientID), OMUSR_PORTID (PortId), OMUSR_PASSTICKET (Password), or other message values.

Code the HWSJAVA0 user message exit to set the OMUSR_FLAG2 flag to OMUSR_TRSTUSR. When the OMUSR_FLAG2 flag is set to OMUSR_TRSTUSR, IMS Connect bypasses the RACF call.

Specifying an OTMA ACEE aging value in the IMS Connect configuration member

You can configure IMS Connect to pass accessor environment element (ACEE) aging values to OTMA. ACEE aging values ensure that OTMA refreshes cached RACF ACEEs at appropriate frequencies for your installation.

The ACEE contains user IDs and other security information. OTMA caches the ACEE to improve the performance of the security checking that OTMA performs. The ACEEs are used to validate the authority of the user IDs included with the input messages passed by IMS Connect to the transactions and commands being requested.

To ensure that this security information is correct, OTMA periodically refreshes the ACEE with the security information that is stored in RACF. The length of time between refreshes of the ACEE is determined by the ACEE aging value, which is set in OTMA.

To define an ACEE aging value for IMS Connect to pass to OTMA, use the OAAV= keyword in the DATASTORE configuration statement of the IMS Connect PROCLIB member, HWSCFGxx. You can specify a value from 300 to 2147483647 seconds. The default is 2147483647 seconds. OTMA requires a value of at least 300 to enable ACEE refreshes.

Related reference:



HWSCFGxx member of the IMS PROCLIB data set (System Definition)

Chapter 15. IMS Connect support for callout requests

IMS Connect is a required component when IMS application programs issue callout requests through OTMA to data or service providers that are external to the IMS installation. For both types of callout request, IMS Connect serves as the TCP/IP gateway between the IMS Connect client that use TCP/IP and the OTMA component of IMS.

IMS Connect supports callout requests without requiring any changes to the IMS Connect configuration statements. You do, however, have to configure the IMS Connect clients to support callout requests.

You can use the following types of IMS Connect clients to support synchronous and asynchronous callout requests:

- IMS TM Resource Adapter
- IMS Enterprise Suite SOAP Gateway
- User-written IMS Connect client

The following IMS Connect exit routines support callout requests:

- HWSJAVA0
- HWSSOAP1
- HWSSMPL0
- HWSSMPL1

IMS provides various sample application programs to test callout support. For more information, see *Samples for the callout function (Installation)*.

For information about how to configure IMS TM Resource Adapter and SOAP Gateway, see:

- Callout programming models
- Enabling an IMS application as a web service consumer

Related information:

Configuring user-written IMS Connect clients for synchronous callout requests

To support synchronous callout requests, user-written IMS Connect clients must be configured to retrieve new callout requests from IMS, acknowledge the receipt of the callout request (ACK or NAK), and to return the synchronous callout responses to IMS through IMS Connect.

Synchronous callout request messages are handled by IMS Connect and OTMA in much the same way as asynchronous output. That is, synchronous callout request messages are retrieved by issuing a RESUME TPIPE call. Many of the same rules and guidelines for retrieving asynchronous output also apply to retrieving synchronous callout request messages.

IMS provides various sample application programs to test callout support. For more information, see *Samples for the callout function (Installation)*.

The high-level steps for configuring a user-supplied IMS Connect client to support synchronous callout messages include coding the client to:

1. Retrieve callout requests by using the RESUME TPIPE call. When retrieving new synchronous callout requests, a user-written IMS Connect client usually issues the RESUME TPIPE call with a very long timeout value or no timeout value.

If a synchronous callout request message is queued when the RESUME TPIPE call is received by IMS, OTMA sends it to the client through IMS Connect. If no callout message is queued, the IMS Connect client waits in a receive state and OTMA sends the next callout request to the client immediately upon arrival at the output queue.
2. Acknowledge the successful (ACK) or unsuccessful (NAK) receipt of synchronous callout requests. To acknowledge the receipt or the rejection of a synchronous callout request, the user-written IMS Connect client must send an ACK or a NAK response to IMS to free the output queue for new callout messages. If an ACK or NAK is not received before the timeout value specified for either the synchronous callout request or the acknowledgment message itself, OTMA discards the callout request and frees the output queue.
3. Return a response message to IMS Connect by using the synchronous callout option of the send-only protocol. The response must include the correlator token that IMS Connect included in the original callout request. The response can be either the data requested by the IMS application or an error message.

In addition to these steps for configuring the IMS Connect client, the following additional steps are required to complete the configuration of your installation for synchronous callout support:

- Code an IMS application program to initiate a synchronous callout request by issuing the DL/I ICAL call.
- Code an OTMA destination descriptor to route the synchronous callout request to the IMS Connect client.

The following figure shows the flow of a synchronous callout request between IMS and a user-written IMS Connect client.




Figure 23. Flow of a synchronous callout request between an IMS application and a user-written IMS Connect client

Format of synchronous callout messages

Synchronous callout messages contain a segment for the correlator token. The correlator token is used to correlate responses with the IMS application programs that issue the callout requests. User-supplied client application programs must account for this token when reading callout request messages and return it with callout response messages.

The presence of the correlation token identifies the message to the IMS Connect client as a synchronous callout message. IMS generates the correlation token and includes it in the outgoing callout request automatically. The IMS Connect client must return the correlation token in the response that it sends back to IMS.

The application data segment in synchronous callout messages can be larger than 32 K and potentially many megabytes in length. IMS Connect includes the length of the application data segment at the beginning of the segment in a four-byte length field, LLLL. The length specified in the LLLL field includes the length of the LLLL field itself. The IMS Connect client must be configured to read all four bytes of the length field.

When the IMS Connect client returns the response to the callout request, the client must include the length of the application data segment in the four-byte length field at the beginning of the segment by using the same LLLL format.

Attention: Be aware that if multiple ICAL calls are issued concurrently, larger message sizes can consume larger amounts of extended private storage.

The following examples illustrate the format of synchronous callout requests and synchronous callout responses. In the examples, the segments enclosed in braces, { and }, are optional.

The structure of the correlation token itself cannot be modified; however, its structure is defined in the HWSIMSCB and HWSOMPFX macros.

The structure of a synchronous callout request message sent by IMS Connect to a client application conforms to the following format:

LLLL | LLZZ COR | {LLZZ RMM} | LLLL data | LLZZ CSM

The structure of a synchronous callout response returned by the client application must conform to the following format:

LLLL | LLZZ IRM | LLLL data | 00040000

For comparison, a message that is not a synchronous callout request and that does not support the 4-byte format of length field the application data segment looks like this:

LLLL | {RMM} | LLZZ data | {LLZZ data} | CSM

Related reference:

“Output message from message exit to client” on page 257

“Format of user portion of IRM for HWSSMPL0, HWSSMPL1, and user-written message exit routines” on page 241

“OTMA user data fields used by IMS Connect” on page 283

Retrieving synchronous callout requests with RESUME TPIPE

When issuing a RESUME TPIPE call to retrieve synchronous callout requests, you can code the RESUME TPIPE call to retrieve only synchronous callout messages or both synchronous callout messages and asynchronous output.

The RESUME TPIPE call for synchronous callout messages requires IRM architecture 3 (IRM_ARCH3) and a value of X'80' (IRM_F0_SYNONLY) or X'40' (IRM_F0_SYNASYN) in the IRM_F0 field. If no option is specified in the IRM_F0 field, the RESUME TPIPE call retrieves only asynchronous output.

IRM_F0_SYNONLY specifies that the IMS Connect client can only retrieve synchronous callout requests. If any asynchronous output messages are present on the tpipe hold queue when a RESUME TPIPE call that specifies IRM_F0_SYNONLY is received, the asynchronous messages are not returned to the client and remain on the tpipe hold queue.

IRM_F0_SYNASYN specifies that the IMS Connect client can retrieve both synchronous callout requests and any other asynchronous output messages on the queue. When IRM_F0_SYNASYN is specified, when a new RESUME TPIPE call is received, OTMA sends all synchronous callout requests first before sending asynchronous output messages.

Define the RESUME TPIPE call for synchronous callout messages with the following field specifications in the IRM prefix:

IRM_ARCH

X'03' (IRM_ARCH3).

IRM_F0

Specify either:

- X'80' (IRM_F0_SYNONLY) to retrieve only synchronous callout messages.
- X'40' (IRM_F0_SYNASYN) to retrieve both synchronous callout messages and asynchronous output messages.

IRM_F4

R character value (IRM_F4_RESUMET).

IRM_F5

Retrieval option for the RESUME TPIPE call, such as X'02' (IRM_F5_AUTO).

IRM_TIMER

The timeout value for the RESUME TPIPE call, such as IRM_TIME_FF, which specifies that the RESUME TPIPE call does not time out.

Related concepts:

“Managing the retrieval of output messages” on page 333

“Timeout specifications on input messages” on page 318

“Socket connections” on page 308

Related reference:

Chapter 17, “IMS Connect message structures,” on page 237

“Output message from message exit to client” on page 257

RESUME TPIPE error scenarios

If the IMS Connect client issues a RESUME TPIPE call for synchronous callout messages with the IRM_F0 field set to either IRM_F0_SYNONLY or IRM_F0_SYNASYN, but IRM_ARCH is not greater than or equal to IRM_ARCH3, IMS connect returns the following return and reason codes to the IMS Connect client in the RSM structure:

- RSM_RETCOD = RSMRTC_EXIT X'04'
- RSM_RSNCOD = RSMRSN_INVBUF X'09'

If IMS Connect retrieves a message on a tpipe that is used for synchronous callout message, but cannot deliver the message to the IMS Connect client, IMS Connect replies to OTMA with a NAK message. How OTMA handles the message after receiving the NAK message depends on the status of the IMS Connect client and whether the undeliverable message is a synchronous callout message or an asynchronous output message.

For asynchronous output messages, OTMA returns the asynchronous output message to the tpipe hold queue.

For synchronous callout messages:

- If IMS Connect could not deliver the synchronous callout message because the client either timed out or otherwise became disconnected, OTMA keeps the synchronous callout message on the queue for retrieval by another RESUME TPIPE call or until the timeout interval specified on the ICAL call expires.
- If IMS Connect could not deliver the synchronous callout message for a reason other than timeout or disconnection, such as TCP/IP errors, OTMA discards the synchronous callout message.

Acknowledging receipt of synchronous callout messages

A user-written IMS Connect client must send a positive (ACK) or negative (NAK) acknowledgment message to IMS after OTMA sends a synchronous callout request message. Until the ACK or NAK message is received or the timeout interval expires for either the acknowledgment or the original request, the tpipe queue remains in a wait state (WAIT_S) and cannot send or receive any other messages.

An ACK message frees the tpipe queue to send and receive other callout messages and the IMS application program continues to wait in the dependent region until either a response to the synchronous callout message is received or the synchronous callout request times out.

A NAK message not only frees the tpipe queue, but also tells OTMA what to do with the synchronous callout request message and whether to maintain or end the current RESUME TPIPE call.

The IMS Connect client can specify one of the following options on a NAK message:

- Discard the rejected synchronous callout request message and terminate the RESUME TPIPE connection.
- Discard the rejected synchronous callout request message, but maintain the RESUME TPIPE connection to continue retrieving other synchronous callout request messages.
- Keep the rejected synchronous callout request message on the tpipe queue, but terminate the RESUME TPIPE connection.

Related concepts:

“Timeout intervals for IMS Connect acknowledgments to OTMA” on page 326

Coding a NAK message to discard the callout message and end RESUME TPIPE call

To code a NAK message that directs OTMA to discard the rejected synchronous callout message and end the RESUME TPIPE call:

1. Specify N in the IRM_F4 field
2. Optional: Include an extended error code by specifying the following:

- IRM_F0 = X'10'
- IRM_NAK_RSNCDE = A two-byte hexadecimal extended error code

The extended error code is returned to the IMS application in the AIBERRXT field of the AIB of the ICAL call.

Upon receiving this NAK message, OTMA issues to the IMS application program a return code of X'100' and a reason code of X'108'.

If the NAK is for an asynchronous output message, the asynchronous output is returned to the tpipe hold queue.

Coding a NAK message to discard a message, but keep a connection

To code a NAK message to direct OTMA to discard the rejected synchronous callout message, but keep the RESUME TPIPE call to retrieve other synchronous callout request messages:

1. Specify the following field values in the IRM of the NAK message:
 - IRM_F4 = N

- IRM_F3 = X'08' (IRM_F3_REROUT)
2. Optional: Include an extended error code by specifying the following:
- IRM_F0 = X'10'
 - IRM_NAK_RSNCDE = A two-byte hexadecimal extended error code
- The extended error code is returned to the IMS application in the AIBERRXT field of the AIB of the ICAL call.

Upon receiving this NAK message, the IMS application program receives a return code of X'100' and a reason code of X'108'.

If the NAK is for an asynchronous output message, the asynchronous output is returned to the reroute tpipe hold queue.

Coding a NAK message to retain message, but end RESUME TPIPE call

To code a NAK message to direct OTMA to keep the rejected synchronous callout message on the tpipe queue, but end the RESUME TPIPE call:

Specify the following field values in the IRM of the NAK message:

- IRM_F4 = N
- IRM_F0 = X'20' (IRM_F0_SYNCNAK)

Upon receiving the NAK message, the RESUME TPIPE call ends, but OTMA keeps the synchronous callout request message on the tpipe queue. The rejected synchronous callout request message can then be retrieved by another RESUME TPIPE call, if the call is received by OTMA before the callout request times out.

If the NAK is for an asynchronous output message, OTMA places the message back on the hold queue and terminates the RESUME TPIPE call.

Returning callout responses to IMS

User-written IMS Connect clients return the callout response messages to IMS by using the send-only protocol for synchronous callout responses.

The send-only message contains no transaction code and can contain either the response data or error information for the IMS application program.

A response to a synchronous callout request must be returned to the same IMS system in which the original DL/I ICAL call was issued. However, the response does not need to be returned by the same IMS Connect client that processed the original callout request. The response message can also be returned to IMS through a different IMS Connect instance than the instance that handled the outbound callout request. The correlation token included in the response message ensures that the response is returned to the correct IMS application program that is waiting for it.

The response message returned by IMS Connect client can be up to the length specified by the MAXSIZE parameter in the HWS configuration statement or the IMS Connect default of 10•MB.

The IMS Connect client must indicate in the IRM_F4 field of the message prefix that the message being returned is a synchronous callout response.

If the client returns a synchronous callout response to IMS Connect, but either the IRM_ARCH is less than IRM_ARCH3 or the correlation token is not included in the IRM_CORTKN field, IMS Connect rejects the response and the synchronous callout request message times out.

To code the callout response:

Specify the following field values in the IRM of the response message:

- IRM_ARCH = X'03' (IRM_ARCH3)
- IRM_F4 = M (IRM_F4_SYNRESP)
- IRM_CORTKN = The 40-byte correlation token (CORTKN) that was included with the original callout request

Related concepts:

“Send-only protocol for synchronous callout responses” on page 307

Related reference:

Chapter 17, “IMS Connect message structures,” on page 237

Returning an error response to IMS

If an error occurs after the user-written IMS Connect client has already returned an ACK message to IMS Connect and either the IMS Connect client or the data or service provider cannot complete the callout request, the IMS Connect client can return an error response to IMS instead of the expected data.

To code an error response to a DL/I ICAL call:

Specify the following field values in the IRM prefix of the error response:

- IRM_ARCH = X'03' (IRM_ARCH3)
- IRM_F0 = X'20' (IRM_F0_SYNCNAK)
- IRM_NAK_RSNCDE = 2-byte hexadecimal extended error code
- IRM_F4 = M (IRM_F4_SYNRESP)
- IRM_CORTKN = the 40-byte correlation token generated by the original DL/I ICAL call.

Related concepts:

“Send-only protocol” on page 305

Related reference:

Chapter 17, “IMS Connect message structures,” on page 237

Chapter 16. IMS Connect XML message conversion

For some IMS Connect client application programs, IMS Connect can convert XML messages into either COBOL or PL/I, so that you do not need to modify existing IMS application programs to process messages submitted to IMS in XML.

IMS Connect conversion support converts the XML user data to the language structures expected by the IMS application program, which can then process the message as a normal IMS transaction message.

The IMS Connect XML conversion support is provided by an XML adapter function and an HWSSOAP1 message exit with a COBOL XML converter.

The actual conversion of XML messages is performed by an XML converter, which is called by IMS Connect.

IMS Connect supports XML conversion into the following languages for the following clients:

COBOL

- IMS Enterprise Suite SOAP Gateway
- IMS Web 2.0 Solution for IBM Mashup Center

PL/I

- IMS Web 2.0 Solution for IBM Mashup Center

IMS Connect also supports the conversion of multi-segment messages between SOAP Gateway and IMS application programs that function as Web service providers and that use multiple language structures. Restrictions apply to the multi-segment message support. These restrictions are detailed in the SOAP Gateway and Rational® Developer for System z® documentation.

The IMS Connect HWSSOAP1 exit routine also supports XML conversion and you can specify an XML adapter name and an XML converter name in OTMA destination descriptors.

In the IMS request message (IRM) header of the input message, the IMS Connect client requests XML conversion support by specifying an XML adapter name and an XML converter name. IMS Connect reads the IRM when it receives an input message and, if XML conversion is required, IMS Connect calls the XML adapter. The XML adapter then calls the XML converter to perform the actual conversion. After the message is converted from its XML format, IMS Connect then sends the resulting message to IMS.

When IMS Connect receives the response message from IMS, IMS Connect calls the XML adapter, which in turn calls the XML converter to convert the user data in the response message into XML. IMS Connect then sends the output message to the IMS Connect client.

Related tasks:

 [Configuring XML conversion support for IMS Connect clients \(System Definition\)](#)

Related reference:

IMS Connect XML converters

When you enable IMS Connect XML conversion support, you must create the XML converters that are used by IMS Connect to convert the data structures from XML to the programming language used by the IMS application program and vice versa.

The XML converters are application programs. You can find an example of a COBOL XML converter by downloading the IMS Enterprise Suite SOAP Gateway Phone Book Sample from the IMS Enterprise Suite SOAP Gateway home page at www.ibm.com/software/data/ims/soa-enterprise-suite/soap/

Recommendation: Use the separately licensed tool IBM Rational Developer for System z to automatically generate the XML converters from the COBOL copybook or the PL/I source code.

For COBOL application programs, the XML converter is based on the COBOL copybook of the IMS COBOL application program that processes the message. For PL/I application programs, the XML converter is based on the source of the PL/I application program. Each IMS application that processes messages converted from XML must have its own unique XML converter.

The XML converters run in an IBM Language Environment for z/OS enclave in the IMS Connect region. The amount of space required depends on the size and number of XML converters required in your environment. The IMS Connect region size must be increased to accommodate this storage requirement.

Related tasks:

 [Configuring XML conversion support for IMS Connect clients \(System Definition\)](#)

Related reference:

 [ADAPTER statement \(System Definition\)](#)

Structure of the XML message

An XML schema defines the XML tags that correspond to the COBOL data structures used by the COBOL IMS application program. The XML schema used for XML conversion is based on the COBOL copybook of the COBOL IMS application program.

The XML schema is not required by IMS Connect, but the application programmer that develops the web service that generates the XML input messages needs it.

If you use Rational Developer for System z to automatically generate the XML converters from the COBOL copybook of the IMS application program, Rational Developer for System z also generates the XML schema for you.

The XML tags defined by the XML schema directly correspond to the fields of the COBOL data structure.

The following is an example of an input and output message data structure defined in the COBOL copybook of the sample phone book application program

that is available at the IMS Enterprise Suite SOAP Gateway download site through www.ibm.com/software/data/ims/soap/.

```

01 INPUT-MSG.
   02 IN-LL          PICTURE S9(3) COMP.
   02 IN-ZZ          PICTURE S9(3) COMP.
   02 IN-TRCD        PICTURE X(10).
   02 IN-CMD          PICTURE X(8).
   02 IN-NAME1        PICTURE X(10).
   02 IN-NAME2        PICTURE X(10).
   02 IN-EXTN         PICTURE X(10).
   02 IN-ZIP          PICTURE X(7).

01 OUTPUT-MSG.
   02 OUT-LL          PICTURE S9(3) COMP.
   02 OUT-ZZ          PICTURE S9(3) COMP.
   02 OUT-MSG         PICTURE X(40).
   02 OUT-CMD          PICTURE X(8).
   02 OUT-NAME1        PICTURE X(10).
   02 OUT-NAME2        PICTURE X(10).
   02 OUT-EXTN         PICTURE X(10).
   02 OUT-ZIP          PICTURE X(7).
   02 OUT-SEGNO        PICTURE X(4)

```

Each field in the copybook has an equivalent XML tag that represents the field in the XML message. XML tags are case-sensitive. The dash symbol '-' in field names in the copybook is represented as an underscore '_' in the corresponding XML tags. The SOAP Gateway client has to build the XML message using that XML schema.

For example, the IN-TRCD field from the above copybook is represented in XML by the opening and closing tags <in_trcd> and </in_trcd>. In the data structure input message, the value of IN-TRCD should be placed in the byte positions 5 to 14. In the XML input message, the same value of IN-TRCD should be placed between the tags <in_trcd> and </in_trcd>.

An input messages from an IMS Connect client that uses the above example COBOL copybook of the phone book application has the following XML tags for the equivalent COBOL data structure fields:

```

<INPUTMSG>
<in_ll> </in_ll>
<in_zz> </in_zz>
<in_trcd> </in_trcd>
<in_cmd> </in_cmd>
<in_name1> </in_name1>
<in_name2> </in_name2>
<in_extn> </in_extn>
<in_zip> </in_zip>
</INPUTMSG>

```

The above COBOL copybook of the phone book application has the following XML tags for the equivalent outbound COBOL data structure fields:

```

<cb1:OUTPUTMSG>
<out_ll> </out_ll>
<out_zz> </out_zz>
<out_msg> </out_msg>
<out_cmd> </out_cmd>
<out_name1> </out_name1>
<out_name2> </out_name2>
<out_extn> </out_extn>
<out_zip> </out_zip>
<out_segno> </out_segno>
</cb1:OUTPUTMSG>

```

The input message XML tags must be wrapped by opening and closing XML tags based on the 01 INPUT-MSG. definition in the COBOL copybook. In the above example, the opening and closing tags are <INPUTMSG> and </INPUTMSG>.

The output message XML tags generated by the outbound converter are wrapped by opening and closing XML tags that correspond to the 01 OUTPUT-MSG. definition. In the above example, the opening and closing tags are <cb1:OUTPUTMSG> and </cb1:OUTPUTMSG>.

Look at the XML schemas to determine what these tags should be for each XML converter.

For each XML message, not all tags have to be specified, just like not all fields are required in the data structure message. The required fields are determined by the COBOL application.

Message conversion example

The code in this example describes the conversion of an XML input message to COBOL and its COBOL response message to XML.

Below is an example of an XML message on input:

```
<INPUTMSG><in_ll>32</in_ll><in_zz>0</in_zz>  
<in_trcd>IVTN0</in_trcd><in_cmd>DISPLAY</in_cmd>  
<in_name1>LAST1</in_name1></INPUTMSG>
```

Below is an example of the same message after the XML data structure has been converted to the COBOL data structure required by the COBOL phone book application program:

```
IVTN0      DISPLAY LAST1
```

Note: In the above example, IVTN0 starts in the fifth byte. The first four bytes are used by IMS Connect, and there are five spaces following LAST1.

The first four bytes of the data structure message, known as LLZZ, will be filled by the XML adapter. The first two bytes, the LL part, will be filled with the length of the data structure message. The second two bytes, the ZZ part, will be filled with zeroes. The XML converter converts the XML message to the application-specific format by taking the value within each XML tag, and placing it in its corresponding field position. The converted message can then be processed by the COBOL application and it returns an output message in its specific data structure. The output message has to be converted to XML before it is returned back to the Client. Below is an example of an output message and the message after XML conversion.

Below is an example of the COBOL data structure of the reply message on output from the IMS data store:

```
ENTRY WAS DISPLAYED                DISPLAY LAST1    FIRST1    8-111-1111D01/R010001
```

In the above example, the text on the first line of the output message begins at the fifth byte, after the four byte LLZZ field.

Below is an example of the same message after the COBOL data structure has been converted to XML as required by the IMS Connect client:

```
<cb1>  
  
<out_ll> 093</out_ll><out_zz> 000</out_zz>  
<out_msg>ENTRY WAS DISPLAYED</out_msg><out_cmd>DISPLAY</out_cmd>  
<out_name1>LAST1</out_name1><out_name2>FIRST1</out_name2>  
<out_extn>8-111-1111</out_extn><out_zip>D01/R01</out_zip>  
<out_segno>0001</out_segno></cb1>
```

In the example above, the values following the opening XML tags `<out_ll>` and `<out_zz>` tags are from the first four bytes of the output COBOL data structure message.

For outbound messages from the IMS application, the conversion is performed in reverse. Each field in the application data structure is wrapped in its corresponding XML tags. This conversion from an XML to a COBOL application data structure message format, and vice versa, is performed by the COBOL XML converter called by the XML adapter.

Chapter 17. IMS Connect message structures

TCP/IP clients and local clients using the z/OS program call interface communicate with IMS Connect by using an IMS request message (IRM) header on each input message. The IRM header is used on input messages from IMS Connect client application programs to communicate protocol options to IMS Connect. The IRM header is mapped by the IMS Connect HWSIMSCB macro.

IMS Connect communicates with OTMA through an z/OS cross-system coupling facility session using the OTMA message headers. One of the primary tasks of the client-specific IMS Connect user-message exit routines is to translate the IRM header into the OTMA header on input.

Clients that use TCP/IP socket calls as their communication vehicle can design a user message exit routine that runs with IMS Connect to convert messages between formats as follows:

- Convert the client message format to OTMA message format.
- Convert the IMS response, in OTMA message format, to client message format.
- Convert XML in the user data portion of messages to COBOL and back again.

These conversions enable the client to retrieve IMS data through a TCP/IP connection. IMS Connect automatically sends and receives messages when they are formatted correctly.

Depending on the type of IMS Connect client and the user message exit used to support the client, the format of the IRM message header differs.

The following user message exit routines require specific message formats:

- HWSSMPL0 and HWSSMPL1
- HWSJAVA0
- HWSCSLO0 and HWSCSLO1
- HWSSOAP1

Related concepts:

Chapter 12, “Overview of IMS Connect exit routines,” on page 183

Related reference:

 [IMS Connect exit routines \(Exit Routines\)](#)

IRM structures for IMS Connect client messages

IMS Connect expects all client messages that it receives to start with a four byte total length field, followed by an IMS request message (IRM) header, followed by the message data segments.

After the four-byte length field, all IRMs begin with a 28 byte fixed-format section that is common to all messages from all IMS Connect client applications. You cannot modify the format of the fixed section of the IRM.

User-written client applications can include a user-defined section in the IRM after the fixed-format section. The user-defined section of the IRM allows you to specify additional options. If you add or delete any fields in the user-defined section of the

IRM, you must also adjust the user message exits that you use, or provide your own user message exits. The user-defined section of the IRM is only supported by the HWSSMPL0 and HWSSMPL1 user message exit routines.

If you are using either the HWSSOAP1 user message exit, which is not provided as source, or the HWSJAVA0 user message exit, you cannot change the user portion of the IRM.

Note: The user message exits HWSCSLO0 and HWSCSLO1 do not have any message definitions because HWSCSLO0 and HWSCSLO1 cannot be modified or replaced.

Format of fixed portion of IRM in messages sent to IMS Connect

The IMS request message (IRM) header contains a 28-byte fixed-format section that is common to all messages from all IMS Connect client applications that communicate with IMS TM.

The following table shows the fixed format preceding the input message sent to IMS Connect from clients. It includes the message field, the field length, and a brief explanation of the message.

The llll field tells IMS Connect how long the message is, and the IRM provides additional information, such as the specific user exit to which the data is to be passed.

Table 17. Fixed portion of IRM prefix.

Field	Length	Offsets		Meaning
		Dec	Hex	
llll	4 bytes	0	0	Length of the total message read as a binary number. The total message length includes the length of the llll field (four bytes), the IRM (variable length, depending on your requirements), the length of the message, and the end of message indicator X'00040000' (four bytes). The minimum length for all messages is X'58'. For IMS TM Resource Adapter, the maximum length is X'00989680' (10,000,000 bytes). For user-written client applications, the maximum length is X'7FFFFFFF'.
The following fields are for the 28 (decimal) byte IRM prefix.				
IRM_LEN	2 bytes	4	4	Length of the IRM structure. The minimum size of the IRM for user written exits is X'24' or binary '00100100'. HWSSMPL0 has a minimum IRM length of X'50' or binary '01010000'.

Table 17. Fixed portion of IRM prefix (continued).

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_ARCH	1 byte	6	6	Specifies the architectural level of the IRM prefix in messages received by IMS Connect from the client. • X'00' Specifies IRM_ARCH0, the base architectural structure of the user portion of the IRM prefix. • X'01' Specifies IRM_ARCH1, the architectural structure of the user portion of the IRM prefix that includes space for: – The IRM_REROUT_NM field – The IRM_RT_ALTCID field • X'02' Specifies IRM_ARCH2, the architectural structure of the user portion of the IRM prefix that includes space for: – The IRM_REROUT_NM field – The IRM_RT_ALTCID field – The IRM_TAG_ADAPT field – The IRM_TAG_MAP field • X'03' Specifies IRM_ARCH3, the architectural structure of the user portion of the IRM prefix that includes space for all of the fields included in IRM_ARCH2 and the following additional fields: – The synchronous callout correlation token fields – The IRM_MODNAME field for MFS MOD names
IRM_F0	1 byte	7	7	• X'80' Specifies IRM_F0_SYNONLY: This is a resume tpipe call that retrieves only synchronous callout messages from IMS application programs running in IMS dependent regions. • X'40' Specifies IRM_F0_SYNASYN: This is a resume tpipe call that retrieves both synchronous callout messages from IMS application programs and asynchronous messages. • X'20' Specifies IRM_F0_SYNCNAK: This is a NAK message from the IMS Connect client that directs OTMA to keep the message that triggered the NAK on the tpipe queue. • X'10' Specifies IRM_F0_NAKRSN: This is a NAK message that includes a reason code for the NAK response. • X'01' Specifies IRM_F0_XMLTD, which indicates a request from an IMS Enterprise Suite SOAP Gateway client to convert an XML tagged message that contains both a transaction code and data into the format expected by the IMS application program. • X'02' Specifies IRM_F0_XML_D, which indicates a request from an SOAP Gateway client to convert an XML tagged message that contains data only into the format expected by the IMS application program.

Table 17. Fixed portion of IRM prefix (continued).

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_ID	8 bytes	8	8	Character string. Specifies the identifier of the user exit that is to be driven after the complete message has been received. Additionally, IMS Connect reads this field to determine whether the incoming message is in ASCII or EBCDIC. The IMS Connect-supplied user message exits reserve and use these IDs: • *HWSCSL*-- for HWSCSLOO • *HWSJAV*-- for HWSJAVA0 • *HWSOA1*-- for HWSSOAP1 • *SAMPL1*-- for HWSSMPL1 • *SAMPLE*-- for HWSSMPL0
IRM_NAK_RSNCDE	2 bytes	16	10	Optional reason code for a NAK response.
IRM_RES1	2 bytes	18	12	Reserved for future use. Initialize to binary zeros.
IRM_F5	1 byte	20	X'14'	Input message type. X'80' OTMA headers built by client. X'40' Translation done by client. X'10' Single message with wait option. Only one message returned following the RESUME TPIPE call. If no message is present, OTMA waits for a message to arrive and then sends that single message to IMS Connect. The timer set on the RESUME TPIPE call can expire before a message is returned to IMS Connect. If that occurs, IMS Connect issues a NAK response to the message when received. X'04' No auto flow of messages. All current messages are returned one at a time. Use the no auto flow option only if the client is a dedicated output client. This value is similar to Auto, except that the IRM_TIMER will cause the last receive to terminate. Set the IRM_TIMER to a small value. Each ACK sent by the client resets the IRM_TIMER value. The IRM_TIMER value set by the RESUME_TPIPE only applies to the first receive state. X'02' Auto flow of messages. All current messages are returned one at a time. Use the auto flow option only if the client is a dedicated output client. Set the IRM_TIMER to a large value. Each ACK sent by the client resets the IRM_TIMER value. The IRM_TIMER value set by the RESUME_TPIPE only applies to the first receive state. X'01' Single message. Only one message returned following the RESUME TPIPE call. If no message is present, OTMA does not wait for a message and the IMS Connect timer causes a timeout to occur based on the timeout value specified. X'08' IRM_F5_XID The message includes an X/Open identifier (XID). X'00' No option flow of messages (see meaning for X'04'). This is the default if no value is specified.

Table 17. Fixed portion of IRM prefix (continued).

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_TIMER	1 byte	21	X'15'	Time delay that IMS Connect will wait for IMS to return data to IMS Connect which, in turn, will be sent to the client. The following functions support the IRM_TIMER settings: • TCP/IP SEND of a RESUME TPIPE call • TCP/IP SEND of an ACK or NAK message • TCP/IP SEND of data • PC SEND of a RESUME TPIPE call • PC SEND of an ACK or NAK message • PC SEND of data
IRM_SOCT	1 byte	22	X'16'	Socket connection type. The client can set this value as follows: X'00' Transaction socket. The socket connection lasts across a single transaction. X'10' Persistent socket. The socket connection lasts across multiple transactions. X'40' Non-persistent socket. The socket connection lasts for a single exchange consisting of one input and one output. Recommendation: Do not use this socket type if you plan on implementing conversational transactions, because multiple connects and disconnects will occur.
IRM_ES	1 byte	23	X'17'	Unicode encoding schema. Initialize to binary zeros. X'01' UTF8 encoding schema. X'02' UCS2 encoding schema. X'02' UTF16 encoding schema.
IRM_CLIENTID	8 bytes	24	X'18'	A string of 1 to 8 uppercase alphanumeric (A through Z, 0 to 9) or special (@, #, \$) characters, left justified, and padded with blanks. IRM_CLIENTID specifies the name of the client ID that is used by IMS Connect. If this string is not supplied by the client, then the user exit must generate it. The client ID is returned to IMS Connect from the exit in the EXIT PARMLIST field, EXPREA_CLID.

Related concepts:

“Timeout specifications on input messages” on page 318

Related reference:

“Format of user portion of IRM for HWSSMPL0, HWSSMPL1, and user-written message exit routines”

Format of user portion of IRM for HWSSMPL0, HWSSMPL1, and user-written message exit routines

Following the 4-byte length field and the 28-byte fixed portion of the IMS request message (IRM) header in IMS Connect client input messages, user-written client applications supported by HWSSMPL0, HWSSMPL1, or user-written message exits can include a user-defined section in the IRM.

Messages submitted by the IMS Connect client might or might not include application or IMS data after the IRM depending on the type of message. For example, the following message types, which are specified in the IRM_F4 field, do not contain a data element after the IRM:

- ACK (A)
- CANCEL TIMER (C)
- DEALLOCATE (D)
- NAK (N)
- RESUME TPIPE (R)

Message types that do contain application data after the IRM include the following types, which are also specified in the IRM_F4 field:

- SENDONLY (S)
- SENDONLYA (K)
- Send-receive (X'40')
- Synchronous callout response (M)

The data that follows the user portion of the IRM must be in the format of LLZZDATA, where LL is the total length of message segment (including the LL field), ZZ is binary zeros, and DATA is the IMS transaction code followed by the transaction data.

Restriction: If you are using IMS Connect client applications such as IMS TM Resource Adapter, IMS Enterprise Suite SOAP Gateway, or IBM Management Console for IMS and DB2 for z/OS, you cannot include a user-defined portion in the IRM.

The following table shows the format of the user portion of the IRM used by HWSSMPL0, HWSSMPL1, and other user-written user message exits.

This topic contains Product-sensitive Programming Interface information.

Table 18. User portion of IRM for HWSSMPL0, HWSSMPL1, and user-written user message exits

Field	Length	Offsets		Meaning
		Dec	Hex	
IIII	4 bytes	0	X'00'	Length of entire message including the IRM, application data, end of message indicator, and this length field
Fixed-format IRM	28 bytes	4	X'04'	The common, fixed portion of the IRM prefix.
I IRM_MAP	0 bytes	32	X'20'	Position of the XID map when it is present.

Table 18. User portion of IRM for HWSSMPL0, HWSSMPL1, and user-written user message exits (continued)

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_F1	1 byte	32	X'20'	The purpose of this field differs depending on its value. X'80' - IRM_F1_MFSREQ The user requests that IMS return an MFS MOD name. X'20' - IRM_F1_UC This is a Unicode message. X'10' - IRM_F1_UCTC This is a Unicode transaction code. X'01' - IRM_F1_TRNEXP The user requests that IMS Connect set the expiration time for the input transaction. X'00' The user requests that no MFS MOD name to be returned. If this value is not supplied by the client, the user exit must use a default value. The MFS MOD name flag is returned to IMS Connect from the exit in the EXIT PARMLIST field, EXPREA_UFLAG1. The Unicode encoding schema (UTF8, UCS2, UTF16) is identified in the IRM_ES field of the fixed portion of the IRM.
IRM_F2	1 byte	33	X'21'	Specifies the commit mode. - X'40' - commit mode '0' (CM0), also known as commit-then-send - X'20' - commit mode '1' (CM1), also known as send-then-commit If this value is not supplied from the client, the user exit must use a default value. The commit mode flag is returned to IMS Connect from the exit in the OTMA header field, OMHDRSYN.

Table 18. User portion of IRM for HWSSMPL0, HWSSMPL1, and user-written user message exits (continued)

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_F3	1 byte	34	X'22'	Depending on the value, this field specifies the sync level, the purge or reroute option for CM0 output, or the serial delivery option for send-only input. • X'00' - sync level is 'NONE' • X'01' - sync level is 'CONFIRM' • X'02' - sync level is 'SYNCPT' • X'04' - purge undeliverable CM0 output (IRM_F3_PURGE) • X'08' - reroute undeliverable CM0 output (IRM_F3_REROUT) • X'10' - send-only with serial delivery (IRM_F3_ORDER). IRM_F3_ORDER invokes the serial (ordered) delivery option for a send-only transaction • X'20' - for multi-segment CM0 output messages, ignore the DL/I PURG call to the TP PCB (IRM_F3_IPURG) until all segments of the output message have been inserted to the TP PCB. This option is useful only if the IMS application issues a PURG call for each segment of a multi-segment message that it inserts to the TP PCB. • X'80' - cancel duplicate client ID (IRM_F3_CANCID). On the first request of a new session, if an existing session is using the same client ID as the new session, IRM_F3_CANCID terminates the existing session to allow the new session to proceed. For CM0, the sync level must be set to confirm. If the synch level is not supplied from the client, the user exit must use a default value. The sync level flag is returned to IMS Connect from the exit in the OTMA header field, OMHDRSLV. The purge not deliverable flag is returned to IMS Connect from the user message exit in the OTMA header field, OMHDCFL, with the setting of OMHDRPND X'10'. If the reroute flag is set, the IRM_REROUT_NM field is optional. You cannot specify the purge not deliverable function and the reroute function at the same time. If both functions are specified, the output messages are not purged or rerouted and OTMA issues message DFS2407W.

Table 18. User portion of IRM for HWSSMPL0, HWSSMPL1, and user-written user message exits (continued)

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_F4	1 byte	35	X'23'	The IRM_F4 flag identifies the type of message being sent by the client. Message types are specified by an ASCII or EBCDIC character value. The value is sent to IMS Connect , passed to the user exit, the exit builds the appropriate OTMA structure and returns it to IMS Connect to be forwarded to IMS. The valid values and the message types that they indicate are as follows: A ACK (IRM_F4_ACK) - An ACK response to output received from IMS Connect. ACK is used by a client to indicate the acceptance of either: • A synchronous callout request message issued by an IMS application program • An output message when the original input message from the client specifies a SYNC level of CONFIRM C Cancel IRM timer (IRM_F4_CANTIMER) - A request to cancel the IRM timer for another connection on which the client, using the same client ID, is waiting for output data. D Deallocate (IRM_F4_DEALLOC) - A request to deallocate the conversation. K Send only requires ACK (IRM_F4_SNDONLYA) - A send-only transaction message that requires an ACK response from IMS Connect. M Synchronous callout response message (IRM_F4_SYNRESP) - A synchronous callout response message that uses the send-only protocol. N NAK (IRM_F4_NACK) - A NAK response from the client that indicates the rejection of one of the following types of output from IMS Connect : • A synchronous callout request message issued by an IMS application program • An output message when the original input message from the client specifies a SYNC level of CONFIRM R RESUME TPIPE (IRM_F4_RESUMET) - A RESUME TPIPE call for asynchronous output data from IMS. A RESUME TPIPE call must execute on a transaction or persistent socket using CM0. S Send only (IRM_F4_SENDOONLY) - A send only transaction message that starts a send-only interaction for a non-response mode, non-conversational transaction. If the host application terminates without issuing an ISRT to the IO PCB, no DFS2082 messages are returned to the client. The SENDONLY interaction must use CM0. blank (X'40') A send-receive interaction for a conversational or non-conversational response mode transaction.
IRM_TRNCOD	8 bytes	36	X'24'	Character string. It specifies the IMS transaction code.

Table 18. User portion of IRM for HWSSMPL0, HWSSMPL1, and user-written user message exits (continued)

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_IMSDESTID	8 bytes	44	X'2C'	Character string. It specifies the data store name (IMS destination ID). This field must be specified by the client. The data store name is returned to IMS Connect from the exit in the OTMA header field, OMUSR_DESTID.
IRM_LTERM	8 bytes	52	X'34'	Character string. It specifies the IMS logical terminal (LTERM) override name. This field can be set to a valid name or to blanks. The LTERM override name is returned to IMS Connect from the exit in the OTMA header field, OMHDLRM. For IMS host applications, the value for this field is set by the user message exit, which either moves this value to the OTMA field OMHDLRMT or sets OMHDLRMT with a predetermined value. If you have specified an LTERM override value, OTMA places that value in the IOPCB LTERM field. If you do not specify an LTERM override value, OTMA instead places the IMS Connect-defined tpipe name in the IOPCB LTERM field. The tpipe name is set to the CLIENT ID if the commit mode is zero; it is set to the PORT ID if the commit mode is one. If you use the LTERM value in the IOPCB to make logic decisions, be aware of the naming conventions of the IOPCB LTERM name.
IRM_RACF_USERID	8 bytes	60	X'3C'	Character string. It specifies the RACF user ID. The client must provide it if RACF is to be used. The RACF user ID name is returned to IMS Connect from the exit in the OTMA header field, OMSECUID.
IRM_RACF_GRPNAME	8 bytes	68	X'44'	Character string. It specifies the RACF group name. The RACF group name is returned to IMS Connect from the exit in the OTMA header field, OMSECGRP. To preserve the expected offset of this field, when you specify IRM_RACF_GRPNAME, you must also specify the following fields with valid values or blanks: IRM_RACF_USERID
IRM_RACF_PW	8 bytes	76	X'4C'	Character string. It specifies the RACF PassTicket or password. The client must provide it if RACF is to be used. The PassTicket or password value is returned to IMS Connect from the user message exit, in the OTMA header field, OMUSR_PASSTICK. To preserve the expected offset of this field, when you specify IRM_RACF_PW, you must also specify the following fields with valid values or blanks: - IRM_RACF_USERID - IRM_RACF_GRPNAME

Table 18. User portion of IRM for HWSSMPL0, HWSSMPL1, and user-written user message exits (continued)

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_APPL_NM	8 bytes	84	X'54'	Character string. It specifies the RACF APPL name, that was defined to RACF on the PTKTDATA definition. If this field is not included or is blank, the RACF APPL name is taken from the APPL parameter of the DATASTORE configuration statement. To preserve the expected offset of this field, when you specify IRM_APPL_NM, you must also specify the following fields with valid values or blanks: • IRM_RACF_USERID • IRM_RACF_GRNAME • IRM_RACF_PW
IRM_REROUT_NM	8 bytes	92	X'5C'	Character string (A through Z, 0 - 9, or special characters, such as @, #, \$). The reroute tpipe name of the client reroute request. This field is optional. Recommendation: Use blanks for the default value. When you specify the IRM_REROUT_NM field, you must also specify: • IRM_F3_REROUT in the IRM_F3 field in the user portion of the IRM • IRM_ARCH1 or IRM_ARCH2 in the IRM_ARCH field in the common fixed portion of the IRM To preserve the expected offset of this field, when you specify IRM_REROUT_NM, you must also specify the following fields with valid values or blanks: • IRM_RACF_USERID • IRM_RACF_GRNAME • IRM_RACF_PW • IRM_APPL_NM IRM_REROUT_NM and IRM_RT_ALTCID use the same offset. Both cannot be specified on the same message.
IRM_RT_ALTCID	8 bytes	92	X'5C'	Alternate client ID for a RESUME TPIPE call. This is an optional field. If the alternate client ID is provided, the IRM_ARCH field must be set to IRM_ARCH1. To preserve the expected offset of this field, when you specify IRM_RT_ALTCID, you must also specify the following fields with valid values or blanks: • IRM_RACF_USERID • IRM_RACF_GRNAME • IRM_RACF_PW • IRM_APPL_NM

Table 18. User portion of IRM for HWSSMPL0, HWSSMPL1, and user-written user message exits (continued)

Field	Length	Offsets		Meaning
		Dec	Hex	
IRM_TAG_ADAPT	8 bytes	100	X'64'	Name of the adapter that IMS Connect calls to convert XML messages to and from select programming languages. To preserve the expected offset of this field, when you include the IRM_TAG_ADAPT field, you must also specify the following fields with valid values or blanks: • IRM_RACF_USERID • IRM_RACF_GRNAME • IRM_RACF_PW • IRM_APPL_NM • IRM_REROUT_NM or IRM_RT_ALTCID
IRM_TAG_MAP	8 bytes	108	X'6C'	Name of the converter the XML adapter calls to perform that actual conversion of XML messages to and from select programming languages. To preserve the expected offset of this field, when you include the IRM_TAG_MAP field, you must also specify the following fields with valid values or blanks: • IRM_RACF_USERID • IRM_RACF_GRNAME • IRM_RACF_PW • IRM_APPL_NM • IRM_REROUT_NM or IRM_RT_ALTCID • IRM_TAG_ADAPT
IRM_MODNAME	8	116	74	The MFS MOD name for input messages.
IRM_CORTKN	0XL40			The correlation token for a synchronous callout message. The correlation token includes information that correlates a reply to a synchronous callout request with the IMS system and IMS application program that issued the request.
IRM_CT_LEN	2	124	7C	The length of the correlation token.
IRM_CT_RESV1	2	126	7E	Reserved.
IRM_CT_IMSID	4	128	80	The IMSID of the IMS system that scheduled the IMS application that issued the synchronous callout request.
IRM_CT_MEMTK	8	132	84	The OTMA TMEMBER token.
IRM_CT_AWETK	8	140	8C	OTMA message token.
IRM_CT_TPIPE	8	148	94	OTMA TPIPE name.
IRM_CT_USERID	8	156	9C	The user ID specified in the ICAL call.
ORG IRM_MAP				
IRM_XID	XL140	32	20	Global transaction ID for two-phase commit processing.

Output structure from client exit

The following table shows the structure (one occurrence per message) of the message returned by the user-written client application exit routine. The table lists the field name, the length of the field, and a brief explanation of the field.

Table 19. Message structure returned by user-written IMS Connect client application exit routine

Field	Length	Meaning
BPE header	64 bytes	Defined in “BPE header format” on page 250.
OTMA structure	Total length of OTMA header	For more information about the HWSOMPFEX macro (full OTMA structure) see, “Macros that support IMS Connect exit routines” on page 186.
LLZZTRANCODEDATA	n bytes	• LL - length of segment • ZZ - set to binary zeros • TRANCODE - IMS 1-8 byte transaction code • DATA - user data
LLZZDATA	n bytes	• LL - length of segment • ZZ - set to binary zeros • DATA - user data The LLZZDATA is repeated to a maximum of 32 KB overall length. If there is more data, then the structures continues as shown in “Other IMS Connect structures.” Only the first segment will contain the IMS transaction code and the following segment will contain the segment data necessary for the transaction to process.
LL	2 bytes	LL - set to binary zeros to denote the end of this structure. The LL field is not part of the segment length.

Related reference:

Chapter 18, “OTMA header fields used by IMS Connect,” on page 265

Other IMS Connect structures

In addition to the IRM, messages handled by IMS Connect contain other recurring structures.

The following table shows other structures that are repeated until all data has been mapped to be returned to IMS Connect. The table lists the field name, the length of the field, and a brief explanation of the field.

Table 20. Other repeated IMS Connect message structures

Field	Length	Meaning
BPE header	64 bytes	Defined in “BPE header format” on page 250.
OTMA structure	32 bytes	The OTMA message control structure that is mapped by the IMS Connect macro HWSOMPFEX.
LLZZDATA	n bytes	• LL - length of segment • ZZ - set to binary zeros • DATA - user data The LLZZDATA field is repeated to a maximum of 32 KB overall length. If there is more data, then the structures continue.

Table 20. Other repeated IMS Connect message structures (continued)

Field	Length	Meaning
LL	2 bytes	LL - set to binary zeros to denote the end of this structure.

Related reference:

“OTMA message-control fields used by IMS Connect” on page 265

BPE header format

The BPE header is a recurring structure in messages that are handled by IMS Connect.

The following table lists the fields in the BPE header layout. **This topic contains Product-sensitive Programming Interface information.**

Table 21. BPE header layout

Field	Length	Meaning
IIII	4 bytes	The length of the total structure and it is set for the first BPE header only. This field is managed by IMS Connect and must not be altered by the exit.
CHAIN PTR	4 bytes	The chain pointer to the next BPE header within this message. The last BPE header in the message must have binary zeros as a chain pointer value to denote the end of the BPE headers within the message. These chain pointers are set by the user-written client application exit routine.
STORAGE TYPE	8 bytes	This field is managed by IMS Connect and should not be modified by the user exit.
TYPE ACCESS	4 bytes	This field is managed by IMS Connect and should not be modified by the user exit.
SUBPOOL	1 byte	This field is managed by IMS Connect and should not be modified by the user exit.
Reserved	43 bytes	This field is managed by IMS Connect and should not be modified by the user exit.

Message structures and IMS Connect user message exit routines

IMS Connect allows up to 254 user exits to be defined in the configuration file. There are two input message structures supported by IMS Connect and two message structures supported on return from a user exit.

Input messages from client

The structure of input messages passed to and returned from a user message exit routine differs based on a number of factors, including whether the user message exit routine supports the IMS TM Resource Adapter or a user-written IMS Connect client.

The following table shows the structure for input messages received from the client by IMS Connect. The table provides information about the input message structure type, if the OTMA header is present or not, if the exit data is translated by the client code, the exit type flag, and the supporting message type.

Table 22. Input message structure

Input message structure type	OTMA header present	Exit data translated by client code	Exit type flag (IRMHDR_FLG5)	Supporting message type
1	Y	Y	11000000	HWSJAVA0
1	Y	N	10000000	HWSSMPL0 and HWSSMPL1 modified not to build OTMA headers when the client/server builds OTMA headers
2	N	Y	01000000	HWSSMPL0 and HWSSMPL1 modified not to translate data
2	N	N	00000000	HWSSMPL0 HWSSMPL1 HWSCSLO0

The following table shows the structure for input messages returned by the exit based on the input structure received by the exit. The table provides information about the input message structure type, the exit output message structure type, the exit type flag, and the supporting message type.

Table 23. Input message structure returned by the exit

Input message structure type	Exit output message structure type	Exit type flag (IRMHDR_FLG5)	Supporting message type
1	1	11000000	HWSJAVA0
1	1	10000000	HWSSMPL0 and HWSSMPL1 modified not to build OTMA headers when the client/server builds OTMA headers
2	3	01000000	HWSSMPL0 and HWSSMPL1 modified not to translate data
2	3	00000000	HWSSMPL0 and HWSSMPL1

Output message to client

The output message from IMS is passed to the user message exit routine that was called from the client.

The user exit normally removes the OTMA headers for output if the exit added the OTMA headers for input. The user exit normally translates the data from EBCDIC to ASCII if it did the translation for input. And the reverse is true if these things were not done for input.

The OTMA header can consist of up to four sections and application data. If the exit is to remove the OTMA header (not present on input), there must be a check for each section. The four sections include:

- Control (always present in the OTMA structure)
- Header (might or might not exist in the OTMA structure)
- Security (might or might not exist in the OTMA structure)
- User (might or might not exist in the OTMA structure)

Output message from IMS-to-IMS Connect

All output messages received by IMS Connect from IMS consist of the same structure, the OTMA header followed by LLZZ DATA. If the message contains multiple segments, then the OTMA header and LLZZ DATA are repeated for the number of segments in the message.

For CM0 multi-segment output messages, if the IMS application program issues a DL/I PURG call after it inserts each segment to the TP PCB, each segment of the output message is sent to the IMS Connect client separately.

If your IMS application issues a PURG call for each segment of multi-segment CM0 output messages, you can instruct OTMA to ignore PURG calls until the last segment of the message is inserted to the TP PCB by specifying X'02' (OMHDRIPG) in the OMHDCFL field of the state data section of the OTMA message header. When OMHDRIPG is specified, OTMA returns the output message to IMS Connect as a single multi-segment message.

User-written IMS Connect clients can set the OMHDRIPG flag by specifying X'20' (IRM_F3_IPURG) in the IRM_F3 field of the user portion of the IRM prefix of input transaction messages.

Output message from IMS returned by the exit back to IMS Connect

The message returned to IMS Connect from the exit consists of one of two structures:

- Messages with OTMA structures imbedded in the message
- Message with no OTMA structures imbedded in the message

Message structures

The required structure of messages that flow through IMS Connect differs depending on whether the message is an input message or an output message, and whether the message has been modified by the user message exit yet. The following topics describe the message structures for the IMS TM Resource Adapter and user-written IMS Connect client applications.

Input message from client and passed to message exit

Input messages from the client consist of the IMS TM Resource Adapter and user-written IMS Connect client application message structures.

This topic contains Product-sensitive Programming Interface information.

IMS TM Resource Adapter message structure - type 1

The following table shows the input message format supported by IMS Connect from an IMS TM Resource Adapter client.

Table 24. IMS Connect message format for IMS TM Resource Adapter client applications

Field	Length	Meaning
llll	4 bytes	Length of entire message, including llll field.
IRM	28 bytes	The common, fixed portion of the IRM.
IRM_MAP	0 bytes	Position of the XID map when it is present.
OTMA HDRs	466 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMACT data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of data segment.
zz	2 bytes	Reserved (set to binary zeros).
DATA	n bytes	User data with the transaction code first.
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMACT data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of second data segment.
zz	2 bytes	Reserved (set to binary zeros).
DATA	n bytes	User data in the second data segment (no transaction code).
...		
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMACT data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of this and last data segment.
zz	2 bytes	Reserved (set to binary zeros).
DATA	n bytes	User data with this data segment (no transaction code).
	ORG IRM_MAP	
IRM_XID	140 bytes	Global transaction ID for two-phase commit processing.

User-written message structure - type 2

The following table shows the input message format supported by IMS Connect for user-written client applications.

The message format is supported by the HWSSMPL0 and HWSSMPL1 user message exit routines, as well as user-written user message exit routines. This message format is not supported by HWSJAVA0, HWSSOAP1, HWSCSLO0, or HWSCSLO1.

If you are using a user-written message exit routine, you can structure the message fields as required by the user exit. You can include fields that are used only by the user message exit or other fields that can be passed in the OTMA headers, such as the MID name.

The following is the data structure for all user-written client applications.

The transaction data in input messages from user-written IMS Connect client applications follows the user-defined IRM and is contained in one or more data segments. The data segments must conform to the format of LLZZDATA, where LL contains the total length of the data segment, ZZ contains binary zeros, and DATA contains the transaction code and transaction data or an IMS command.

IMS command input must be submitted in a single data segment followed by EOM.

Table 25. IMS Connect message format for user-written client applications

Field	Length	Meaning
llll	4 bytes	Length of entire message, including llll field.
Fixed IRM prefix	28 bytes	The common, fixed portion of the IRM mapped by the HWSIMSCB macro in the IMS.SDFSMAC data set. For a description of the fields see "Format of fixed portion of IRM in messages sent to IMS Connect" on page 238.
User-defined IRM	Variable length	User-defined part of the IRM mapped by the HWSIMSCB macro in the IMS.SDFSMAC data set. For a description of the fields see "Format of user portion of IRM for HWSSMPL0, HWSSMPL1, and user-written message exit routines" on page 241. HWSJAVA0, HWSSOAP1, HWSCSLO0, and HWSCSLO1 do not support the user-defined part of the IRM.
LL	2 bytes	Length of first data segment.
zz	2 bytes	Reserved (set to binary zeros).
DATA	n bytes	User data with the transaction code first.
LL	2 bytes	Length of second data segment.
zz	2 bytes	Reserved (set to binary zeros).
DATA	n bytes	User data second data segment (no transaction code).
...		
LL	2 bytes	Length of this data segment.
zz	2 bytes	Reserved (set to binary zeros).
DATA	n bytes	User data segment (no transaction code).
LL	2 bytes	End of message (set to binary 0000 0000 0000 0100).
zz	2 bytes	Reserved (set to binary zeros).

Related reference:

"Format of fixed portion of IRM in messages sent to IMS Connect" on page 238

Input message returned from message exit

Input messages from the message exit consist of the IMS TM Resource Adapter and user-written IMS Connect client application message structures.

IMS TM Resource Adapter message structure - type 1

The IMS TM Resource Adapter exit output message format that is supported by IMS Connect is the same message format of the input message.

The total length of the message can be 10,000,000 bytes. The length of each segment (from the BPE header to the next BPE header) within the message can be a maximum of 32 KB, excluding the BPE and OTMA headers.

User-written IMS Connect client applications message structure - type 3

The following table shows the output message format supported by IMS Connect from the supplied HWSSMPL0 exit (user-written IMS Connect client application exit routine). The table provides information about the field, length, and meaning. Variable length OTMA headers are supported, and therefore, the OTMA header length can be other than 466 bytes. The following example contains 466 bytes as used by the supplied exits.

Table 26. Supported output message format for the HWSSMPL0 exit

Field	Length	Meaning
BPE HEADER	64 bytes	For the format of the BPE header, see "BPE header format" on page 256.
OTMA HDRs	466 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMA data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of first data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data with the transaction code first
Repeat of ll,zz,DATA		
LL	2 bytes	Length of this data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data (no transaction code)
yy	2 bytes	Binary value of zero
BPE HEADER	64 bytes	For the format of the BPE header, see "BPE header format" on page 256.
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMA data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of this data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data segment (no transaction code)
Repeat of LL,zz,DATA		
LL	2 bytes	Length of this data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data (no transaction code)
...		
LL	2 bytes	Length of this data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data (no transaction code)
yy	2 bytes	Binary value of zero

Table 26. Supported output message format for the HWSSMPL0 exit (continued)

Field	Length	Meaning
BPE HEADER	64 bytes	For the format of the BPE header, see “BPE header format.”
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFS MAC data set. For a description of the fields see Chapter 18, “OTMA header fields used by IMS Connect,” on page 265.
LL	2 bytes	Length of this data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data (no transaction code)
...		
LL	2 bytes	Length of this data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data (no transaction code)
yy	2 bytes	Binary value of zero

Restriction: The length of data from one BPE header to the next BPE header cannot exceed 32K, excluding the BPE header and the OTMA header.

BPE header format

The following table describes length and meaning of the fields in the BPE header format.

Restriction: User message exit routines must not modify any fields in the BPE header except the chain pointer field. User message exit routines modify the chain pointer field to chain the BPE headers together with the last BPE chain pointer set to binary zeros.

This topic contains Product-sensitive Programming Interface information.

Table 27. BPE header format

Field	Length	Meaning
llll	4 bytes	Length of field of entire buffer
CHAIN PTR	4 bytes	Chain pointer to next BPE header
STORAGE TYPE	8 bytes	Storage type
TYPE ACCESS	4 bytes	Type access
SUBPOOL	1 byte	Subpool
RESV	43 bytes	Reserved

Related reference:

“Input message from client and passed to message exit” on page 252

Output message passed to message exit

Output messages from IMS Connect are passed to the user message exit in a certain message structure, regardless of whether the exit is for an IMS TM Resource Adapter client or a user-written IMS Connect client application.

The following table shows the message format from IMS Connect to the exit for the client. The table provides information about the length and meaning of the fields in the output message.

Table 28. Format of output messages from IMS Connect to the user message exit

Field	Length	Meaning
OTMA HDRs	Length of total OTMA headers	Mapped by the HWSOMPFX macro in the IMS.SDFSMA data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMA data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of this segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data
...		
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMA data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of this segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data

Related reference:

 [IMS TM Resource Adapter user message exit routine \(HWSJAVA0\) \(Exit Routines\)](#)

Output message from message exit to client

Depending on the type of IMS Connect client and the user message exit used to support the client, the format of the message structure differs.

This topic contains Product-sensitive Programming Interface information.

The output message from the message exit that is sent to IMS Connect client applications is in one of the following formats:

- "IMS TM Resource Adapter message structure"
- "User-written application message structure" on page 258

IMS TM Resource Adapter message structure

The following table shows the message format of the output message from the user message exit HWSJAVA0. The table provides information about the length and meaning of the fields in the output message.

Table 29. Output message format from the user message exit HWSJAVA0

Field	Length	Meaning
LLLL	4 bytes	Total message length
Id	8 bytes	*HWSJAV*
OTMA HDRs	466 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMAc data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMAc data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of this data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data
...		
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMAc data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of this segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data
OTMA CTL HDR	32 bytes	Mapped by the HWSOMPFX macro in the IMS.SDFSMAc data set. For a description of the fields see Chapter 18, "OTMA header fields used by IMS Connect," on page 265.
LL	2 bytes	Length of data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data
...		
LL	2 bytes	Length of this data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data

User-written application message structure

The user-written IMS Connect client application message structure can consist of one or more TCP/IP message structures.

Message structures returned to the client by the sample user message exit routine HWSSMPL1 (*SAMPL1*) start with LLLL, a 4-byte field that contains the total length of the message. The sample user message exit routine HWSSMPL0 (*SAMPLE*) does not add the LLLL field at the beginning of the message structures that it returns.

The output message from the message exit that is sent to user-written IMS Connect client applications is in one of the following formats:

Message format when the MFS MOD name is requested

The following table shows the format of the output message sent to the client application when the MFS MOD name is requested. The table provides information about the length and meaning of the fields in the output message.

Table 30. Output message format containing RMM, DATA, and CSM

Field	Length	Meaning
LLLL (HWSSMPL1 only)	4 bytes	The total length of an output message returned by HWSSMPL1; output messages returned by HWSSMPL0 do not include this field
RMM header (optional)	20 bytes	The first four bytes of the RMM are LLzz, the two byte length of the RMM and two bytes of binary zeros
LL	2 bytes	Length of data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data
LL	2 bytes	Length of second data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data second data segment
...		
LL	2 bytes	Length of nth segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data nth data segment
CSM	12 bytes	Complete status message

Message format when the MFS MOD name is not requested

The following table shows the format of the output message sent to the client application when the MFS MOD name is not requested and only data and the CSM are being sent. The table provides information about the length and meaning of the fields in the output message.

Table 31. Output Message format containing output data and CSM only

Field	Length	Meaning
LLLL (HWSSMPL1 only)	4 bytes	The total length of an output message returned by HWSSMPL1; output messages returned by HWSSMPL0 do not include this field
CORTKN header (applicable to synchronous callout request)	52 bytes	Correlation token for this synchronous callout request

Table 31. Output Message format containing output data and CSM only (continued)

Field	Length	Meaning
LL	2 bytes	Length of data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data
LL	2 bytes	Length of second data segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data second data segment
...		
LL	2 bytes	Length of nth segment
zz	2 bytes	Reserved (set to binary zeros)
DATA	n bytes	User data nth data segment
CSM	12 bytes	Complete status message

Format of an IMS command response message

The output message sent to a client submitting commands to the Operations Manager (OM) component of the IMS Common Service Layer (CSL) is shown in the following table. The table provides information about the length and meaning of the fields in the output message.

Table 32. Output message format sent to OM command client

Field	Length	Meaning
LLLL	4 bytes	Length of output
DATA	n bytes	IMS command output

Related information:

Format of the request mod message:

The IMS Connect request mod message (RMM) is returned as the first structure of an output message from the message exit if the MFS MOD name is requested and the data output is present.

The following table shows the output message format of the request mod message built by the user message exits HWSSMPL0 and HWSSMPL1. The table includes the field name, field length, and field meaning.

Table 33. Request mod message output message format

Field	Length	Meaning
LL	2 bytes	Length of RMM message
zz	2 bytes	Reserved (set to binary zeros)
ID	8 bytes	Character value of *REQMOD*
MOD	8 bytes	Character value of the requested MFS MOD name

Format of the complete status message:

The IMS Connect complete status message (CSM) is returned as the last structure of an output message from the message exit if the input message is processed successfully.

The following table shows the output message format of the CSM. The table includes the field name, field length, and field meaning.

Table 34. Complete status message output message format

Field	Length	Meaning
CSM_LEN	2 bytes	Length of CSM message
CSM_FLG1	1 byte	Flag byte: X'80' Asynchronous message queued in IMS X'40' Conversational output message X'20' ACK/NAK required
Reserved	1 byte	Reserved (set to binary zeros)
CSM_ID	8 bytes	Character value of *CSMOKY*

Format of the request status message:

The IMS Connect request status message (RSM) is returned as the only structure of an output message from the message exit if IMS Connect or the message exit determined an error occurred.

The following table shows the output message format of the request status message for an error condition. The table includes the field name, field length, field meaning.

Table 35. Request status message output message format

Field	Length	Meaning
RSM_LEN	2 bytes	Length of RSM message
RSM_FLG1	1 byte	Flag byte: X'80' Asynchronous message queued in IMS X'40' Conversational output message X'20' ACK/NAK required
RSM_RACFRC or RSM_OTMRSN	1 byte	If this RSM is generated by a RACF security error, this is the RSM_RACFRC field, which contains the return code from a RACROUTE REQUEST=VERIFY command. Otherwise, this is the RSM_OTMRSN field, which contains a reason code from OTMA.
RSM_ID	8 bytes	Character value of *REQSTS*
RSM_RETCOD	4 bytes	Return code
RSM_RSNCOD	4 bytes	Reason code

Format of the correlation token for synchronous callout messages from IMS applications:

The output messages from the IMS Connect message exit for synchronous callout requests include a correlation token to correlate the reply to the callout request with the IMS application that issued the request.

The correlation token is the first structure in a synchronous callout request and can be followed by the optional request mod message (RMM) structure and the application data.

The following table shows the structure of the correlation token on output from IMS Connect, as mapped by the CORMask DSECT in the HWSIMSCB macro:

Table 36. Format of the correlation token for synchronous callout requests in output messages

Field	Length	Meaning
COR_Len	2 bytes	Length of correlation structure, including the structure header and the correlation token
COR_Rsvd	2 bytes	Reserved (set to binary zeros)
COR_Id	8 bytes	Character value of *CORTKN*
COR_LL	2 bytes	Length of the correlation token
COR_RESV1	2 bytes	Reserved
COR_IMSID	4 bytes	ID of the IMS system in which the IMS application program that issued the callout request is running
COR_MEMTK	8 bytes	OTMA TMEMBER token
COR_AWETK	8 bytes	OTMA message token
COR_TPIPE	8 bytes	OTMA Tpipe name
COR_USERID	8 bytes	User ID included in the ICAL call by the IMS application that issued the callout request

Related tasks:

“Configuring user-written IMS Connect clients for synchronous callout requests” on page 223

Examples of message structures in a simple interaction

The following examples show the message structures in a simple interaction between IMS Connect and a user-supplied client application program.

The flow represented in each example is a simple send-receive transaction that uses commit-mode 1 and sync level=confirm, as defined by the following IRM values:

- IRM_F2 = X'20' (for CM1)
- IRM_F3 = X'01' (for sync level = confirm)
- IRM_F4 = X'40' (for a send-receive interaction)

HWSSMPL0 message structures in a simple interaction

The example flow shown in the following table is a simple send-receive transaction between a user-supplied IMS application program and IMS Connect when the sample user message exit routine HWSSMPL0 is used (IRM_ID = *SAMPLE*).

In the table:

- The flow starts at the top of the table and progresses down.
- The vertical bars separate the various structures within each message.
- The input transaction and transaction response can have one or more LLZZ data structures.
- The LL RMM structure shown in the transaction response is optional.
- 0004 0000 is the 4 byte end-of-message structure.

Table 37. Example message structures for HWSSMPL0 in a simple interaction

Client state	Example message structure
Client send (input transaction)	LLLL LL IRM LLZZ data LLZZ data 0004 0000
Client receive (transaction response)	LL RMM LLZZ data LLZZ data LL CSM
Client send (ACK)	LLLL LL IRM 0004 0000
Client receive (deallocate confirm)	LL RSM

HWSSMPL1 message structures in a simple interaction

The example flow shown in the following table is a simple send-receive transaction between a user-supplied client application program and IMS Connect when the sample user message exit routine HWSSMPL1 is used (IRM_ID = *SAMPL1*).

In the table:

- The flow starts at the top of the table and progresses down.
- The vertical bars separate the various structures within each message.
- The input transaction and transaction response can have one or more LLZZ data structures.
- The LL RMM structure shown in the transaction response is optional.
- 0004 0000 is the 4 byte end-of-message structure.

Table 38. Example message structures for HWSSMPL1 in a simple interaction

Client state	Example message structure
Client send (input transaction)	LLLL LL IRM LLZZ data LLZZ data 0004 0000
Client receive (transaction response)	LLLL LL RMM LLZZ data LLZZ data LL CSM
Client send (ACK)	LLLL LL IRM 0004 0000
Client receive (deallocate confirm)	LLLL LL RSM

Chapter 18. OTMA header fields used by IMS Connect

IMS Connect uses fields in the headers of messages sent to OTMA to communicate processing options and other information to IMS.

OTMA defines the format and valid values of the message headers. The fields of the OTMA message header are grouped into sections based on the purpose of the fields.

Depending on the type of message being sent or the options you select, the format of a given section of the OTMA header can change. For example, depending on the type of message being submitted, the format of the message-control section of the OTMA header differs.

The tables in the following subsections document the IMS Connect requirements for setting the OTMA header fields or integrating them with the IMS Connect user message exits. The tables correspond to the sections of the OTMA header. Additional notes for the fields listed in each table are denoted by a number in the final column. The corresponding note text is listed by number in “Notes to OTMA header tables” on page 289.

Related reference:

Chapter 47, “OTMA message prefix,” on page 855

OTMA message-control fields used by IMS Connect

The table in this topic defines the fields of the OTMA message control header and the order of those fields.

The numbered notes in the table are explained in Notes to OTMA header tables (Communications and Connections).

Table 39. HWSOMCTL DSECT - OTMA message-control header. Control data common section for all messages

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMCTLALV	1	0		ARCHITECTURE LEVEL Set to X'01' architecture level 1. Set for all messages.	1
OMCTLMGT	1	1	OMCTLDTA X'80'	MESSAGE TYPE=Data Set for conversational transactions but not on first input. If EXPREA FLAG1 is set to EXPREA_ CONVERS then set OMCTLMGT to OMCTLDTA. EXPREA FLAG1 is not set to EXPREA_ CONVERS on the first input for conversation.	1

Table 39. HWSOMCTL DSECT - OTMA message-control header. Control data common section for all messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
			OMCTLTXN X'40'	MESSAGE TYPE=Transaction Set for first transaction input. That is, first input for conversation or nonconversation, EXPREA_FLAG1 is not set to EXPREA_CONVERS.	1
			OMCTLRSP X'20'	MESSAGE TYPE=Response Set for: • ACK response to msg sent to client • NAK response to msg sent to client Required for: • Commit Mode 0 (synch level=CONFIRM) • Commit Mode 1 (synch level=CONFIRM)	1
			OMCTLCMD X'10'	MESSAGE TYPE=Command Set for - RESUME TPIPE calls	1
			OMCTLCMT X'08'	MESSAGE TYPE = Commit Confirmation Set for SEND ONLY or DEALLOCATE. SEND ONLY or DEALLOCATE is indicated in IRM from client.	1

Table 39. HWSOMCTL DSECT - OTMA message-control header. Control data common section for all messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMCTLRSI	1	2	OMCTLACK X'80'	RESPONSE INDICATOR RESPONSE = ACK Set for ACK. ACK is indicated in IRM.	
			OMCTLNAK X'40'	RESPONSE = NAK Set for NAK. NAK is indicated in IRM.	1
			OMCTLRQ X'20'	RESPONSE = Response requested If set, then conversational transaction and the IMSEA_RSNCODE must be set to 96 (X'60') to signal client application that conversation continues.	1
			OMCTLERQ X'10'	RESPONSE=Extended response requested NEITHER TESTED NOR SET BY EXIT.	4
			OMCTLSYR X'08'	Response to a synchronous callout message	
			OMCTLDAN X'02'	Support for delayed ACK or NAK response	
OMCTLCCI	1	3	OMCTLCTD X'80'	COMMIT CONFIRMATION INDICATOR Confirm=Committed If set, then the IMS application has terminated the conversation, and the IMSEA_RSNCODE must be set to 97 (X'61') to signal client application that the IMS application terminated successfully.	
			OMCTLABT X'40'	Confirm=Aborted NEITHER TESTED NOR SET BY EXIT.	4
OMCTLTYP	1	4	OMCTLBID X'04'	COMMAND TYPE COMMAND=Client Bid NEITHER TESTED NOR SET BY EXIT.	4
			OMCTLAVL X'08'	COMMAND=Server Available NEITHER TESTED NOR SET BY EXIT.	4

Table 39. HWSOMCTL DSECT - OTMA message-control header. Control data common section for all messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
			OMCLTRSN X'0C'	Command=Resynch NEITHER TESTED NOR SET BY EXIT.	4
			X'10'	Reserved for future use. Neither tested nor set by exit.	4
			OMCTLSPA X'14'	Command=Suspend I/P for all tpipes. Neither tested nor set by exit.	4
			OMCTLRSA X'18'	Command=Resume I/P for all tpipes. Neither tested nor set by exit.	4
			OMCTLSPN X'1C'	Command=Suspend I/P for named tpipe. Neither tested nor set by exit.	4
			OMCTLRSM X'20'	Command=Resume I/P for named tpipe. Neither tested nor set by exit.	4
			OMCTLRTP X'24'	Command=Resume O/P for named tpipe without options. Set for RESUME TPIPE call without options.	1
			OMCTLRID X'28'	Command=Resume single tpipe with options. Set for RESUME TPIPE call with options.	1
			OMCTLMTR X'3C'	Command=Resource state protocol command. OTMA sends this command to notify IMS Connect and other OTMA clients about how well the IMS system is processing the input received from OTMA clients. IMS Connect updates the exit interface block data store entry (HWSXIBDS) with the information sent in this command.	

Table 39. HWSOMCTL DSECT - OTMA message-control header. Control data common section for all messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMCTLPFG	1	5	OMCTLLPG X'80'	PROCESSING FLAG	4
				Load Program	
				NEITHER TESTED NOR SET BY EXIT.	
			OMCTLSYP X'40'	Synchronized tpipe.	4
				NEITHER TESTED NOR SET BY EXIT.	
			OMCTLASY X'20'	Asynchronous/ unsolicited queued messages.	4
			OMCTLERR X'10'	NEITHER TESTED NOR SET BY EXIT.	4
				There is an error message with the NAK. NEITHER TESTED NOR SET BY EXIT.	
			OMCTLQUE X'08'	Asynchronous message is in IMS Hold Queue.	4
				If set, set CSM_FLG1 to CSM_AMSG if sending CSM, orset RSMFLG1 to RSM_AMSG if sending RSM.	
			OMCTLOME X'01'	SCI not present error message.	
OMCTLTNM	8	6		Tpipe name. NEITHER TESTED NOR SET BY EXIT.	4
OMCTLCHN	1	E	OMCTLFIC X'80'	CHAIN STATE FLAG	1
				First in chain. Set for first message segment in chain.	
			OMCTLMIC X'40'	Middle in chain. Set for not first and, or, not last message segment in chain.	1
			OMCTLLIC X'20'	Last in chain. Set for last message segment in chain.	1
			OMCTLCAN X'10'	Cancel this message. Neither tested nor set by exit.	4
OMCTLPFL	1	F	OMCTLSTD X'80'	PREFIX FLAG	1
				State Data is present. Set if State Data Header present in OTMA Headers being built.	
			OMCTLSEC X'40'	Security data is present. Set if Security Data Header present in OTMA Headers being built.	1

Table 39. HWSOMCTL DSECT - OTMA message-control header. Control data common section for all messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
			OMCTLUSR X'20'	User data is present. Set if User Data Header present in OTMA Headers being built.	1
			OMCTLAPP X'10'	Application data is present. Set if Application Data Header present in OTMA Headers being built.	1
OMCTLSSN	4	10		SEND SEQUENCE NUMBER. NEITHER TESTED NOR SET BY EXIT.	4
OMCTLSNS	4	14		SENSE CODE. See OMCTLSNC and OMCTLRSC, which follow.	
		ORG OMCTLSNS			
OMCTLSNC	2	14		SENSE CODE If nonzero value, then build a NAK RSM to send to the client application, pass the sense code in the RSM as the reason code, and set the return code to X'0C'.	1
OMCTLRSC	2	16		REASON CODE NEITHER TESTED NOR SET BY EXIT.	4
OMCTLRSQ	4	18		RECOVERABLE MESSAGE SEQUENCE NUMBER NEITHER TESTED NOR SET BY EXIT.	4
OMCTLSEQ	2	1C		SEGMENT SEQUENCE NUMBER Set to 1 in first OTMA Control Header and count maintained in user work area. Increment by 1 for each subsequent OTMA Control Header within a single message being sent to IMS.	1
	1	1E		RESERVED.	3
	1	1F		RESERVED.	3
		ORG OMCTLRSQ			
OMCTLUID	8	18		The user ID of the client that is submitting a RESUME TPIPE call	

Related reference:

“Message-control information section” on page 855

OTMA state data fields used by IMS Connect

The tables in this topic describe the fields of the OTMA state data header and the order of those fields.

The numbered notes in each table are explained in Notes to OTMA header tables (Communications and Connections).

Server Available and Client Bid command format

Table 40. HWSOMHDR DSECT - OTMA state data header. State data common section for Server Available and Client Bid command format

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHDRLEN	2	0		STATE DATA LENGTH Set to State Data length.	1
OMHDRORG		2	State data for 'Server Available' and 'Client Bid' commands.		
OMHDRONM	16	2		MEMBER NAME OF ORIGINATING SERVER. NEITHER TESTED NOR SET BY EXIT.	4
OMHDROMT	8	12		MEMBER TOKEN OF COMMAND ORIGINATOR. NEITHER TESTED NOR SET BY EXIT.	4
OMHDRDMT	8	1A		MEMBER TOKEN OF COMMAND DESTINATION. NEITHER TESTED NOR SET BY EXIT.	4
OMHDRUEN	8	22		UNRESOLVED DESTINATION EXIT NAME. NEITHER TESTED NOR SET BY EXIT.	4
OMHDRMBS	2	2A		XCF TRANSMISSION MAX BLOCKSIZE. NEITHER TESTED NOR SET BY EXIT.	4

Table 40. HWSOMHDR DSECT - OTMA state data header. State data common section for Server Available and Client Bid command format (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHRRQE	1	2C	OMHRCMQ X'80'	CREATE HOLD MESSAGE QUEUE. NEITHER TESTED OR SET BY EXIT.	4
OMHDRFL2	1	2D	X'80'	OMHDRMXI - Specifies that OTMA limit the number of active messages that IMS can process concurrently. The maximum number is specified in the OMHDRTIB field.	3
			X'40'	Reserved	
			X'20'	OMHDRTMO - CM1 connections time out if OTMA does not receive an ACK message in the interval specified in the OMHDRTIO field.	
			X'10'	Reserved	
			X'08'	OMHDRSMI - This client is a member of the super member group specified in the OMHDRSMN field.	
			X'04'	OMHDR2SY - This client processes synchronous callout messages from IMS application programs.	
			X'02'	OMHDRTIOQ - CM0 ACK Timeout Queue Name	

Table 40. HWSOMHDR DSECT - OTMA state data header. State data common section for Server Available and Client Bid command format (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHDRUAV	4	2E		SAF USER ID TABLE AGING VALUE. NEITHER TESTED OR SET BY EXIT.	4
OMHDRHTS	4	32		MESSAGE RE-ASSEMBLY HASH TABLE SIZE. NEITHER TESTED OR SET BY EXIT.	4
OMHDRSMN	4	36			
I OMHDROSY	2	3A			
	2	3C			
OMHDRTIB	2	3E			
	1	40			
OMHDRTO	1	41			
I OMHDRTQN	8	42			
		2	ORG OMHDRORG		

Resume output for single named TPIPE for the asynchronous option of NO OPTION selection

Table 41. HWSOMHDR DSECT - OTMA state data header. State data common section for resume output for single named TPIPE for the asynchronous option of NO OPTION selection

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHDCRSM_COUNT	2	2		NUMBER OF TPIPES IN THE ARRAY. Set to number of tpipes to retrieve output from a RESUME TPIPE call. Only valid value is one (1).	1
OMHDCRSM_TPIPEN	n	4		TPIPE ARRAY. Set to the name of the tpipe to retrieve output from a RESUME TPIPE call. Only one name is valid.	1
		2	ORG OMHDRORG		

Resume output for single named TPIPE for options of NOAUTO, SINGLE, SINGLE with WAIT, and AUTO

Table 42. HWSOMHDR DSECT - OTMA state data header. State data common section for resume output for single named TPIPE for options of NOAUTO, SINGLE, SINGLE with WAIT, and AUTO

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHRRRHQ	1	2	Message return options for resume tpipe calls		
			OMHRRRHQ_NOAUTO - X'00'	Return all messages that are on the IMS queue when the call is received. Any messages received after the resume tpipe call is received are not sent. If no messages are on the queue, do not wait. Default option of HWSIMS00, HWSIMS01, HWSSMPL0, HWSSMPL1.	1
			OMHRRRHQ_AUTO1 - X'04'	Return one message only from the queue. If no message is on the queue, wait and send the next message to arrive. This option requires that IRM_TIMER be set to X'E9' on ACK to IMS Connect from client, to wait for next output from IMS Connect.	
			OMHRRRHQ_AUTO - X'02'	Return all messages that are currently on the IMS queue. After all current messages are sent, wait and send each new message as it arrives. This option requires that IRM_TIMER be set to X'E9' on ACK to IMS Connect from client, to wait for next output from IMS Connect.	1
			OMHRRQ_ONE - X'01'	Return one message only from the queue. If no messages are on the queue, cancel resume tpipe request. A new RESUME TPIPE Receive sequence is required to get any subsequent messages.	1

Table 42. HWSOMHDR DSECT - OTMA state data header. State data common section for resume output for single named TPIPE for options of NOAUTO, SINGLE, SINGLE with WAIT, and AUTO (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
 	1	3	Processing mode for synchronous callout messages		
			OMHRRHQ_SYNC - X'80'	This client processes only synchronous callout messages.	
			OMHRRHQ_SYAS - X'40'	This client processes both synchronous callout messages and asynchronous messages	
OMHDCRHQ_TPIPEN	8	4		Tpipe name. Set to the name of the tpipe to retrieve output from a RESUME TPIPE call. Only one name is valid.	1
		2	ORG OMHDRORG		

Transaction messages

Table 43. HWSOMHDR DSECT - OTMA state data header. State data common section for transaction messages

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
 	1	2	OMHDCRCNV X'80'	IMS STATE FLAG Conversational State.	2
			OMHRRSP X'40'	Response Mode. Neither tested nor set by exit.	4
			OMHDMHQ X'20'	Message from Hold Queue.	
			OMHDRROT X'08'	Message rerouted	
			OMHDXP1 X'01'	Transaction expiration time provided in user data section at the offset specified in OMHDRSXP	
			OMHDRTOE X'01'	Indicates that a resume tpipe token is present in the user data section.	

Table 43. HWSOMHDR DSECT - OTMA state data header. State data common section for transaction messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHDRSYN	1	3	X'80'	SYNCHRONIZATION FLAG. Reserved.	3
			OMHDRCM0 X'40'	Commit Mode 0 Set if default for exit or the IRM requests Commit Mode 0, field IRM_F2 is set to "IRM_CMODE0."	1
			OMHDRCM1 X'20'	Commit Mode 1 Set if default for exit or the IRM requests Commit Mode 1, field IRM_F2 is set to "IRM_CMODE1."	1
			OMHDRNTX X'10'	Notify of transfer. NEITHER TESTED NOR SET BY EXIT.	4
			OMHDRSYC X'08'	This is a synchronous callout message.	
OMHDRSLV	1	4	OMHDRSL0 X'00'	SYNCH LEVEL Synchlevel=0 (None) Set if default for exit or the IRM requests synch level none, field IRM_F3 is not set to "IRM_CONFIRM." Synchlevel=0 is only valid for Commit Mode 1.	1
			OMHDRSL1 X'01'	Synchlevel=1 (Confirm) Set if default for exit or the IRM requests synch level confirm, field IRM_F3 is set to "IRM_CONFIRM." Synchlevel=1 is valid for Commit Modes 0 and 1.	1
			OMHDRSL2 X'02'	Synchlevel=2 (Syncpt) NEITHER TESTED NOR SET BY EXIT.	

Table 43. HWSOMHDR DSECT - OTMA state data header. State data common section for transaction messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHDCFL	1	5	OMHDRSOM X'80'	Set for SENDONLY. SENDONLY is indicated in the IRM from the client.	1
			OMHDRAGN X'40'	Specifies either: - An aging value is in effect for the input message. - This is a NAK response message to a synchronous callout message that instructs OTMA to retain the undelivered message on the tpipe hold queue until it is retrieved by another resume tpipe or the message times out.	
			OMHDRRRQ X'20'	Set for REROUTE REQUEST for undeliverable CM0 output.	
			OMHDRPND X'10'	Set for PURGE NOT DELIVERABLE. SENDONLY and PURGE NOT DELIVERABLE are mutually exclusive.	
			OMHDREWC X'04'	Obsolete. EWLM is no longer supported by IMS. If this flag is specified, it is ignored by OTMA.	
OMHDRMAP	8	6		MAP NAME If client application requested that the MODname be returned, then the exit must build an RRM in front of the data being returned to the client. The MODname (OMHDRMAP) would be moved to the RRM to field RRM_MODNAME by the exit.	1
OMHDR TOK	16	E		SERVER TOKEN. NEITHER TESTED NOR SET BY EXIT.	4
			ORG OMHDR TOK + 4		
OMHDRSXP	2	12		OFFSET IN USER DATA PREFIX TO THE TRAN EXPIRATION UTC	

Table 43. HWSOMHDR DSECT - OTMA state data header. State data common section for transaction messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHDRSXQ	2	14		Offset in the user data section at which OTMA placed the correlation token for this synchronous callout message. This offset is specified by IMS Connect and passed to OTMA in the OMHDROSY field in the client bid.	
OMHDRUID	8	16		USERID RETURNED FOR RT REQ	
OMHDMCOR	16	1E		CORRELATOR. NEITHER TESTED NOR SET BY EXIT.	4
		ORG OMHDMCOR + 8			
OMHDMTIM	8	26		Time stamp correlator ID. On send and receive messages, this field contains the time stamp that is used to correlate responses to original input messages.	
OMHDMTKN	8	2E		Resume tpipe token. On resume tpipe requests and associated responses, this field contains the resume tpipe token.	
		ORG OMHDMTKN			
OMHDMCID	16	2E		CONTEXT ID. NEITHER TESTED NOR SET BY EXIT.	4
OMHDMRLTM	8	3E		OVERRIDE LTERM NAME. Set from IRM LTERM field "IRM_LTERM."	1
OMHDMRLIU	2	46		LENGTH OF IMS HEADER USER DATA. Set to the user header data length that follows. The length of this field is not included in the total length.	5
OMHDMRIUD	n	48		IMS HEADER USER DATA. Variable length, set by the user.	5
		ORG OMHDMRTOK+4			

Table 43. HWSOMHDR DSECT - OTMA state data header. State data common section for transaction messages (continued)

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMHDRSXP	2	12	Offset at which the transaction expiration time is located in the user data section of the OTMA header	This field is set by IMS Connect	2

Resource state protocol command

The resource state protocol command is used by OTMA for two purposes: to notify the OTMA client about the current state of the IMS server resources or to send a heartbeat message to the client. IMS Connect currently ignores heartbeat messages.

The resource state protocol command is identified by X'3C' in the command type field (OMCTLTYP) of the message control information section of the OTMA header.

When IMS Connect receives a resource state protocol command, IMS Connect updates the exit interface block data store entry (HWSXIBDS with the information contained in the resource state protocol command message.

The following table summarizes the format of state data for the resource state protocol command. The summary includes byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 44. HWSOMHDR DSECT - OTMA state data header. State data common section for resource state protocol command format

Content	Length	Dec offset	Hex offset	Value
OMHDRLEN	2	0	X'00'	Length of the state-data section, including the length field itself.
OMHDRSIM_STATUS	2	2	X'02'	The state of OTMA resources in the IMS system server.
X'03' - Normal state				
IMS is available and processing messages normally. OTMA issues a normal state protocol command:				
- When an OTMA client establishes a new tpie connection - When an OTMA client reestablishes a tpie connection - As a heartbeat message at 60 second intervals				
X'02' - Degraded state				
IMS is processing messages slowly. OTMA issues a degraded state protocol command when one or more conditions indicate that IMS is not processing messages as quickly as it should.				
X'01' - Unavailable state				
IMS can no longer accept messages for processing. OTMA issues the unavailable state protocol command to alert the OTMA client that one or more severe conditions prevent IMS from processing OTMA messages.				
OMHDRSIM_SVRFLG1	1	4	X'04'	Reserved

Table 44. HWSOMHDR DSECT - OTMA state data header. State data common section for resource state protocol command format (continued)

Content	Length	Dec offset	Hex offset	Value
OMHDRSIM_SVRFLG2	1	5	X'05'	Reserved
OMHDRSIM_SVRFLG3	1	6	X'06'	Reserved
OMHDRSIM_SVRFLG4	1	7	X'07'	Flags for resource in the fourth group of unavailable resources. X'01' - OMHDRSIM_S4FLOOD The server is flooded with OTMA messages that are waiting to be processed and is no longer available.
OMHDRSIM_WRNFLG1	1	8	X'08'	Flags for resources in the first group of degraded resources. X'80' - OMHDRSIM_W1FLOOD Global flood warning for all OTMA clients X'40' - OMHDRSIM_W1MTP Global tpipes warning. The total number of tpipes that are being monitored by OTMA on the server is equal to or greater than the highest specified limit for any one OTMA client.
OMHDRSIM_WRNFLG2	1	9	X'09'	Reserved
OMHDRSIM_WRNFLG3	1	10	X'0A'	Reserved
OMHDRSIM_WRNFLG4	1	11	X'0B'	Flags for resources in the fourth group of degraded resources. X'08' - OMHDRSIM_W5MTP The number of tpipes for this OTMA client has reached the maximum allowable number of tpipes set for this client on the MAXTP parameter of the OTMA client descriptor. No new tpipes can be created for this OTMA client until the number of tpipes drops. X'04' - OMHDRSIM_W4MTP The number of tpipes for this OTMA client has reached 80% of the maximum allowable number of tpipes set for this client on the MAXTP parameter of the OTMA client descriptor. X'02' - OMHDRSIM_AWE Message AWE reaches 80% flood. X'01' - OMHDRSIM_W4FLOOD Flood warning for this client only. The number of OTMA messages waiting to be processed on the server is at eighty percent of the maximum allowable number defined for the server.
OMHDRSIM_NRSFLGS	1	12	X'0C'	Other flags for non-resource related indicators. X'80' - OMHDRSIM_HB60S Identifies this message as a heartbeat message. The server is available and resource usage is within normal limits. Heartbeat messages are sent every 60 seconds.
Reserved	3	13	X'0D'	
OMHDRSIM_SRVNAME	16	16	X'10'	The 16 character z/OS cross-system coupling facility (XCF) member name of the OTMA server.
OMHDRSIM_CLTNAME	16	32	X'20'	The 16 character XCF member name of the OTMA client.
Reserved	20	48	X'30'	
OMHDRSIM_UTC	12	68	X'44'	UTC time for this message

Related reference:

“State data section” on page 867

OTMA security data fields used by IMS Connect

The tables in this topic define the fields of the OTMA security data header and the order of those fields.

The numbered notes in each table are explained in Notes to OTMA header tables (Communications and Connections).

Common security data section for all messages

Table 45. HWS0MSEC DSECT - OTMA security data header. Security data common section for all messages

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMSECLN	2	0		SECURITY DATA LENGTH	
OMSECFLG	1	2	OMSECNON C'N'	SECURITY FLAG No RACF checking. Set to 'N' if no OTMA RACF calls are to be made.	1
			OMSECCHK C'C'	Check for transaction and command. NEITHER TESTED NOR SET BY EXIT.	4
			OMSECFUL C'F'	Check for transaction, command, and MPR Set to 'F' is OTMA is to issue RACF call.	1
OMSECFLN	1	3		LENGTH OF FOLLOWING FIELDS Set to length of USERID and GROUPID section. • Set to X'0A' if only USERID. • Set to X'14' if USERID and GROUPID. • Set to X'00' if neither USERID or GROUPID present.	1

USERID security data section for all messages

Table 46. HWSECUDS DSECT - OTMA USERID definition. Security data USERID section for all messages

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMSECULN	1	0		LENGTH OF USERID FIELDS Set to length of USERID fields. The length includes this field. Set to X'09' if USERID present.	1
OMSECUTY	1	1	OMSECUXX X'02'	FIELD TYPE USERID type. Set to X'02' to identify USERID present.	1
OMSECUID	8	2		USERID Set to USERID from IRM field IRM_RACF_USERID.	1

GROUPID security data section for all messages

Table 47. HWSECGDS DSECT - OTMA GROUPID definition. Security data GROUPID section for all messages

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMSECGLN	1	0		LENGTH OF GROUPID FIELDS Set to length of GROUPID fields. The length includes this field. Set to X'09' if GROUPID present.	1
OMSECGTY	1	1	OMSECGXX X'02'	FIELD TYPE GROUPID type. Set to X'03' to identify GROUPID present.	1
OMSECGRP	8	2		RACF GROUPID Set to GROUPID from IRM field IRM_RACF_GROUPID or from default GROUPID from IMS Connect configuration file.	1

UTOKEN security data section for all messages

Table 48. HWSECFDS DSECT - OTMA RACF UTOKEN definition. Security data UTOKEN section for all messages

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMSECRLN	1	0		LENGTH OF UTOKEN FIELDS Set to length of UTOKEN fields. The length includes this field. Set to X'51' if user security exit issued RACF call.	1
OMSECRTY	1	1	OMSECRXX X'02'	FIELD TYPE UTOKEN type. Set to X'00' to identify UTOKEN present.	1
OMSECPRF	80	2		UTOKEN Set to UTOKEN from user security exit.	1

Related reference:

"Security data section" on page 880

OTMA user data fields used by IMS Connect

The format of the user data fields in the OTMA header is defined by the HWSOMUSR DSECT in the HWSOMPFX macro and is common to all IMS Connect messages.

The numbered notes in the table are explained in Notes to OTMA header tables (Communications and Connections).

Table 49. HWSOMUSR DSECT - user data header. User data common section for all messages.

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMUSRLN	2	0		USER DATA LENGTH. Set to length of User Data Header.	1
	2	2		ALIGN FULLWORD.	
OMUSR_DESTID	8	4		DESTINATION ID. Set to destination ID (data store) from IRM field IRM_IMSDESTID.	1
OMUSR_ORIGID	8	C		ORIGIN ID. Neither tested nor set by exit.	4
OMUSR_PORTID	8	14		PORTID. Neither tested nor set by exit.	4

Table 49. HWSOMUSR DSECT - user data header. User data common section for all messages (continued).

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMUSR_LTOKEN	8	1C		LOGON TOKEN. Neither tested nor set by exit.	4
OMUSR_RETCODE	4	24		Communication return code. Neither tested nor set by exit.	4
OMUSR_RESCODE	8	28		Communication reason code. Neither tested nor set by exit.	4
OMUSR_RTOKEN	4	30		RESPONSE TOKEN - A(SVT) COMMUNICATION RETURN CODE. Neither tested nor set by exit.	4
OMUSR_PASSTICK	8	34		RACF PASSWORD/ PASSTICKET. Set to password from IRM field IRM_RACF_PW. This field must be cleared before passing message back to IMS Connect.	1

Table 49. HWSOMUSR DSECT - user data header. User data common section for all messages (continued).

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMUSR_FLAG1	1	3C	OMUSR_CONV_OPT EQU X'80'	For IMS TM Resource Adapter, specifies IMS conversational transaction support for service oriented architecture composite business applications (process choreography).	4
			OMUSR_NPSOCKET X'40'	Non-persistent Socket. Neither tested nor set by exit.	4
			OMUSR_CANCID X'20'	Cancel client ID When establishing a new session, cancel client ID requests that if an existing session on the same port is using the same client ID as the new session, the existing session be terminated to allow the new session to connect.	
			OMUSR_PSOCKET X'10'	Persistent Socket. Set if persistent socket specified in IRM field IRM_SOCT has been set to IRM_SOCT_PER.	4
			OMUSR_TRAN X'00'	FLAG 1. Transaction Socket. Set if transaction socket specified in IRM field IRM_SOCT has been set to IRM_SOCT_TRAN.	

Table 49. HWSOMUSR DSECT - user data header. User data common section for all messages (continued).

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMUSR_FLAG2	1	3D	OMUSR_PWDTEXT X'01'	FLAG 2. PASSTCKT field is text password. Neither tested nor set by exit.	4
			OMUSR_PWDBIN X'02'	PASSTCKT field is binary password. Neither tested nor set by exit.	4
			OMUSR_HWSPLSET X'10'	For IMS TM Resource Adapter, set if the IMS Connect protocol level is specified in the OMUSR_PROLEV field.	
			OMUSR_TRSTUSR X'80'	Trusted user can be set by exit.	
			OMUSR_F2_CIDREQ X'40'	For IMS TM Resource Adapter clients, specifies that if the ID of the current connection is a duplicate of an ID of an existing connection, IMS Connect generate a unique client ID for this connection.	
			OMUSR_F2_CIDGEN X'20'	Set when a reply message sent to IMS TM Resource Adapter by IMS Connect contains a generated client ID.	

Table 49. HWSOMUSR DSECT - user data header. User data common section for all messages (continued).

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMUSR_FLAG3	1	3E	OMUSR_HDRCM0 X'40'	ORIGINAL SYNCHRONIZATION. Commit Mode 0. Neither tested nor set by exit.	4
			OMUSR_HDRCM1 X'20'	Commit Mode 1 (Send Commit). Neither tested nor set by exit.	4
			OMUSR_SOA X'08'	Send only protocol with the acknowledgment option.	
			OMUSR_SOO X'04'	Send only protocol with the serial (ordered) delivery option.	
			OMUSR_ALT_ANAK X'02'	ACK or NAK requests output to be rerouted to an alternate client ID	
			OMUSR_OM_MSG X'01'	OM message present	
OMUSR_TIMER	1	3F	OMUSR_ZERO X'E9' - X'nn' specifies no wait. OMUSR_FF specifies wait forever. See "Timeout specifications on input messages" on page 318 for a range of timer values.	OMUSR_TIMER specifies the amount of time that IMS Connect will wait for a reply from OTMA. Value is supplied by the IRM_TIMER byte, the TIMEOUT= keyword of the IMS Connect TCPIP configuration statement, or the internal IMS Connect logic.	1
OMUSR_USTAT	4	40		USTAT ADDRESS. Neither tested nor set by exit.	4
OMUSR_APPL_NM	8	44		PassTicket APPLname set to blanks or IRM value.	1
OMUSR_RRS_RCD	4	4C		z/OS Resource Recovery Services return code	

Table 49. HWSOMUSR DSECT - user data header. User data common section for all messages (continued).

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMUSR_ARCLEV	1	50	Architectural level of the message structure of IMS TM Resource Adapter messages.		
			OMUSR_AL00 - X'00': Indicates the base architecture level IMS TM Resource Adapter		
			OMUSR_AL02 - X'02': Indicates architecture level 2, which is required to support: - Reroute of undeliverable CM0 output - RESUME TPIPE call with an alternate client ID		
OMUSR_PROLEV	1	51	When OMUSR_HWSPLSET is set, this byte specifies the protocol level of this IMS Connect instance.		
			OMUSR_PR00 - X'00'	Base protocol level	
			OMUSR_PR01 - X'01'	Currently not used	
			OMUSR_PR02 - X'02'	CM0 NOWAIT ACK support	
OMUSR_RES4	2	52		Reserved for IMS Connect usage.	3
				Neither tested nor set by exit.	
OMUSR_RES5	4	54		Reserved for IMS Connect usage.	3
				Neither tested nor set by exit.	
OMUSR_RES6	4	58		Reserved for IMS Connect usage.	3
				Neither tested nor set by exit.	
OMUSR_REROUT_NM	8	5C		Client Reroute name. Set to reroute name specified in IRM_REROUT_NM. This field occupies the same offset as the OMUSR_RT_ALTCID field.	1

ORG

OMUSR_REROUT_NM

Table 49. HWSOMUSR DSECT - user data header. User data common section for all messages (continued).

Field	Length	Hexadecimal offset	Field value	Description and settings	Note
OMUSR_RT_ALTCID	8	5C		Alternate client ID. Tested by exit. Also requires OMUSR_ARCLEV to be set to OMUSR_AL02. This field occupies the same offset as the OMUSR_REROUT_NM field.	6
OMUSR_ADPTNM	8	64		Adapter name.	
OMUSR_DRVNM	8	6C		Input driver name.	
OMUSR_RES7	48	74		Reserved for IMS Connect usage.	
	2	A4		Reserved. Align to fullword.	
OMUSR_UTC_TO	16	A6		Transaction expiration time in UTC format	
OMUSR_CORTKN	0XL40			Correlation token for synchronous callout messages	
OMUSR_CT_LEN	2	B6		Length of the correlation token	
OMUSR_CT_RESV1	2	B8		Reserved	
OMUSR_CT_IMSID	4	BA		IMS identifier	
OMUSR_CT_MEMTK	8	BE		OTMA TMEMBER token	
OMUSR_CT_AWETK	8	C6		OTMA message token	
OMUSR_CT_TPIPE	8	CE		OTMA TPIPE name	
OMUSR_CT_USERID	8	D6		User ID from ICAL call	
OMUSR_RES8	8	DE		Reserved for IMS Connect usage.	
Reserved	2	E6		Align to fullword.	

Related reference:

“User data section” on page 882

Notes to OTMA header tables

The following notes correspond to the numbers shown in the tables that describe the fields of the OTMA header.

- Note 1: Set by READ routine of user-written exit.
- Note 2: Set by IMS Connect.
- Note 3: Reserved fields.
- Note 4: Set by IMS Connect and not analyzed by user exit.
- Note 5: User-defined area.
- Note 6: Analyzed by user exit.

Chapter 19. IMS Connect protocols

IMS Connect provides several different transaction protocols.

The protocols include:

- Conversational support
- Send only
- RESUME TPIPE/Receive for asynchronous output
- Socket connections
- Asynchronous output support

Transaction restrictions and limitations

The restrictions and limitations for transactions can differ depending on the specific transaction type.

The following is a list of restrictions and limitations of specific transactions:

- IMS Fast Path, conversational, and nonrecoverable transactions must be issued using commit mode 1. This is a restriction of IMS OTMA.
- Non-response transactions can be sent to IMS Connect using the SENDONLY option and must be issued using commit mode 0 on a transaction or persistent socket.

Commit mode and synchronization level definitions

IMS Connect supports the following commit modes and synchronization levels.

Commit mode 0

Commit mode 0 (CM0) is also called commit-then-send. CM0 is supported on both persistent and transaction sockets and supports only synch level CONFIRM.

Commit mode 1

Commit mode 1 (CM1) is also called send-then-commit. CM1 is supported on both persistent and transaction sockets and supports synch levels NONE, CONFIRM, and SYNCH.

For CM1 input messages, always specify a timeout value for IMS to wait for the ACK and NAK responses from IMS Connect. Until an ACK or NAK response is received by IMS, the IMS dependent regions remains occupied and continues to hold the necessary database locks. Valid timeout intervals are from 0 to 255 seconds and are specified on the ACKTO parameter of the DATASTORE configuration statement.

You can view the current timeout values by issuing the IMS Connect commands VIEWHWS or VIEWDS, or you can issue the MVS command QUERY MEMBER.

Synch Level=NONE

The synchronization level specifies the level of acknowledgment for each transaction. If a transaction is specified with Synch Level=NONE, no acknowledgment is required from the client. The database changes are still committed if the output message is sent to IMS Connect, but not to the

client. However, if OTMA is unable to deliver the output message to IMS Connect, the input and output message are discarded, the database changes are backed out, and the IMS application terminates and returns with a 119 ABEND.

Synch Level=CONFIRM

If a transaction is specified with Synch Level=CONFIRM, the client is required to send an acknowledgment to signal to IMS Connect whether the output message was successfully (ACK) or unsuccessfully (NAK) processed by the client.

The processing of CONFIRM is dependent on the type of commit mode that you specify:

- If Synch Level=CONFIRM is requested with CM0, and the client responds with ACK, IMS dequeues the output message. If the client returns a NAK response, the output message is requeued in IMS for later retrieval by the client.
- If Synch Level=CONFIRM is requested with commit mode 1, and the client responds with ACK, the database changes are committed. If the client responds with NAK, the database changes are backed out and the output message is discarded by IMS.

Synch Level=SYNCH

If a transaction is specified with Synch Level=SYNCH, two phase commit processing is required. Use Synch Level=SYNCH when multiple participants are involved in sync point processing. Synch Level=SYNCH is managed through z/OS Resource Recovery Services.

Related concepts:

“Socket connections” on page 308

IMS Connect conversational support

A conversational program is a message processing program (MPP) that processes transactions made up of several steps. The MPP does not process the entire transaction at once.

The conversational support for IMS Connect includes having conversational transactions that let you retain uninterrupted connection (continuity) for messages coming from a given client. Typically, a conversation is terminated when the message is sent and dequeued and the application program has placed blanks in the SPA, or the conversation is terminated when a COMMIT CONFIRMED message is received from the client. For conversational support for IMS Connect, conversations require a send-then-commit mode and are nonrecoverable.

Requirement: IMS support for SOA composite business applications requires that all iterations of a given IMS conversation be processed by the same IMS Connect and same IMS. More specifically, all iterations of a given IMS conversation must use the same conversational ID, port number, IMS Connect, and data store. This can be accomplished by using the same connection factory for every iteration of a conversation. Using the same connection factory for each iteration ensures that the same host name, port number, and data store will be used. The IMS TM Resource Adapter client must also save and reuse the same IMSInteractionSpec convID property to ensure that the same Conversational ID is used for each iteration of that conversation.

In addition, different physical socket connections, represented by `IMSTCPIPManagedConnection` instances, can be used for each iteration of an IMS conversation, as long as these socket connections meet the above criteria. Having originated from the same connection factory, these connections would be members of the same managed connection pool.

These requirements preclude the use of Sysplex Distributor or the `SHAREPORT` keyword on the `PORT` statement of the z/OS TCP/IP profile when you use the IMS support for SOA composite business applications.

OTMA conversational protocols

The following examples show the OTMA conversational protocols.

Send-then-commit, sync level=none

The send-then-commit flow sends IMS output before IMS completes synchronization-point (hereafter referred to as sync-point) processing.

To use the send-then-commit flow, specify commit Mode 1 in the state-data section of the message prefix.




Figure 24. Send-then-commit, sync level=none flow for OTMA conversational protocols

The sample flow shown in the above figure assumes the following:

- The transaction pipe is not synchronized.
- The synchronization level is specified as NONE in the state-data section of the message prefix. Therefore, IMS does not request a response (an ACK) when sending output.

Send-then-commit, sync level=confirm

The send-then-commit flow assumes no synchronization for the transactions as they are processed by IMS.

The following figure shows a flow in which all transactions are confirmed as they are received (each message requests a response).




Figure 25. Send-then-commit, sync level=confirm flow for OTMA conversational protocols

The sample flow shown in the above figure assumes the following:

- Commit mode 1 is specified in the state-data section of the message prefix.
- The transaction pipe is not synchronized.
- The synchronization level is specified as Confirm in the state-data section.

IMS Connect conversational protocols

IMS Connect conversation protocols use the send-then-commit protocol. The following examples show variations of the conversational protocols with sync levels of none and confirm.

Send-then-commit, sync level=none, transaction terminated from the program

The send-then-commit flow sends IMS output before IMS completes sync-point processing.

To use the send-then-commit flow, specify commit mode 1 in the state-data section of the message prefix.




Figure 26. Send-then-commit, sync level=none (transaction terminated from program) flow

The sample flow shown in the above figure assumes the following:

- The transaction pipe is not synchronized.
- The synchronization level is specified as NONE in the state-data section of the message prefix. Therefore, IMS does not request a response (an ACK) when sending output.
- The transaction is terminated from the program.
- IMS Connect will close the socket as soon as Commit confirmed has been sent by IMS.

Send-then-commit, sync level=none, transaction terminated from the client

The send-then-commit flow sends IMS output before IMS completes sync-point processing.

To use the send-then-commit flow, specify commit mode 1 in the state-data section of the message prefix.




Figure 27. Send-then-commit, sync level=none (transaction terminated from client) flow

The sample flow shown in the above figure assumes the following:

- The transaction pipe is not synchronized.
- The synchronization level is specified as NONE in the state-data section of the message prefix. Therefore, IMS does not request a response (an ACK) when sending output.
- The transaction is terminated from client.

Send-then-commit, sync level=confirm, ACK response

This send-then-commit flow assumes no synchronization for the transactions as they are processed by IMS.

The following figure shows a flow in which all transactions are confirmed as they are received (each message requests a response).




Figure 28. Send-then-commit, sync level=confirm (ACK response) flow

The sample flow shown in the preceding figure assumes the following:

- Commit mode 1 is specified in the state-data section of the message prefix.
- The transaction pipe is not synchronized.
- The synchronization level is specified as Confirm in the state-data section.
- ACK can be replied to by a remote workstation before the check response requested bit.

Send-then-commit, sync level=confirm, NAK response

This send-then-commit flow assumes no synchronization for the transactions as they are processed by IMS.

The following figure shows a flow in which all transactions are confirmed as they are received (each message requests a response).




Figure 29. Send-then-commit, sync level=confirm (NAK response) flow

The sample flow shown in the preceding figure assumes the following:

- Commit mode 1 is specified in the state-data section of the message prefix.
- The transaction pipe is not synchronized.
- The synchronization level is specified as Confirm in the state-data section.
- NAK can be replied to by either IMS Connect or a remote workstation before the check requested bit.
- If the client forgets to send the NAK/ACK before it closes the socket, IMS Connect will send the NAK to IMS and it will cause a U0119 abend.

Purging undeliverable commit-then-send output

You can configure OTMA to purge commit-then-send (commit mode 0) IOPCB output when the output cannot be returned to the OTMA client application that initiated the transaction.

When configured, if the OTMA client requests the purge function, OTMA dequeues and discards the undeliverable commit mode 0 (CM0) IOPCB output from the IMS output queue. The purge function is requested on a message-by-message basis.

When the purge function is not specified, IMS stores undeliverable commit-then-send IOPCB output on the asynchronous hold queue of the OTMA tpipe associated with the client application that submitted the original input message. The output message remains on the hold queue for later retrieval by using a RESUME TPIPE call.

You can specify the purge function on either CM0 or commit mode 1 (CM1) input messages. However, when specified on CM1 input messages, IMS purges only CM0 IOPCB output, such as might be generated by a program-to-program switch. For example, if you specify the purge function on a CM1 transaction input that does a program switch to a second transaction and the first transaction does an insert to the IOPCB, the purge function applies only to the subsequent transactions that insert to the IOPCB.

Both user-written applications and IMS TM Resource Adapter applications on either persistent sockets or transaction sockets can request the purge function.

Specifying the purge function for undeliverable commit-then-send output

If you are using either the HWSSMPL0 or the HWSSMPL1 IMS Connect user message exit, you can enable the purge function for undeliverable commit-then-send output by specifying the IRM_F3_PURGE flag (X'04') in the IRM_F3 field on certain types of input messages.

You can specify the purge function on the following input messages from a client application:

- A SEND of a commit-then-send (CM0) transaction.
- A SEND of a send-then-commit (CM1) transaction. A purge request on CM1 input applies only to any CM0 output that the CM1 input generates.
- A SEND of a NAK response to CM0 output.

Restriction: You cannot specify the purge function and the reroute function at the same time. If both functions are specified, the output messages are neither purged nor rerouted from the original output queue and OTMA issues message DFS2407W.

When IMS purges undeliverable commit-then-send output

If the purge function is specified, IMS purges commit-then-send (commit mode 0) output in certain circumstances only.

When the purge function is specified, IMS purges commit-then-send (commit mode 0) output when:

- IMS Connect receives a NAK response from the client application
- IMS Connect cannot deliver the output to the client application
- OTMA cannot deliver the output to IMS Connect

When IMS Connect receives a NAK response from the client or cannot deliver the output to the client, IMS Connect notifies OTMA to discard the output message from the IOPCB queue. When OTMA cannot deliver the output to IMS Connect, OTMA discards the output without waiting for notification from IMS Connect.

If an IMS Connect STOPCLNT command is issued for a client ID that specifies the purge function, the reply message is purged.

IMS does not support the purge function for the following types of output:

- IMS application output to ALTPCBs. Even if the purge function is specified for the IMS application CM0 output, the purge function does not apply to inserts to ALTPCBs.

- CM1 output. Any output from a CM1 transaction that is undeliverable is already discarded and the transaction is backed out.
- Output associated with a send-only transaction. Output for send-only transactions is routed directly to the asynchronous hold queue.
- Output associated with a RESUME TPIPE call input message. A RESUME TPIPE call, by definition, requires and guarantees that the output is delivered.

The purge function, multiple-message output, and NAKs

After receiving a NAK response from a client application for one of multiple related messages for which the purge function is specified, IMS Connect issues purge requests for the remaining output messages without attempting to deliver them to the client application and OTMA discards them from the IOPCB.

How IMS Connect issues the purge requests differs slightly depending on whether the multiple messages are generated by a single application program or by program-to-program switches.

If there are multiple commit then send output messages for the same application program on the IOPCB and the client issues a NAK response for one of the messages, IMS Connect sends the purge request to OTMA. IMS Connect also sends purge requests to OTMA for the remaining output messages on the IOPCB without attempting to deliver the output to the client application.

If there are multiple commit then send output messages generated by program-to-program switches on the IOPCB and the client issues a NAK response for one of the messages, IMS Connect passes a purge request to OTMA and then generates additional purge requests for any other related output messages currently on the IOPCB queue. If program-to-program switches generate related output messages after the initial NAK response was received, IMS Connect issues purge requests for them as well without passing the output to the client application.

Rerouting commit-then-send output

You can configure IMS to reroute commit-then-send (commit mode 0) IOPCB output to an alternate OTMA tpipe hold queue for retrieval.

Normally, if IMS cannot return commit mode 0 (CM0) output to the application client, the output is routed to the tpipe hold queue associated with the client application that submitted the original message; however, if you request the reroute function, IMS reroutes the output to either a user-specified tpipe hold queue or the default tpipe hold queue HWS\$DEF. Whether the reroute tpipe is a user-specified tpipe or the default tpipe, the reroute tpipe is always associated with the tmember of the original tpipe.

The reroute function can be used for managing output generated by send-only transactions and for managing output that cannot be delivered to the original client because the connection timed out or failed.

The reroute function is also useful when IMS TM Resource Adapter (formerly known as IMS Connector for Java) is used with shareable persistent sockets. The IMS TM Resource Adapter automatically generates the client ID when connecting to IMS Connect. Consequently, the client ID is unknown to the client applications, which need the client ID to retrieve the CM0 output.

You can specify the reroute function in either CM0 or CM1 input messages. However, in the case of CM1, IMS can only reroute the CM0 output, such as might be generated by a program to program switch.

Both user-written applications and IMS TM Resource Adapter applications on either persistent sockets or transaction sockets can request the reroute function.

Restrictions: The reroute function is not supported for:

- CM1 output messages
- Output resulting from a RESUME TPIPE call
- Output resulting from an insert to an ALTPCB

Specifying the reroute function for commit-then-send output

You enable the reroute function for commit-then-send output by setting a flag in the IRM header of your input message or by coding your user-written user message exit to set the appropriate flag in the OTMA state data.

If you are using either the HWSSMPL0 or the HWSSMPL1 IMS Connect user message exit, you can enable the reroute function for commit-then-send output by specifying the IRM_F3_REROUT flag (X'08') in the IRM_F3 field for the following input messages from a client application:

- A SEND of a commit-then-send (CM0) transaction
- A SEND of a CM0 send-only transaction from a user-written client application
- A SEND of a send-then-commit (CM1) transaction (A reroute request on CM1 input applies only to any CM0 output that the CM1 input generates)
- A SEND of a NAK response to CM0 output

Restriction: You cannot specify the purge function and the reroute function at the same time. If both functions are specified, the output messages are neither purged nor rerouted from the original output queue and OTMA issues message DFS2407W.

Specifying a destination for rerouted output

You can define the reroute destination by specifying a reroute request name. Output can only be rerouted to tpipes associated with the same tmember. Specifying a destination for rerouted output is optional.

If a client application requests that output be rerouted, but does not identify a reroute destination by specifying a tpipe name, the default reroute destination is tpipe HWS\$DEF.

You can define the reroute destination by specifying a reroute request name in one or more of the following places:

- The RRNAME= keyword in the IMS Connect DATASTORE configuration file
- An IMS Connect user message exit
- The IRM_REROUT_NM in the fixed IRM format of an input message associated with a SEND/RECEIVE request from a client application
- The IRM_REROUT_NM in the fixed IRM format of a NAK message from a client application

You can specify a different reroute destination in a NAK response message than is specified by an initial input message; however, doing so can cause problems for

transactions that generate multiple output messages. If a different reroute destination is specified in a NAK response message and multiple output messages are generated by the initial input message, OTMA reroutes only the message that triggered the NAK response to the destination specified in the NAK response message. After receiving the NAK response, OTMA automatically reroutes any subsequent output messages for the same transaction to the destination that was specified on the initial input message.

For more information about:

- The RRNAME= keyword of the IMS Connect DATASTORE configuration file, see *IMS Version 11 System Definition*.
- Coding IMS Connect user message exits, see *IMS Version 11 Exit Routines*.

Related reference:

Chapter 17, “IMS Connect message structures,” on page 237

 DATASTORE statement (System Definition)

When IMS reroutes commit-then-send output

If the reroute function is specified, IMS reroutes commit-then-send (CM0) output only in certain circumstances.

When the reroute function is specified, IMS reroutes commit-then-send (CM0) output if:

- IMS Connect cannot deliver the output to the client application.
- IMS Connect receives a NAK response from the client application.
- OTMA cannot deliver the output to IMS Connect.
- IMS inserts output for send-only transactions to the IOPCB.

When IMS Connect cannot deliver the output or when IMS Connect receives a NAK response, IMS Connect notifies OTMA to reroute the output message to the alternate destination.

In the case of a NAK response from the client application, OTMA reroutes the output if the NAK is in response to an IMS application insert to the IOPCB. If the NAK is in response to the output related to a RESUME TPIPE call, IMS does not reroute the output.

In the event of a communication failure between OTMA and IMS Connect, OTMA reroutes the commit-then-send output only if the original input message requested the reroute function.

For send-only transactions, when the reroute function is specified OTMA always reroutes the output.

If IMS Connect receives a disconnect notification on the TCP/IP READ for an ACK or NAK response to an output message, IMS Connect requests that OTMA reroute the CM0 output only if the input message specified the reroute function. If the client disconnects or times out prior to IMS Connect receiving the output message from OTMA, the CM0 output message is rerouted only if the input message specified the reroute function.

The reroute function, multiple-message output, and NAK responses

If a transaction produces multiple output messages, OTMA reroutes the first output message that triggers a NAK response and any subsequent output messages for the same transaction that are on the IOPCB at the time the NAK response message is received.

OTMA reroutes subsequent output messages for the same transaction that arrive to the IOPCB after the initial reroute only after triggering another NAK response.

If a transaction initiates a program-to-program switch and IMS Connect receives a NAK response to the first output message from an application program, OTMA reroutes output messages sent by a secondary application program after a program-to-program switch only if they are already on the output queue when the initial NAK response message is received or if the original input message specifies a reroute destination. If the original input message does not specify a reroute destination, OTMA does not reroute undeliverable output generated after a program-to-program switch.

If a send-then-commit (CM1) transaction message does a program-to-program switch to a second CM0 transaction message and the first transaction does an insert to the IOPCB, IMS reroutes only the second or subsequent CM0 messages that insert to the IOPCB. IMS reroutes the CM0 output of a CM1 input message only if the client application requests reroute in the original CM1 transaction message.

Recoverable IMS transactions

This topic contains scenarios for running recoverable transactions in the IMS Connect environment.

For each of the following scenarios:

- OTMA will have deleted the input message.
- Requeuing of the input message will not occur.
- For CM1 (send-then-commit), none of the output is placed (ENQUEUED) in the IMS queue.

Only commit mode 0 (CM0, also referred to as commit-then-send) is treated as recoverable; CM1 is not recoverable. With the use of CM0, IMS Connect creates a separate TPIPE for each client that uses CM0. This TPIPE remains in IMS, so a fixed client name is highly recommended for each client that intends to use CM0.

The combination of commit mode and Sync level is critical. The following scenarios describe the different uses and the results.

- With CM1 and SYNC LEVEL = NONE:
The input message is processed by IMS and an output message is sent back to IMS Connect, IMS Connect sends the message to the client, and any ACK/NAK from the client in response to the output message would become an error because the ACK/NAK are not expected and IMS Connect would have received a message from the client with no application data.
- With CM1 and SYNC LEVEL = CONFIRM:

The input message is processed by IMS and an output message is sent back to IMS Connect, IMS Connect sends it to the client, and an ACK from the client will result in the successful completion of the application. This scenario works as expected.

The input message is processed by IMS and an output message is sent back to IMS Connect, IMS Connect sends it back to the client, and a NAK from the client will result in an IMS MPP 119 abend and an IMS message, DFS555. The 119 abend will back out the database changes, and both the input and output messages are discarded. The result would be as if the system had never seen the transaction, and a reentry of the transaction would be necessary.

- With CM0 and SYNC LEVEL = CONFIRM:

The input message is processed by IMS, the application program commits the changes, and an output message is sent back to IMS Connect. IMS Connect sends it to the client, and an ACK response from the client results in IMS dequeuing the output, and represents the successful completion of transaction processing. CM0 forces the Synch level to Confirm. This scenario works as expected.

The input message is processed by IMS and an output message is sent back to IMS Connect, IMS Connect sends it back to the client, and a NAK from the client will result in the database changes not being backed out. The input message is discarded and the output message is requeued to the IMS queue for representation. These output messages will be moved to the hold asynchronous queue by OTMA, and will be retrievable only with the RESUME TPIPE, RECEIVE and ACK process.

Recommendation: To run recoverable transactions in the IMS Connect environment, use CM0 and SYNC LEVEL = CONFIRM, and use a single unique CLIENT_ID for each client that uses CM0 and SYNC LEVEL = CONFIRM

Send-only protocol

Client application programs use the send-only protocol to submit commit then send (CM0) input messages to IMS in rapid succession without requiring the client application to wait for a response. The send-only protocol is designed for fast, high volume input.

Clients configured to support synchronous callout requests also use the send-only protocol to return responses from an external data or service provider to an IMS application program that issued the synchronous callout request and that is waiting for the response in an IMS dependent region.

The output generated by IMS in response to send only input is stored on an asynchronous hold queue associated with tpipe used by the client application and can be retrieved later by issuing a RESUME TPIPE call.

If send-only input messages must be processed by IMS in the order in which they are sent by the client, the send-only protocol offers two options that can help ensure that IMS receives the messages in the order in which IMS Connect receives them from the client application:

- Send-only protocol with the acknowledgment option
- Send-only protocol with the serial delivery option

When the send-only protocol is used to return responses from a data or service provider external to the IMS installation, the send-only message does not contain a

transaction code, but rather the data requested by the IMS application or, in the event there was a problem processing the callout request, error codes.

The following figure shows an example of the CM0 send-only protocol flow when the commit confirmed flag is on. The CM0 flow, also known as the IMS standard flow, enqueues IMS output before sending it to the client. However, in this case for non-response transactions, the client does not expect any output from IMS.




Figure 30. Send-only protocol flow

The preceding sample flow assumes the following:

- CM0 is specified in the state-data section of the message prefix.
- The transaction bit and the commit confirmed bit is specified in the control-data section of the message prefix.

Send-only with acknowledgment protocol

Use the send-only protocol with acknowledgment (SNDONLYA) to ensure that send only transaction inputs are enqueued by IMS in the same order that they are submitted by the client application program.

SNDONLYA messages execute a send-only interaction for a non-response mode, non-conversational transaction. If the host application terminates without issuing an ISRT to the IO PCB, no DFS2082 messages are returned to the client. SNDONLYA also specifies that IMS Connect must indicate in its reply that an ACK or NAK of the output is required from the client. The SNDONLYA interaction must use CM0.

When SNDONLYA is specified, the client application receives an ACK response message from OTMA for each input message successfully enqueued by IMS. All other output generated by the send only transaction is sent to the asynchronous hold queue.

Before sending the next send only input message, the client application must wait for and process the ACK response. Because subsequent input messages sent by the client application to the same tpipe are not sent until the preceding input message has been enqueued, transaction messages are enqueued in IMS in the order in which they were sent.

You can select the send-only protocol with the acknowledgment option by specifying K in the IRM_F4 field of the HWSSMPL0 and HWSSMPL1 user message exit IRM format.

The send-only protocol with the acknowledgment option and the send-only protocol with the serial delivery option are mutually exclusive. If both are specified, the send only with acknowledgment protocol takes effect.

The send-only protocol with the acknowledgment option is not supported by HWSJAVA0.

Send-only protocol with serial delivery protocol

When the send-only protocol with the serial delivery option is specified, IMS Connect ensures that the order in which it submits send-only transactions to OTMA is in fact the order in which IMS receives the transactions.

You can select the send-only protocol with the serial delivery option by specifying X'10' for IRM_F3_ORDER in the IRM_F3 field of the HWSSMPL0 and HWSSMPL1 user message exit IRM format.

The send-only protocol with the serial delivery option and the send-only protocol with the acknowledgment option are mutually exclusive. If both are specified, the send only with acknowledgment protocol takes effect.

The send-only protocol with the serial delivery option is not supported by HWSJAVA0.

The following requirements apply to the send-only protocol when the serial delivery option is used.

- All transactions that use the send-only protocol with the serial delivery option must be defined as SCHDTYP=SERIAL. If the transactions are defined as SCHDTYP=PARALLEL, OTMA cannot guarantee that IMS will process them serially.
- Send-only transactions that must be delivered serially relative to one another must use the same tmember and tpipe connection, including transactions of the same transaction type that must be delivered serially with respect to all other transactions of the same type. If a transaction is sent on multiple tmember and tpipe connections, OTMA cannot guarantee that IMS will process them serially.

Send-only protocol for synchronous callout responses

IMS Connect clients return responses to synchronous callout requests from IMS application programs by using the send-only protocol.

The IMS Connect client specifies the send-only protocol for synchronous callout responses in the IRM_F4 field of the IRM prefix of the response message.

The send-only message contains no transaction code and can contain either the response data or error information for the IMS application program. Send-only messages used for callout responses do not generate any type of output.

For user-written IMS Connect clients, code the callout response by specifying the following field values in the IRM of the response message:

- IRM_ARCH = X'03' (IRM_ARCH3)
- IRM_F0 = X'10' (IRM_F0_NAKRSN)

- IRM_NAK_RSNCDE = 2 byte hexadecimal extended error code
- IRM_F4 = M (IRM_F4_SYNRESP)
- IRM_CORTKN = 40 byte correlation token (CORTKN) from original callout request

Socket connections

IMS Connect provides three kinds of client TCP/IP connection protocols, which are called *sockets*. The TCP/IP sockets define how IMS Connect manages client TCP/IP connections when IMS Connect sends a disconnect message.

The three socket types provided by IMS Connect are:

- Persistent
- Transaction
- Non-persistent

Persistent sockets

A *persistent* socket is a connection between the client and IMS Connect that remains connected until either the client or IMS Connect specifically make a disconnect request. A persistent socket can exist across multiple transactions.

There are two ways that the client can force a termination:

- By sending IMS Connect a disconnect request.
- By changing the socket type to "transaction" for the last transaction entered, such as a logoff transaction.

IMS Connect can also terminate the connection when an error occurs.

The IMS Connect user message exits HWSSMPL0, HWSSMPL1, HWSSOAP1, and HWSJAVA0 support the use of persistent sockets.

A persistent socket supports both commit mode 1 (CM1 or send then commit) and commit mode 0 (CM0 or commit then send) processing.

In the following cases, IMS Connect generates a unique client ID for a connection from IMS TM Resource Adapter on a persistent socket:

- The IMS TM Resource Adapter requests that IMS Connect check for duplicate client IDs and the client ID provided on a new connection request is a duplicate of the client ID of an existing connection.
- The IMS TM Resource Adapter passes a blank client ID to IMS Connect.

Related concepts:

"Socket processing for transactions" on page 310

Related tasks:

"Setting socket types for IMS TM clients" on page 309

Related reference:

 User message exit routines HWSSMPL0 and HWSSMPL1 (Exit Routines)

 IMS TM Resource Adapter user message exit routine (HWSJAVA0) (Exit Routines)

 SOAP Gateway exit routine (HWSSOAP1) (Exit Routines)

Transaction sockets

A *transaction* socket is a connection between the client and IMS Connect that remains connected for a single transaction or IMS conversation. The connection can be terminated only by IMS Connect, either when IMS itself terminates, or when an error occurs.

A transaction socket supports both commit-mode-1 (CM1 or send then commit) and commit-mode-0 (CM0 or commit then send) processing.

Related concepts:

“Socket processing for transactions” on page 310

Related tasks:

“Setting socket types for IMS TM clients”

Non-persistent sockets

A *non-persistent* socket maintains a connection for a single input-and-output pair to IMS Connect.

IMS Connect terminates the connection after sending the output to the client for non-conversational and conversational transactions. If three exchanges of input and output occur, the disconnect is issued three times, one for each output from IMS Connect.

Restriction: The HWSSMPL0, HWSSMPL1, HWSSOAP1, and HWSJAVA0 user message exit routines do not support non-persistent sockets, nor does IMS TM Resource Adapter.

Related tasks:

“Setting socket types for IMS TM clients”

Setting socket types for IMS TM clients

For IMS Connect clients that access IMS TM, client code controls the socket settings, and the IMS Connect user message exits and the user initialization exit enforce the socket settings.

The client selects the socket connection type by setting a flag in the IMS request message (IRM) header, in the field IRM_SOCT. The IRM_SOCT flag values are shown in the following table.

Table 50. IRM_SOCT flags

Flag	Definition	Socket type
IRM_SOCT_PER	X'10'	Persistent
IRM_TRAN	X'00'	Transaction
IRM_SOCT_NONPER	X'40'	Non-persistent

The IRM_SOCT flag must be set for each message that is sent to IMS Connect from an IMS Connect client that accesses IMS TM.

Recommendation: Set all messages that are associated with a single transaction to the same socket type. If you do not, unexpected results can occur, as described in the following examples:

- If the first message of a conversational transaction is set to persistent, and the last message is set to transaction, then the socket connection will be terminated following the last message.
- If one of the messages in the middle of the conversational transaction set the socket type to transaction, and the IMS transaction terminates for some reason, then IMS Connect will disconnect the socket. This is because "transaction" was the last known socket type.

The user message exits determine the socket type, then move the socket type information to the user data section of the OTMA message header that they return to IMS Connect. To transfer the socket type information to the OTMA message header, the user exits set the OMUSR_FLAG1 field, which is mapped by the HWSOMUSR DSECT of the HWSOMPFX macro, with one of the following flags as shown in the following table:

Table 51. OMUSR_FLAG1 flags

Flag	Definition	Socket type
OMUSR_PSOCKET	X'10'	Persistent
OMUSR_TRAN	X'00'	Transaction
OMUSR_NPSOCKET	X'40'	Non-persistent

Related reference:

“IRM structures for IMS Connect client messages” on page 237

“OTMA user data fields used by IMS Connect” on page 283

Socket processing for transactions

For a transaction on either a transaction socket or persistent socket, the client application must always issue a TCP/IP READ following all TCP/IP SENDs.

The exceptions are for a TCP/IP SEND of SENDONLY or a TCP/IP SEND of an ACK with IRM_TIMER set to NO_WAIT (X'E9' char Z), which is issued in response to a READ of a RESUME_TPIPE single request.

The following scenarios describe transactions on a transaction socket. For transactions on a persistent socket, the process is the same as transactions on a transaction socket. However, the client application and IMS Connect do not disconnect. Also, the client application will receive a return code of X'28' if there is a timeout. The return code states a disconnect is not required.

For a Commit mode 0, Synch Level Confirm, non-conversational transaction on a transaction socket, the following scenario occurs:

1. The client application issues a SEND to send the transaction data to IMS Connect.
2. IMS Connect returns the output to the client application.
3. The client application receives the output, sends an ACK, and must issue a READ to receive the next output or the timeout notification.
4. IMS Connect issues a timeout notification with the return code of either X'20' or X'24' for a transaction socket, or an X'28' for a persistent socket. IMS Connect will disconnect the socket for the X'20' and X'24' return codes, and will keep the connection for the X'28' return code.
5. The client application issues a disconnect for return codes X'20' and X'24'. The client can issue a disconnect for return code X'28' or send in the next input.

For a CM1, Synch Level Confirm, non-conversational transaction on a transaction socket, the following scenario occurs:

1. The client application issues a SEND of the transaction data to IMS Connect.
2. IMS Connect returns the output to the client application.
3. The client application receives the output, sends an ACK or NAK, and issues a READ.
4. After an ACK is sent, the client receives one of the following responses:
 - Deallocate commit if the IMS transaction completes successfully.
 - A DFS message if the IMS transaction failed.
 - A timeout notification with a return code of X'20' or X'24' for a transaction socket or a return code of X'28' for a persistent socket. The client application is required to issue a disconnect for return codes X'20' and X'24'.
5. The client application issues a disconnect.

For a CM1, Synch Level Confirm, conversational transaction on a transaction socket, the following scenario occurs:

1. The client application issues a SEND to send the transaction data to IMS Connect.
2. IMS Connect returns the output to the client application.
3. The client application receives the output, sends an ACK, and issues the next input. The client continues SEND, READ, ACK until the transaction is complete.
4. IMS Connect issues an RSM deallocate commit, deallocate abort, or a timeout notification. The timeout notification returns either X'20' or X'24', which indicates that IMS Connect will disconnect.
5. The client application issues a disconnect.

Managing the number of sockets

IMS Connect supports from 50 to 65,535 sockets at one time. Within that range, you can define both a maximum number of allowable sockets and the point at which IMS Connect issues warnings when the number of sockets approaches the maximum.

The total number of sockets that IMS Connect supports across all TCP/IP ports at the same time is set by the MAXSOC parameter in the TCPIP configuration statement of the IMS Connect configuration member. If the number of sockets reaches the maximum, IMS Connect refuses any new connections and issues message HWSS0771W. After the number of connections falls below the MAXSOC value, IMS Connect resumes accepting connections.

Because IMS Connect uses one socket on each TCP/IP port for listening, the maximum number of physical connections that IMS Connect supports is the MAXSOC value less the number of TCP/IP ports. For example, if you specify MAXSOC=80 and have five TCP/IP ports, 75 physical connections can be made.

By default, the MAXSOC parameter sets the maximum number of sockets at 50.

As the number of sockets approaches the maximum allowable number, IMS Connect issues warnings at a warning threshold set by the WARNSOC parameter and, if the number of sockets continues to rise, at increments set by the WARNINC

parameter thereafter. Both the WARNSOC parameter and the WARNINC parameter are specified on the TCPIP configuration statement in the IMS Connect HWSCFGxx PROCLIB member.

You can display the current values of the warning threshold and the warning increment, as well as the current number of open sockets for an instance of IMS Connect, by issuing the IMS Connect command VIEWHWS or the IMS Connect z/OS command QUERY MEMBER.

You can also display the current number of sockets on a given port by issuing either the IMS Connect command VIEWPORT or the IMS Connect z/OS command QUERY PORT.

Even when there are no active client connections on a port, when displaying the current number of sockets, IMS Connect ports always show one open socket: the port listen socket.

The IMS Connect SSL port shows at least two sockets: the port listen socket and a socket that represents the SSL-related file descriptor. When the IMS Connect SSL socket receives its first connection, an additional SSL-related file descriptor is created, so that the IMS Connect SSL socket shows four sockets for the first client connection: the port listen socket, two file descriptors, and the client connection socket.

Related concepts:

 [IMS Connect definition and tailoring \(System Definition\)](#)

Related reference:

 [HWSCFGxx member of the IMS PROCLIB data set \(System Definition\)](#)

Limits on the number of sockets set by z/OS UNIX System Services

The number of sockets supported by IMS Connect is affected by the z/OS UNIX System Services parameter MAXFILEPROC, which limits the number of sockets IMS Connect can open on each port.

For example, if the value of the UNIX System Services MAXFILEPROC parameter is 100, IMS Connect can open no more than 100 sockets on one port. In contrast, the IMS Connect MAXSOC parameter limits the total number of sockets that IMS Connect supports, regardless of the number of ports that are used.

Important: The value of the UNIX System Services parameter MAXFILEPROC must be equal to or greater than the value of the IMS Connect parameter MAXSOC. Otherwise, IMS Connect cannot open any ports.

You can ensure that the value of the MAXFILEPROC parameter is set appropriately by granting UNIX System Services superuser privileges to IMS Connect, which allows IMS Connect to change the value of the MAXFILEPROC parameter automatically. You can grant UNIX System Services superuser privileges to IMS Connect by using the RACF command ALTERUSER to assign an OMVS segment with a UID of 0 to the user ID of the IMS Connect-started task. Alternatively, your UNIX System Services administrator can adjust the value of the MAXFILEPROC parameter in the BPXPRMxx member of the z/OS SYS1.PARMLIB data set.

If IMS Connect does not have superuser privileges, and the MAXSOC parameter value is greater than the MAXFILEPROC parameter value, IMS Connect issues the message HWSP1415E TCP/IP SOCKET FUNCTION CALL FAILED; F=SETRLIMI, R=-1, E=139, M=SDOT and does not open any ports.

You can check the value of the MAXFILEPROC parameter for IMS Connect by issuing the UNIX command D OMVS,L,PID=, where PID is the process ID for IMS Connect. You can determine the PID for IMS Connect by issuing the UNIX command: D OMVS,V.

As the number of sockets on an IMS Connect port approaches the MAXFILEPROC parameter value, UNIX System Services issues message BPXI040I. For example, BPXI040I PROCESS LIMIT MAXFILEPROC HAS REACHED 85% OF ITS CURRENT 404. The BPXI040I message is displayed by UNIX System Services only if LIMMSG is set to SYSTEM or ALL in the SYS1.PARMLIB(BPXPRMxx) data set or using the SETOMVS command.

When the MAXFILEPROC parameter value is reached, IMS Connect issues the following messages:

- HWSP1415E TCP/IP SOCKET FUNCTION CALL FAILED; F=ACCEPT4, R=-1, E=124, M=SDCO
- HWSS0771W LISTENING ON PORT=*portid* FAILED; R=*rc*, S=*sc*, M=*mc*

When the MAXSOC parameter limit is reached, IMS Connect issues only the HWSS0771W message. When the MAXFILEPROC parameter limit is reached, IMS Connect issues both the HWSP1415E and the HWSS0771W messages.

For more information about UNIX System Services superusers and the UNIX System Services parameter MAXFILEPROC, see *z/OS UNIX System Services Planning*.

Socket number warnings

IMS Connect issues warning message HWSS0772W when the number of sockets reaches the default warning threshold of 80 % of the maximum allowable number of sockets set by the MAXSOC parameter.

Socket number warnings

If the number of sockets continues to increase past 80 %, IMS Connect issues a new HWSS0772W message at a default increment of every 5 % increase in the number of sockets. When the number of sockets reaches the maximum allowable number, IMS Connect refuses any new connections and issues message HWSS0771W.

You can set the warning threshold by using the WARNSOC parameter on the TCPIP statement of the IMS Connect HWSCFGxx PROCLIB member. You can set the warning increment by using the WARNINC parameter on the TCPIP statement.

To prevent insignificant fluctuations in the number of sockets from flooding the console with messages, until the warning mechanism is reset, IMS Connect issues a HWSS0772W messages only once for the warning threshold and only once for each warning increment thereafter.

IMS determines the reset threshold by subtracting either twice the value of the WARNINC parameter or 5 % from the value of the MAXSOC parameter, whichever results in a lower reset threshold. When the reset threshold is reached, IMS Connect resets the warning mechanism and issues message HWSS0773I.

By default, IMS Connect resets the warning mechanism when the number of sockets falls to 70 % of the value of the MAXSOC parameter, which is the default value of the MAXSOC parameter less 10 %, which is two times the default value of the WARNINC parameter.

Related reference:

➞ HWSCFGxx member of the IMS PROCLIB data set (System Definition)

Related information:

➞ HWSS0771W (Messages and Codes)

➞ HWSS0772W (Messages and Codes)

➞ HWSS0773I (Messages and Codes)

Resolving duplicate client IDs

If a failure occurs between IMS Connect and the IMS Connect client, either because the connection failed or the client terminated unexpectedly, and IMS Connect does not detect the failure before the client attempts to reconnect with the same client ID, IMS Connect prevents the client from reconnecting because the client ID is a duplicate of the ID associated with the original failed connection.

To resolve duplicate client ID conditions after a failure you can have the incoming connection cancel the original client ID or you can cancel the IMS Connect message timer for the socket connection that the client was originally connected to. Canceling the client ID is the easier of these two options.

To decrease the chances of a duplicate client ID condition occurring in the first place, you can specify a small KeepAlive interval for all the socket connections on a given port. Specifying a small KeepAlive value helps IMS Connect detect and clean up failed client connections earlier, potentially before the client attempts to reconnect; however, a small KeepAlive value can also increase network traffic.

IMS TM Resource Adapter clients can also have IMS Connect generate unique client IDs for incoming connections, thereby avoiding the possibility that an incoming connection from IMS TM Resource Adapter might specify an ID that is a duplicate of an ID that is being used by an existing connection.

Canceling connections that have duplicate client IDs

When an incoming connection to IMS Connect specifies a client ID that is a duplicate of a client ID used by another existing connection, you can cancel the existing connection so that IMS Connect accepts the new connection. When a client ID is canceled, IMS Connect discards the existing connection, regardless of its state, and issues message HWSS0743I DUPLICATE CLIENT ID TERMINATED.

Canceling a connection with a duplicate client ID is supported for clients that use either persistent sockets or transactions sockets with either commit-then-send (CM0) or send-then-commit (CM1) transactions.

If a duplicate client ID is canceled, the previous session, if one exists, is cleaned up and the new request is completed. Then, if the socket type is a persistent socket, the session is maintained and if the socket is a transaction socket, the session is terminated.

For user-written IMS Connect client applications, you can cancel duplicate client ID connections in the following ways:

- By coding the IMS Connect client application to specify X'80' (IRM_F3_CANCID) in the IRM_F3 field of the user section of the IRM message header mapped by the HWSIMSCB macro when the client establishes a new connection with IMS Connect.
- By coding any of the following exit routines to specify X'20' (OMUSR_CANCID) in the OMUSR_FLAG1 field of the user data section of the OTMA message header mapped by the HWSOMPFX macro.

If IMS Connect has already detected the error and terminated the original connection, IMS Connect ignores any specification to cancel a duplicate client ID.

Canceling the message timer to resolve a duplicate client ID

You can resolve a duplicate client ID condition by canceling the IRM timer on the original failed client connection; however, the steps required to cancel a timer, which include issuing ACKs, SENDs, and READs, and connecting and disconnecting, make canceling the duplicate client ID the recommended method for resolving duplicate client ID conditions.

Auto-generated IDs for IMS TM Resource Adapter

For users of the IMS TM Resource Adapter, IMS Connect can generate a unique user ID if the user ID provided by IMS TM Resource Adapter on a new connection request is a duplicate of a user ID already in use by an existing connection.

After IMS Connect generates a user ID, the generated ID is used in all replies returned to IMS TM Resource Adapter.

When an IMS TM Resource Adapter client requests that IMS Connect generate a user ID, the OMUSR_FLAG2 field is set to X'40' (OMUSR_F2_CIDREQ) in the OTMA message prefix. When IMS Connect returns a reply that includes a generated user ID, OMUSR_FLAG2 is set to X'20' (OMUSR_F2_CIDGEN).

Related tasks:

“Canceling a message timer” on page 325

Related reference:

“Format of user portion of IRM for HWSSMPL0, HWSSMPL1, and user-written message exit routines” on page 241

“OTMA user data fields used by IMS Connect” on page 283

IMS Connect override for the z/OS TCP/IP KeepAlive interval

IMS Connect can override the default KeepAlive interval that is defined in z/OS for TCP/IP socket connections.

The KeepAlive function, which is a function provided by the TCP/IP protocol, can detect certain socket error conditions by sending a KeepAlive packet on sockets that have been inactive for a specified interval. For example, the KeepAlive function can detect sockets that are no longer valid because the client has abruptly disconnected without informing IMS Connect.

By default, IMS Connect accepts the specification set in the z/OS layer for TCP/IP sockets; however, if a KeepAlive interval set by z/OS is large, as might be the case for installations trying to reduce network traffic, the large interval can delay the detection of an invalid socket by IMS Connect.

To help detect error conditions earlier on IMS Connect sockets, you can specify a smaller KeepAlive interval for IMS Connect ports. Each port defined to IMS Connect can specify a different KeepAlive interval. The KeepAlive interval specified for each port applies to all sockets that use that port.

Typically, a network manager determines if IMS Connect sockets require a different KeepAlive interval from the value defined by z/OS for the TCP/IP stack. After the KeepAlive intervals for IMS Connect are determined, an IMS System Programmer updates either the PORT or DRDAPORT parameters in the IMS Connect HWSCFGxx PROCLIB member with the KeepAlive intervals.

For IMS Connect clients that access IMS TM, KeepAlive intervals can be specified by using the KEEPAV parameter of the PORT keyword in the TCPIP configuration statement in the HWSCFG PROCLIB member.

For IMS Connect clients that access IMS DB, KeepAlive intervals can be specified by using the KEEPAV parameter of the DRDAPORT keyword in the ODACCESS configuration statement in the HWSCFG PROCLIB member.

The range of valid values for the KEEPAV parameter is defined by the TCP/IP protocol and is from 1 to 2 147 460 seconds. A KEEPAV value of zero accepts the KeepAlive interval set in the z/OS layer for the TCP/IP stack. KEEPAV=0 is the default.

You can display the current KeepAlive interval for by issuing any of the following commands:

- VIEWHWS
- VIEWPORT
- QUERY MEMBER
- QUERY PORT

An IMS Connect-defined KeepAlive interval is not supported by ports defined by using the PORTID parameter, by “LOCAL” ports, or by SSL ports. If the KEEPAV parameter is specified for any of these ports, IMS Connect abends during initialization.

Related reference:

➡ HWSCFGxx member of the IMS PROCLIB data set (System Definition)

TCP/IP failures

If active IMS Connect ports unexpectedly lose their connection to the TCP/IP services provided by the TCP/IP network or the z/OS TCP/IP stack, IMS Connect automatically reconnects the ports when the TCP/IP services becomes available again.

To detect when the TCP/IP services become available, IMS Connect continues listening on all active ports. When communication with the TCP/IP services is resumed, IMS Connect reconnects the ports and issues message HWSS0780I.

If communication between a port and the TCP/IP services was terminated by a command, IMS Connect does not automatically resume communication with the TCP/IP services if they should become available.

IMS Connect timeout specifications

You can specify timeout intervals to IMS connect for various stages of the communication process and for various types of interactions.

How you specify each timeout interval differs depending on whether you are using an IMS Connect client that accesses IMS DB or an IMS Connect client that access IMS TM.

Related concepts:

“IMS Connect client support” on page 172

Timeout specifications for IMS DB clients

For IMS Connect clients that access IMS DB, you can specify timeout values on two parameters in the ODACCESS configuration statement.

The ODBMTMOT parameter controls the amount of time that IMS Connect waits for both:

- A response message on connections with ODBM
- An initial input message after a socket connection is established on connections with a client application

The PORTTMOT keyword controls how long IMS Connect keeps an existing connection open after the client stops sending input.

Related reference:

 ODACCESS statement (System Definition)

Timeout specifications for IMS TM clients

For IMS Connect clients that access IMS TM, you can specify timeout values in the TCP/IP and DATASTORE configuration statements, as well as in the IMS Request Message (IRM) header of input messages.

On the TCPIP configuration statement, the TIMEOUT keyword controls two aspects of communication wait time:

- How long IMS Connect keeps a connection open if the client does not send any input after the connection is first established
- For input messages that do not specify a timeout interval in the IRM, how long ICON waits for a response from IMS before notifying the client of the timeout and returning the socket connection to a RECV state

On the DATASTORE configuration statement, the ACKTO keyword controls how long OTMA waits for acknowledgements from IMS Connect before rerouting the output to a tpipe hold queue.

The timeout values specified on the TCPIP and DATASTORE configuration statements can also be specified in the IRM message header on messages received from IMS TM clients.

The IRM can also be used to specify how long the client waits for output after issuing a resume tpipe request.

Related reference:

 DATASTORE statement (System Definition)

TCPIP statement (System Definition)

Timeout specifications on input messages

Each and every input message from the IMS Connect client can set a different timeout value in the IRM_TIMER field of the fixed portion of the IMS request message (IRM) header.

Set the IRM_TIMER value to an appropriate wait time for IMS to return data to IMS Connect.

The settings for the IRM_TIMER is enforced as described in the following list:

1. If the IRM_TIMER is set at X'00', the following default values are used:
 - The default for all RESUME_TPIPE is two seconds.
 - The default for all RESUME_TPIPE non-single ACK is .25 seconds.
 - The value of the TIMEOUT parameter in the IMS Connect TCPIP configuration statement for all others.
2. X'FF' and X'01' - X'9E' are used only when requested.
3. X'E9' (char Z) NO_WAIT means do not wait for any IMS output. NO_WAIT is not valid on some Client SENDs. Because IMS Connect does not wait for output from IMS, on a transaction socket connection, IMS Connect disconnects the socket; and on a persistent socket connection, IMS Connect requests the next input from the client rather than disconnect the socket. If NO_WAIT is used, it is enforced as follows:
 - There is a two second delay for:
 - RESUME_TPIPE request
 - conversational transaction code
 - conversational data
 - ACK or NAK associated with a conversational transaction
 - non-conversational transaction code
 - A .25 second delay for each of the following is used:
 - an ACK or NAK associated with a non-conversational transaction commit mode one confirm
 - an ACK or NAK associated with a RESUME_TPIPE with Asynch output options AUTO or NOAUTO
 - an ACK or NAK associated with non-conversational transaction commit mode zero confirm
 - NO_WAIT can be used for the following:
 - a SENDONLY
 - an ACK or NAK associated with RESUME_TPIPE with Asynch output option SINGLE

Misuse of X'E9' can result in one of the following problems:

1. The socket disconnects.
2. An output message to the client on a transaction socket is lost.
3. A hang condition occurs between the client and IMS Connect or IMS Connect and OTMA. For example, the client can be in a READ state waiting for output from IMS Connect while IMS Connect is in a READ state waiting for input from the client and OTMA is in READ state waiting for acknowledgment.
4. The deallocate commit or deallocate abort notification for CM1 SynchLevel=Confirm is lost.

5. Other unpredictable conditions occur.

To determine the appropriate wait time for IMS to return data to IMS Connect, consider the following guidelines:

- For a client SEND of transaction code and data or data only, the IRM_TIMER value should be set to reflect the amount of time IMS Connect should wait for the output from IMS.

Recommendation: Do not use X'E9'.

- If the client application knows that the last message received is the last output message to the client for the transaction, set the IRM_TIMER to X'01' (.01 of a second) for a client SEND of ACK or NAK. The IRM_TIMER of X'01' is the smallest value that can be set for non-RESUME TPIPE ACK messages. However, if the ACK message is associated with an output from a RESUME TPIPE call, do **not** set the IRM_TIMER value to X'E9' (character Z).
- For a client SEND of a RESUME TPIPE call, the timer value can be set as follows:

AUTO option

X'FF' for dedicated output device, or any X'00' to X'9E' values for non-dedicated output device

NOAUTO option

any value other than X'FF' or X'E9'

SINGLE or SINGLE with WAIT option

any value other than X'FF' or X'E9'

Timer interval specifications:

You can specify timer values in several incremental ranges.

The values in each range are selected by entering a hexadecimal value in the IRM_TIMER field in the IRM of the input message from the client. The hexadecimal values that can be specified and the time intervals that they represent are:

- Increments of one one-hundredth of a second are represented by values of X'01' to X'19', as shown in Table 52.
- Increments of five one-hundredths of a second are represented by values of X'1A' to X'27', as shown in Table 53 on page 321.
- Increments of one second are represented by values of X'28' to X'63', as shown in Table 54 on page 321.
- Increments of one minute are represented by values of X'63' to X'9E', as shown in Table 55 on page 323.
- Default timer values, the no-timer option, and the indefinite wait option, specified by X'00', X'E9', and X'FF' respectively, are shown in Table 56 on page 325.

The following table lists the IRM_TIMER values and their corresponding time in increments of one one-hundredth of a second.

Table 52. IRM_TIMER values in one one-hundredth of a second

Time	Hexadecimal value
.01 of a second	X'01'

Table 52. IRM_TIMER values in one one-hundredth of a second (continued)

Time	Hexadecimal value
.02 of a second	X'02'
.03 of a second	X'03'
.04 of a second	X'04'
.05 of a second	X'05'
.06 of a second	X'06'
.07 of a second	X'07'
.08 of a second	X'08'
.09 of a second	X'09'
.10 of a second	X'0A'
.11 of a second	X'0B'
.12 of a second	X'0C'
.13 of a second	X'0D'
.14 of a second	X'0E'
.15 of a second	X'0F'
.16 of a second	X'10'
.17 of a second	X'11'
.18 of a second	X'12'
.19 of a second	X'13'
.20 of a second	X'14'
.21 of a second	X'15'
.22 of a second	X'16'
.23 of a second	X'17'
.24 of a second	X'18'

Table 52. IRM_TIMER values in one one-hundredth of a second (continued)

Time	Hexadecimal value
.25 of a second	X'19'

The following table lists the IRM_TIMER values and their corresponding time in five one-hundredths of a second.

Table 53. IRM_TIMER values in five one-hundredths of a second

Time	Value
.30 of a second	X'1A'
.35 of a second	X'1B'
.40 of a second	X'1C'
.45 of a second	X'1D'
.50 of a second	X'1E'
.55 of a second	X'1F'
.60 of a second	X'20'
.65 of a second	X'21'
.70 of a second	X'22'
.75 of a second	X'23'
.80 of a second	X'24'
.85 of a second	X'25'
.90 of a second	X'26'
.95 of a second	X'27'

The following table lists the IRM_TIMER values for time increments of one second each.

Table 54. IRM_TIMER time values in seconds

Time	Value
1 second	X'28'
2 seconds	X'29'
3 seconds	X'2A'
4 seconds	X'2B'

Table 54. IRM_TIMER time values in seconds (continued)

Time	Value
5 seconds	X'2C'
6 seconds	X'2D'
7 seconds	X'2E'
8 seconds	X'2F'
9 seconds	X'30'
10 seconds	X'31'
11 seconds	X'32'
12 seconds	X'33'
13 seconds	X'34'
14 seconds	X'35'
15 seconds	X'36'
16 seconds	X'37'
17 seconds	X'38'
18 seconds	X'39'
19 seconds	X'3A'
20 seconds	X'3B'
21 seconds	X'3C'
22 seconds	X'3D'
23 seconds	X'3E'
24 seconds	X'3F'
25 seconds	X'40'
26 seconds	X'41'
27 seconds	X'42'
28 seconds	X'43'
29 seconds	X'44'
30 seconds	X'45'
31 seconds	X'46'
32 seconds	X'47'
33 seconds	X'48'
34 seconds	X'49'
35 seconds	X'4A'
36 seconds	X'4B'
37 seconds	X'4C'
38 seconds	X'4D'
39 seconds	X'4E'
40 seconds	X'4F'
41 seconds	X'50'
42 seconds	X'51'
43 seconds	X'52'
44 seconds	X'53'

Table 54. IRM_TIMER time values in seconds (continued)

Time	Value
45 seconds	X'54'
46 seconds	X'55'
47 seconds	X'56'
48 seconds	X'57'
49 seconds	X'58'
50 seconds	X'59'
51 seconds	X'5A'
52 seconds	X'5B'
53 seconds	X'5C'
54 seconds	X'5D'
55 seconds	X'5E'
56 seconds	X'5F'
57 seconds	X'60'
58 seconds	X'61'
59 seconds	X'62'
60 seconds	X'63'

The following table lists the IRM_TIMER values and their corresponding time increments of one minute each.

Table 55. IRM_TIMER time values in minutes

Time	Value
1 minute	X'63'
2 minutes	X'64'
3 minutes	X'65'
4 minutes	X'66'
5 minutes	X'67'
6 minutes	X'68'
7 minutes	X'69'
8 minutes	X'6A'
9 minutes	X'6B'
10 minutes	X'6C'
11 minutes	X'6D'
12 minutes	X'6E'
13 minutes	X'6F'
14 minutes	X'70'
15 minutes	X'71'
16 minutes	X'72'
17 minutes	X'73'
18 minutes	X'74'
19 minutes	X'75'

Table 55. IRM_TIMER time values in minutes (continued)

Time	Value
20 minutes	X'76'
21 minutes	X'77'
22 minutes	X'78'
23 minutes	X'79'
24 minutes	X'7A'
25 minutes	X'7B'
26 minutes	X'7C'
27 minutes	X'7D'
28 minutes	X'7E'
29 minutes	X'7F'
30 minutes	X'80'
31 minutes	X'81'
32 minutes	X'82'
33 minutes	X'83'
34 minutes	X'84'
35 minutes	X'85'
36 minutes	X'86'
37 minutes	X'87'
38 minutes	X'88'
39 minutes	X'89'
40 minutes	X'8A'
41 minutes	X'8B'
42 minutes	X'8C'
43 minutes	X'8D'
44 minutes	X'8E'
45 minutes	X'8F'
46 minutes	X'90'
47 minutes	X'91'
48 minutes	X'92'
49 minutes	X'93'
50 minutes	X'94'
51 minutes	X'95'
52 minutes	X'96'
53 minutes	X'97'
54 minutes	X'98'
55 minutes	X'99'
56 minutes	X'9A'
57 minutes	X'9B'
58 minutes	X'9C'
59 minutes	X'9D'

Table 55. IRM_TIMER time values in minutes (continued)

Time	Value
60 minutes	X'9E'

The following table lists additional options that you can specify in the IRM_TIMER field and the value you use to specify them.

Table 56. Additional IRM_TIMER options

Timer option	Value
Use default values. For RESUME TPIPE calls and associated ACK messages, the default is .25 seconds. For all other SENDs, the default is the configuration file TIMEOUT value.	X'00'
Do not wait.	X'E9' C'Z'
Wait indefinitely. This setting is intended to support the auto option of the asynchronous output function.	X'FF'

Canceling a message timer:

User-written IMS Connect client applications can cancel the active message timer when waiting on output from the data store.

The cancel timer feature prevents IMS Connect clients that have specified a large timeout interval from being lost in the event that the data store does not send a reply. Without the cancel timer feature an IMS Connect STOPCLNT command would have to be issued to clear the socket connection. When a request to cancel the timer is submitted, IMS Connect notifies the client.

The cancel timer feature is supported by user-written message exit routines and the sample user message exit routines HWSSMPL0 and HWSSMPL1.

Note that if a client is waiting in a CONN state after issuing a RESUME TPIPE call and the data store is closed by IMS or a STOPDS command, then the client receives an RSM message with RC= X'2C' (or decimal 44).

A cancel timer request is specified by a C in the IRM_F4 field and can be submitted from either a single instance of a client or two instances of the client with the same client ID.

To submit a cancel timer request from a single client instance:

1. Issue SEND of ACK.
2. Set local timer.
3. Issue READ for Response. Timer pops rather than receipt of data.
4. Issue Disconnect.
5. Issue Connect.
6. Issue SEND with Cancel Timer set in the IRM.
7. Issue Read for Cancel Timer. The user message exit issues return code 8 with a reason code of X'3B' (or decimal 59) in the RSM.
8. Issue Disconnect.

To submit a cancel timer request from two instances of the same client with the same client ID:

1. From first client
 - a. Issue SEND of ACK (with, for example, a client ID of ICON01).
 - b. Issue READ for Response.
2. From second client instance
 - a. Issue Connect.
 - b. Issue SEND with Cancel Timer set in the IRM (with client ID of ICON01).
 - c. Issue Read for Cancel Timer. The user message exit issues return code 8 with a reason code of X'3B' (or decimal 59) in the RSM.
 - d. Issue Disconnect.
3. First client receives an RSM with a return code of X'2C' (or decimal 44) and a reason code set to the value of the timer value. This instance remains connected and is in RECV state.
 - a. Issue Disconnect or continue processing.

Timeout intervals for IMS Connect acknowledgments to OTMA

You can specify a timeout interval that determines how long OTMA waits for an acknowledgment from IMS Connect. You can also specify a timeout tpipe queue to hold commit-then-send (CM0) output after the timeout interval has expired.

You can specify a timeout interval by using the ACKTO parameter in the DATASTORE configuration statement. The timeout value specified on the ACKTO parameter applies to acknowledgments sent to OTMA for the following message types:

- CM0 output
- Send-then-commit (CM1) output

For CM0 output, when the timeout interval expires, OTMA removes the output from the tpipe queue and reroutes the output to either:

- A specified reroute tpipe queue
- A specified timeout tpipe queue
- The default OTMA timeout tpipe queue DFS\$\$TOQ

You can specify the name of a timeout tpipe queue for CM0 output on the CM0ATOQ parameter in both the HWS and DATASTORE configuration statements.

On the HWS configuration statement, the CM0ATOQ parameter defines a default timeout tpipe queue for all data store connections defined to an instance of IMS Connect. On the DATASTORE statement, the CM0ATOQ parameter defines the timeout tpipe queue to be used by only the defined data store connection. Specifications for CM0ATOQ on the DATASTORE statement override the specification for CM0ATOQ made on the HWS configuration statement.

CM1 output is not rerouted to a timeout queue, because OTMA discards CM1 output if the timeout interval expires.

When a timeout occurs, and the IMS application does not reply to the IOPCB or complete a message switch to another transaction, OTMA issues a DFS2082 message for CM1 input messages, regardless of the transaction response mode.

Related information:

IMS Connect transaction expiration support

IMS Connect can adjust the expiration time for IMS transactions to match the timeout value of the socket connection on which the transaction is submitted.

If an expiration time is specified in IMS, transactions can expire and be discarded if IMS does not process them before the expiration time is exceeded. The expiration time is set in the IMS definition of the transaction by the EXPRTIME parameter on either the TRANSACT stage-1 system definition macro or either of the dynamic resource definition type-2 commands CREATE TRAN or UPDATE TRAN. A transaction expiration time set by IMS Connect overrides any transaction expiration time specified in the definition of the transaction in IMS.

The following IMS Connect user message exits support setting a transaction expiration time:

- HWSSMPL0
- HWSSMPL1
- HWSSOAP1

IMS Connect clients that use HWSSMPL0, HWSSMPL1, or a user-written exit routine can instruct IMS Connect to set the expiration time for transactions by specifying X'01' (IRM_F1_TRANEXP) in the user section of the IRM prefix. IMS Enterprise Suite SOAP Gateway clients can modify the HWSSOAP1 user message exit routine to set the IRM_F1_TRANEXP field.

To match the expiration time to the timeout value of the socket connection, IMS Connect calculates when the client socket connection will time out by reading the value of either the TIMEOUT parameter of the TCPIP configuration statement or the IRM_TIMER field in the fixed section of the IRM prefix of input messages. IMS Connect then places that time in the user data section of the OTMA message prefix in store clock (STCK) format and notifies OTMA of the new expiration time by setting fields OMHDRXP1 and OMHDRSXP in the state data section of the OTMA message prefix.

When transaction expiration is enabled for transactions submitted by IMS Connect, OTMA monitors the processing of the transaction and, if the transaction expiration time has passed, OTMA discards the transaction and returns a message to IMS Connect.

Depending on when the transaction times out, OTMA might return either a NAK response to IMS Connect or message DFS3688I.

For detailed information about when OTMA checks for expired transactions, see OTMA support for transaction expiration (Communications and Connections).

If IMS Connect does not set a transaction expiration time, the transaction can still expire if the EXPRTIME parameter was specified when the transaction was originally defined to IMS.

Related concepts:

“OTMA support for transaction expiration” on page 797

Related reference:

 HWSCFGxx member of the IMS PROCLIB data set (System Definition)

➡ TRANSACT macro (System Definition)

Setting a transaction expiration time with IMS Connect

You can instruct IMS Connect to set an expiration time for transactions submitted to IMS that matches the timeout value of the socket connection on which the transaction is submitted.

For IMS Connect to set a transaction expiration time, a timeout value for the socket connection must first be specified in either the TIMEOUT parameter of the TCPIP configuration statement or the IRM_TIMER field in the fixed section of the IRM prefix.

You can instruct IMS connect to set a transaction expiration time for a transaction by using one of the following methods:

- Specify X'01' (IRM_F1_TRANEXP) in the IRM_F1 field of the user section in the IRM prefix of the input message. When IRM_F1 is set to X'01', the user message exit sets the OMHDRIST field to X'01' (OMHDRXP1) in the state data section of the OTMA header.
- Modify any of the following user message exit routines to set the IRM_F1 field to X'01' (IRM_F1_TRANEXP).
 - HWSSMPL0
 - HWSSMPL1
 - HWSSOAP1

When IRM_F1 is set to X'01', the user message exit sets the OMHDRIST field to X'01' (OMHDRXP1) in the state data section of the OTMA header.

Related reference:

“IRM structures for IMS Connect client messages” on page 237

“OTMA state data fields used by IMS Connect” on page 271

➡ HWSCFGxx member of the IMS PROCLIB data set (System Definition)

Retrieval of output on OTMA tpipe hold queues

OTMA uses tpipe hold queues to send output to IMS Connect that must be queued for delivery. To retrieve output from the tpipe hold queue, the IMS Connect client issues a RESUME TPIPE request that specifies the name of the tpipe on which the output messages are queued.

The types of output messages that OTMA sends via the tpipe hold queue include:

- Synchronous callout messages that IMS application programs send by issuing the DL/I ICAL call.
- Output, such as asynchronous callout messages, that IMS application programs send by issuing the DL/I ISRT call to an alternate PCB (ALTPCB).
- CM0 IOPCB output messages for which the receiving OTMA client returned a NAK.
- Response messages to CM0 SendOnly input.

For CM0 transactions sent by a single client that uses a send and receive interaction, the name of the tpipe is typically the client ID. The output from the CM0 transaction can be retrieved in the following ways:

- By the original client with the client ID that matches the tpipe name.

- If the original client terminated, by another client that uses the same client ID that matches the tpipe name.
- By another client that specifies the tpipe name as an alternate client ID in the RESUME TPIPE request.

For synchronous callout messages, the tpipe name is typically defined in IMS on an OTMA destination descriptor. For asynchronous callout messages, the tpipe name can be defined by either an OTMA destination descriptor or an OTMA routing exit routine. For both types of callout messages, the tpipe name is usually specified as an alternate client ID in the RESUME TPIPE request that retrieves the callout messages. The alternate client ID must match the tpipe name that is specified in the OTMA destination descriptor or in the OTMA routing exit routines.

Note: Synchronous callout request messages are handled by OTMA and IMS Connect in much the same way as asynchronous output. That is, synchronous callout request messages are retrieved by issuing a RESUME TPIPE call. Many of the same rules and guidelines for retrieving asynchronous output also apply to retrieving synchronous callout request messages.

IMS Connect communicates the presence of asynchronous output to the client from a CM0 output response message in one of the following ways:

- By returning the flag CSM_AMSG in the CSM_FLG1 field in the CSM (complete status message).
- By returning the flag RSM_AMSG in the RSM_FLG1 field in the RSM (request status message).

If you do not use IMS Connect to retrieve output from an OTMA tpipe hold queue, your client application does not need to analyze the CSM or the RSM. IMS Connect communicates the presence of asynchronous output regardless of whether a client application requests the asynchronous output.

Use the RESUME TPIPE call to retrieve the asynchronous output from the client. You can retrieve asynchronous output on both persistent and transaction sockets.

Restriction: The IMS TM Resource Adapter supports only the asynchronous option, SINGLE.

RESUME TPIPE/receive protocol

IMS Connect clients use the RESUME TPIPE protocol to retrieve commit-then-send (CM0) output or synchronous callout requests from a tpipe hold queue in IMS.

Synchronous callout requests are issued by IMS application programs running in an IMS dependent region. Callout requests are for data or services from a provider that is external to the IMS installation. In a callout scenario, the IMS application is the client and the external provider is the server.

Asynchronous output from IMS can include response messages from an IMS transaction and asynchronous callout messages.

For user-written IMS Connect client application programs, the RESUME TPIPE call is coded by specifying the following field values in the IRM prefix:

IRM_F4

R character value (IRM_F4_RESUMET).

IRM_F5

Hexadecimal value that specifies the retrieval option for the RESUME TPIPE call. The retrieval option specifies how many messages a single RESUME TPIPE call can retrieve and how long the IMS Connect client waits for additional messages when no messages remain on the tpipe hold queue.

IRM_TIMER

The timeout value for the RESUME TPIPE call. The value of IRM_TIMER determines how long the client waits for new messages if no output is currently on the tpipe hold queue.

When retrieving synchronous callout messages, in addition to the above fields, you must also specify the following fields to set the IRM architecture level (IRM_ARCH3) and specify whether the IMS Connect client can retrieve synchronous callout message only, or asynchronous output as well:

IRM_ARCH

X'03' (IRM_ARCH3). IRM_ARCH3 is required when retrieving synchronous callout messages.

IRM_F0

One of the following values must be specified when retrieving synchronous callout messages:

- X'80' (IRM_F0_SYNONLY). When IRM_F0_SYNONLY is specified, only synchronous callout messages are returned to the IMS Connect client.
- X'40' (IRM_F0_SYNASYN). When IRM_F0_SYNASYN is specified, both synchronous callout messages and asynchronous output are returned to the IMS Connect client.

Related reference:

“Format of fixed portion of IRM in messages sent to IMS Connect” on page 238

“Format of user portion of IRM for HWSSMPL0, HWSSMPL1, and user-written message exit routines” on page 241

“Timer interval specifications” on page 319

Examples flows for the RESUME TPIPE protocol

The following figures show examples of the RESUME TPIPE protocol to receive asynchronous commit-then-send (CM0) output.

In the following figure, the CM0 flow enqueues IMS output before sending it to the client with the client application sending a positive acknowledgment (ACK) for both outputs. The ACK removes the output from the IMS queue.




Figure 31. Commit-then-send, receive asynchronous output (client waits for output) flow

The sample flow in the previous figure assumes the following:

- The client sends the OTMA call RESUME TPIPE to ask IMS OTMA to post the named Tpipe (the client name).
- The client issues a RECEIVE request to receive the output from IMS.
- The client sends ACK to IMS (required for commit-then-send).
- The client receives the next output from IMS.
- The client sends ACK to IMS.
- The client waits for the next output from IMS, or for Time out notification.

In the following figure, the commit-then-send flow enqueues IMS output before sending it to the client and the client application sends a positive acknowledgment (ACK) for the first output (removing the output from the IMS queue) and a NAK to the second output (which results in the output remaining in the queue).

Requirement: Use this protocol with the timeout function. Otherwise, the client hangs if there are no more messages to send.

1. Issue a CONNECT command.
2. A TCP/IP SEND of an OTMA RESUME TPIPE call, immediately followed by a TCP/IP READ function from the primary client application.
3. A TCP/IP SEND of an ACK or NAK response on the receipt of the output message. If the ACK was sent with a timer value of NOWAIT (NOWAIT is only valid for a RESUME TPIPE call with SINGLE or SINGLE with WAIT option), go to step 5. If NAK was sent, go to step 5.
4. A TCP/IP READ function from the primary client application. Repeat steps 2 and 3 until either all messages have been received, until the end user has received all of the messages that they want, until an error occurs, or until time out notification occurs.
5. Issue a DISCONNECT command, if you are using transaction sockets. If you are using persistent sockets, the connection is still connected.

Enabling end user asynchronous output requests

You can easily implement the CONNECT, RESUME TPIPE, READ, ACK/NAK, and DISCONNECT functions on the client application's screen with the buttons on the graphical user interface.

- Create a CONNECT button.
- Create a RESUME TPIPE button to send a RESUME TPIPE call to IMS Connect. IMS Connect will then send a RESUME TPIPE request to OTMA.
- Create a READ button to issue a TCP/IP READ request. OTMA will send a message to IMS Connect following the RESUME TPIPE or ACK response.
- Create an ACK/NAK button.
- You can also combine the READ and ACK requests into a single button that issues the READ request, then sends an ACK on receiving the message.
- Create a DISCONNECT button.

Managing the retrieval of output messages

When retrieving either asynchronous output messages or synchronous callout request messages with the RESUME TPIPE call, you have options regarding how messages are returned.

You specify your retrieval options in the IRM of the RESUME TPIPE request message and the user message exit sets the options in the OTMA header.

The retrieval options that you can specify include:

- Single
- Single with wait
- Noauto
- Nooption
- Auto

The IMS Connect user message exits HWSSMPL1 and HWSSMPL0 support all these options. To choose a type of message control, the client code sets the IRM field IRM_FLG5 to be one of the following values:

IRM_F5_ONE

Retrieves a single message (single).

IRM_F5_SWAIT

Waits for a single message if none are currently present in the IMS message queue (single with wait).

IRM_F5_NOAUTO

Retrieves all messages that have been queued (noauto).

IRM_F5_AUTO

Retrieves all messages that have been queued, then retrieves any additional messages that are queued later (auto).

IRM_F5

Makes a RESUME TPIPE call function like NOAUTO (nooption) when set to X'00'.

The HWSSMPL0 and HWSSMPL1 user message exits default to the noauto type of asynchronous output message management.

The following subsections describe the asynchronous output message control options in detail.

Single message control

When using the single message control option (by setting field IRM_F5 to IRM_F5_ONE), the client can receive only a **single message**.

If there are no messages in the IMS OTMA Asynchronous Queue for the client ID when the request is made, no message will be returned and a time out will occur. Using the single message control option will force the following sequence of events to occur:

1. Client issues CONNECT function.
 - a. Following the CONNECT function, if the socket type is persistent socket, one or more transactions can be sent and the responses received before RESUME TPIPE call processing.
 - b. If the socket type is a transaction socket, the RESUME TPIPE call must be issued after the CONNECT function.
2. Client issues RESUME TPIPE call with the correct IRM settings.
3. Client issues RECEIVE function to receive the Asynchronous output.
4. Client sends ACK or NAK to IMS Connect.
 - a. The ACK or NAK can be sent with a user-specified numeric timeout value.
Or
 - b. Specify NOWAIT for the timeout value.
5. If a numeric timeout value is specified, the client must issue a RECEIVE function to receive the timeout notification. If the NOWAIT option is specified, no timeout notification is sent. Therefore, the client must not issue a RECEIVE function if NOWAIT is specified.
6. IMS Connect disconnects the Socket from the Host end if the socket connection is a transaction socket. If the socket connection is a persistent socket, IMS Connect does not disconnect the socket.
7. Client must issue a DISCONNECT function if the socket connection is a transaction socket. If the socket is a persistent socket, the client can either DISCONNECT the socket or choose to send in a new request such as SENDONLY, SEND of transaction code and data, or issue another RESUME TPIPE call.

If the client responds with a NAK rather than an ACK, the message that has been NAKed will be put back on the OTMA Asynchronous Hold Queue, and can be re-retrieved later. IMS Connect will continue to process as described in events five through seven when a NAK is sent to IMS Connect by the Client.

Single with wait message control

When using the single with wait message control option (by setting IRM_F5 to IRM_F5_SWAIT), the client can receive only one single message; however, unlike single message control, the single with wait message control can receive a message that is placed in the IMS OTMA Asynchronous Queue for the client ID.

Using the single with wait message control option will force the following sequence of events to occur:

1. Client issues CONNECT function.
 - a. Following the CONNECT function, if the socket type is persistent socket, one or more transactions can be sent and the responses received before RESUME TPIPE call processing.
 - b. If the socket type is transaction socket, then the RESUME TPIPE call processing must be issued after the CONNECT function.
2. Client issues RESUME TPIPE call, with the correct IRM settings.
3. Client issues RECEIVE function to receive the Asynchronous output.
4. Client sends ACK or NAK to IMS Connect.
 - a. The ACK or NAK can be sent with a user-specified timeout notification.
Or
 - b. Specify NOWAIT for the timeout value.
5. If a numeric timeout value is specified, the client must issue a RECEIVE function to receive the timeout notification. If the NOWAIT option is specified, no timeout notification is sent. Therefore, the client must not issue a RECEIVE function if NOWAIT is specified.
6. IMS Connect disconnects the Socket from the Host end if the socket connection is a transaction socket. If the socket connection is a persistent socket, IMS Connect does not disconnect the socket.
7. Client must issue a DISCONNECT function if the socket connection is a transaction socket. If the socket is a persistent socket, the client can either DISCONNECT the socket or choose to send in a new request such as SENDONLY, SEND of transaction code and data, or issue another RESUME TPIPE call.

If the client responds with a NAK rather than an ACK, the message that has been NAKed will be put back on the OTMA Asynchronous Hold Queue, and can be re-retrieved later. IMS Connect will continue to process as described in events five through seven when a NAK is sent to IMS Connect by the Client.

Noauto message control

When using the noauto message control option (by setting field IRM_F5 to IRM_F5_NOAUTO), the client can receive **all** of the messages on the OTMA Asynchronous Queue.

Using the noauto message control option will force the following sequence of events to occur:

1. Client issues CONNECT function.
 - a. Following the CONNECT function, if the socket type is persistent socket, one or more transactions can be sent and the responses received before RESUME TPIPE call processing.
 - b. If the socket type is transaction socket, then the RESUME TPIPE call processing must be issued after the CONNECT function.
2. Client issues RESUME TPIPE call.

3. Client issues RECEIVE function to receive the asynchronous output.
4. Client sends ACK message to IMS Connect.
5. Client repeats events three and four until event six occurs.
6. IMS Connect disconnects the socket from the host end, unless the socket is persistent. If the socket is persistent, IMS Connect sends a timeout message and the socket remains connected, which allows the client to either disconnect or send another request.
7. Client issues DISCONNECT function.

Using the noauto message control option, the client can always terminate by issuing a DISCONNECT function after sending an ACK message to IMS Connect.

If the client responds with a NAK message rather than an ACK, the message that received the NAK response is put back on the OTMA asynchronous hold queue and can be retrieved later. IMS Connect terminates the socket as described in event six, unless the socket is persistent.

Nooption message control

When using the nooption message control option (by setting field IRM_F5 to X'00'), the client can receive **all** of the messages on the OTMA Asynchronous Queue.

Using the nooption message control option will force the following sequence of events to occur:

1. Client issues CONNECT function.
 - a. Following the CONNECT function, if the socket type is persistent socket, one or more transactions can be sent and the responses received before RESUME TPIPE call processing.
 - b. If the socket type is transaction socket, then the RESUME TPIPE call processing must be issued after the CONNECT function.
2. Client issues RESUME TPIPE call.
3. Client issues RECEIVE function to receive the asynchronous output.
4. Client sends ACK message to IMS Connect.
5. Client repeats events three and four until event six occurs.
6. IMS Connect disconnects the socket from the host end, unless the socket is persistent. If the socket is persistent, IMS Connect sends a timeout message and the socket remains connected, which allows the client to either disconnect or send another request.
7. Client issues DISCONNECT function.

Using the nooption message control option, the client can always terminate by issuing a DISCONNECT function after sending an ACK to IMS Connect.

If the client responds with a NAK message rather than an ACK, the message that received the NAK response is put back on the OTMA asynchronous hold queue, and can be retrieved later. IMS Connect terminates the socket as described in event six above, unless the socket is persistent.

Auto message control

When using the auto message control option (by setting field IRM_F5 to IRM_F5_AUTO), the client can receive all of the messages on the OTMA Asynchronous Queue, and any messages that are placed on the OTMA Asynchronous Queue after the current messages are all removed.

Using the auto message control option will force the following sequence of events to occur:

1. Client issues CONNECT function.
2. Client issues RESUME TPIPE call.
3. Client issues RECEIVE function to receive the asynchronous output.
4. Client sends ACK message to IMS Connect.
5. Client repeats events three and four.

If all messages have been removed from the queue, event three will remain active (that is, in receive state) until the user-specified timer supplied in the IRM has expired. IMS Connect will then terminate the socket, unless the socket is persistent.

Recommendation: If event three or event five receives a disconnect of the socket, the client should disconnect and then wait for a time interval before repeating events one through five.

Using the auto message control option, the client can always terminate a non-persistent socket connection in either one of the following ways:

- Respond to the output message with a NAK response
- Send a DEALLOCATE request rather than an ACK response

The message being processed is put back on the IMS output queue, and IMS Connect terminates the socket, unless the socket is persistent.

If the client responds with an ACK message, then issues a DISCONNECT, the connection is only terminated between the client and TCP/IP; the client remains in a CONN state with IMS Connect. When IMS Connect attempts to send the next asynchronous output message, IMS Connect is notified that the connection has been lost. IMS Connect does not acknowledge (NAK) OTMA, and the message is put back on the IMS output queue. IMS Connect then terminates the socket. If the client issues an ACK message and then issues a DISCONNECT, followed by a connect and transmittal of data, IMS Connect responds with a duplicate client ID error and disconnects the socket connection.

If the client responds with a NAK message rather than an ACK in events three or five, the message that received the NAK message is put back on the OTMA asynchronous hold queue, IMS Connect terminates the socket, and the messages can be retrieved later.

Note: The IMS Connect AUTO support is based on the premise that the socket connection is dedicated as an output-only device. Combining RESUME TPIPE calls (with auto asynch option specified) with transactions on the same socket connection or SENDONLY on a persistent socket, can yield unpredictable results. If you want to change from RESUME TPIPE auto option mode to a mode that will allow for transaction processing, you must change the auto asynch option by performing one of the following options:

1. NAK one of the RESUME TPIPE outputs. This will change the asynch mode from auto to noauto. To return to auto mode, a RESUME TPIPE call with auto must be specified.
2. On a timeout notification associated with the RESUME TPIPE AUTO, the client application can disconnect, reconnect, and issue a RESUME TPIPE call with single, single with wait, noauto, or nooption with a very short IRM_TIMER value. The IRM_TIMER value should be small, so that a timeout notification can be returned immediately. Issuing a RESUME TPIPE call with one of the

four asynch mode options, changes the mode from auto to one of the specified options. After the RESUME TPIPE call is issued and a timeout notification is returned, the client application can send in a transaction.

3. On a timeout notification associated with the RESUME TPIPE AUTO, the client application can disconnect, reconnect, and issue a RESUME TPIPE call with single, single with wait, noauto, or nooption with any valid IRM_TIMER value. Upon receiving an output message, send an ACK with IRM_TIMER set to nowait or a valid value. If the IRM_TIMER value is set to nowait, the client can then send in a transaction. If the IRM_TIMER is set to a valid value, after receiving the timeout notification, the client application can then send in a transaction.

Related concepts:

“Timeout specifications on input messages” on page 318

Execution time out during RESUME TPIPE call processing with auto message control option

If you are using RESUME TPIPE calls with the auto message control option and the IRM_TIMER value times out, you might experience some unpredictable results.

If the auto option is selected on the RESUME TPIPE call and a timeout occurs, to get the timeout notification and send transactions again, you must change the auto option processing mode to noauto. To get out of the auto option processing mode, you can choose one of the following options:

- Issue a RESUME TPIPE call with the auto option and set a large IRM_TIMER value to ensure that the client application will NAK the output. When the output is NAK, OTMA will change the asynchronous mode from auto to noauto to stop the sending of asynchronous output. The client application then issues READ to retrieve the timeout notification. Upon receiving the timeout notification, the client can begin sending transactions to IMS Connect.
- Issue a RESUME TPIPE call with the noauto option and set any value in the IRM_TIMER field. After receiving ACK output, repeat READ of asynch output and SEND of ACK until a timeout notification is received. (Issuing a RESUME TPIPE call with noauto changed the processing mode from auto to noauto. This also reset the asynchronous mode in OTMA to noauto where OTMA no longer supports the automatic sending of asynch output when the IMS Message Queue is empty.) The client application then issues READ to retrieve the timeout notification. Upon receiving the timeout notification, the client can begin sending transactions to IMS Connect.
- Issue a RESUME TPIPE call with noauto option and set any value in the IRM_TIMER field. If you receive NAK output, the processing mode and OTMA Asynch mode is reset to noauto. Resetting the OTMA Asynch mode to noauto stops the sending of asynch output and the NAK output terminates the process. The client application then issues READ to retrieve the timeout notification. Upon receiving the timeout notification, the client can begin sending transactions to IMS Connect.
- Issue a RESUME TPIPE call with single option and set any value in the IRM_TIMER field. The OTMA Asynch mode is reset from auto to single and no more asynchronous messages are sent. After you receive ACK or NAK output with an IRM_TIMER setting that is anything other than NO_WAIT, the single option has been completed and the client application can issue a READ to get the timeout notification. Upon receiving the timeout notification, the client can begin sending transactions to IMS Connect.
- Issue a RESUME TPIPE call with single option and set any value in the IRM_TIMER field. The OTMA Asynch mode is reset from auto to single and no

more asynchronous messages are sent. After you receive ACK or NAK output with an IRM_TIMER setting of NO_WAIT, the single option has been completed and the client application does not have to issue a READ to get timeout notification. The client application can start sending transactions to IMS Connect.

Values for asynchronous output processing

To retrieve asynchronous output from the OTMA tpipe hold queue, the IMS Connect client application must issue a RESUME TPIPE request. To issue a RESUME TPIPE call, you need to specify the socket type, commit mode, sync level, timer setting, and RESUME TPIPE call options.

To issue a RESUME TPIPE call, specify the following values:

- Socket Type Transaction or Persistent
- Commit Mode Zero
- Sync level Confirm
- Timer setting The timeout range required by your enterprise.
- RESUME TPIPE options Single, single with wait, auto, noauto, or nooption.

For example, if you want to create a dedicated output client that only receives unsolicited output, start a client application to complete the following sequence:

1. The client application performs a connection sequence.
2. The client application sends a RESUME TPIPE call with the correct settings in the IRM.

Recommendation: Set the IRM_TIMER value to X'FF', which causes IMS Connect to override the TIMEOUT value in the configuration file and wait forever.

3. The client application sends a TCP/IP READ to receive the output message.
4. The client application sends an acknowledgment (ACK—Set the IRM_TIMER value to the same value you set on the RESUME TPIPE call.) and returns to the TCP/IP READ.

The timer interval that is set in IRM_TIMER is a different timer value from the one that is set in the IMS Connect configuration file (that value is TIMEOUT=).

The IRM_TIMER value is the wait value to wait for a RECEIVE issued from the client following a RESUME TPIPE call, or an ACK to the RECEIVES following the RESUME TPIPE call.

Retrieving output from alternate OTMA tpipe hold queues

Client applications can retrieve the asynchronous output or callout messages from an alternate tpipe hold queue by specifying the name of the alternate tpipe as an alternate client ID a RESUME TPIPE call.

When IMS Connect passes a RESUME TPIPE call that specifies an alternate client ID to OTMA, OTMA returns to the caller any messages that are queued to the tpipe that matches the alternate client ID.

Specifying an alternate client ID is used to retrieve output in the following scenarios:

- By IMS TM resource adapter when client IDs are unknown because they are automatically generated at runtime.

- In Sysplex Distributor environments, where client applications typically do not know on which tpipe hold queue their output is queued.
- In callout environments, where the tpipe name is usually defined in IMS by an OTMA destination descriptor or, for asynchronous callout only, an OTMA routing exit. In this case, the callout messages are retrieved by specifying the tpipe name as an alternate client ID.

By using an alternate client ID, the client application programs can retrieve the output through any instance of IMS Connect by specifying the tpipe name in the alternate client ID field of either the IRM or the OTMA header.

When retrieving output by using an alternate client ID, a /DISPLAY TMEMBER TPIPE command displays the alternate client ID instead of the actual ID of the client that submitted the RESUME TPIPE request. Consequently, identifying which client submitted a particular RESUME TPIPE request can be a challenge if you need to diagnose a problem.

- User-provided client applications can retrieve asynchronous output from an alternate tpipe by specifying the following in the IRM of the RESUME TPIPE call:

X'01' in the IRM_ARCH field of the fixed IRM section

The name of the target tpipe in the IRM_RT_ALTCID field of the user-defined section of the IRM

- For IMS TM resource adapter client applications, to retrieve asynchronous from an alternate tpipe the client application must specify the following values in the OTMA header:

The name of the target tpipe in the OMUSR_RT_ALTCID field of the OTMA header

OMUSR_AL02 in the OMUSR_ARCLEV field

Defining groups for shared asynchronous output

You can define groups of data store connections, IMS Connect instances, or both, in which the members of the group can retrieve the asynchronous output for any other member of the group. These groups use the OTMA super member function.

When a super member group is not defined, asynchronous output can be retrieved only by using the client ID of the originating client and the originating IMS Connect instance. The dependency on the original client ID and IMS Connect instance can present problems if, for example, either the originating client ID is unknown, as might be the case when the ID is automatically generated, or the original instance of IMS Connect is unknown, as might be the case when a product like the z/OS Sysplex Distributor is used.

Even when the client ID and IMS Connect instance are known, there is still a risk presented by the IMS Connect instance being a single point of failure; if the IMS Connect instance should fail, the asynchronous output could not be retrieved until the IMS Connect instance was brought back up.

You create a super member group by defining IMS Connect instances and individual data store connections as participants of the super member group. Each participant within the same super member group must specify the same super member name on the SMEMBER parameter in the IMS Connect configuration member in the IMS.PROCLIB data set. For IMS Connect, the SMEMBER parameter is on the HWS statement. For data store connections, the SMEMBER parameter is on the DATASTORE statement.

By default, the data store connections defined to an instance of IMS Connect are participants in the super member group that the instance is a participant in, if any; however, by specifying the SMEMBER parameter on the DATASTORE statement, a data store connection can belong to a different super member group or to no super member group at all. Specifications for super member groups made on the DATASTORE statement, override any super member group specification made on the HWS statement.

Note that if a DATASTORE statement does include the SMEMBER parameter but the HWS statement does, when the attributes of the data store connection are displayed they will not include the super member group name that, by default, the data store is a participant in. You must display the attributes of the HWS statement to see the super member group name.

- To define an IMS Connect instance as a participant in a super member group, specify a super member name on the SMEMBER parameter on the HWS configuration statement. The super member name specified on the HWS configuration statement becomes the default super member name for all data store connections defined to the instance of IMS Connect that do not specify a super member group of their own.
- To define a data store connection as a participant in a super member group other than the super member group specified on the HWS configuration statement, specify a super member name on the SMEMBER parameter on the DATASTORE configuration statement. A super member name specified on the DATASTORE statement overrides for this data store connection any super member name that is defined on the HWS configuration statement
- To prevent a data store connection from participating in a default super member group defined by an IMS Connect instance, specify SMEMBER='####' on the DATASTORE configuration statement that defines the data store connection.

Example

The following example of an IMS Connect configuration statement includes specifications for super member groups. The SMEMBER parameter on the HWS configuration statement defines the IMS Connect instance HWS1 as participating in the super member group SHRD and sets SHRD as the default super member group for all data store connections defined in this configuration member.

In the example, the data stores defined to the IMS Connect instance HWS1 each participate in super member groups as follows:

- IMS1 participates in the SHRD super member group defined on the HWS configuration statement because its DATASTORE statement does not include the SMEMBER parameter.
- IMSY and IMSN both participate in the super member group SHR1
- IMSX participates in super member SHR2, possibly with other data store connections from another instance of IMS Connect
- IMSA does not participate in any super member group

```
HWS=(ID=HWS1,XIBAREA=100,RACF=Y,SMEMBER=SHRD)
TCPIP=(HOSTNAME=TCPIP,PORTID=(9999,LOCAL),RACFID=GOFISHIN,TIMEOUT=0,
IPV6=N,SSLPORT=(9998),SLENVAR=HWSCFSSL,
EXIT=(HWSSMPL0,HWSSMPL1,HWSCSL00,HWSCSL01))
DATASTORE=(ID=IMS1,GROUP=XCFGRP1,MEMBER=HWS1,TMEMBER=IMS1)
DATASTORE=(ID=IMSY,GROUP=XCFGRP1,MEMBER=HWSY,TMEMBER=IMS1,SMEMBER=SHR1)
DATASTORE=(ID=IMSN,GROUP=XCFGRP1,MEMBER=HWSN,TMEMBER=IMS1,SMEMBER=SHR1)
```



```

DATASTORE=(ID=IMSX,GROUP=XCFGRP1,MEMBER=HWSX,TMEMBER=IMS1,SMEMBER=SHR2)
DATASTORE=(ID=IMSA,GROUP=XCFGRP1,MEMBER=HWSA,TMEMBER=IMSA,SMEMBER=####)
IMSPLEX=(MEMBER=HWSPLEX1,TMEMBER=PLEX1)
IMSPLEX=(MEMBER=HWSPLEX2,TMEMBER=PLEX2)

```

Related tasks:

“Sharing asynchronous commit-then-send output: the OTMA super member function” on page 843

Related reference:

 HWSCFGxx member of the IMS PROCLIB data set (System Definition)

Asynchronous output message flow

Implementing asynchronous output support forces a commit-then-send (commit mode 0) message flow. This flow requires an acknowledgment (ACK/NAK) from the client.

If an IMS transaction running in commit-then-send message flow sends a message to the client, and that message cannot be delivered, OTMA will react as though a NAK had been sent to OTMA from IMS Connect, and the message will be placed on the OTMA Hold Queue. OTMA will behave in this manner for whatever reason that the NAK gets sent (for example, because the z/OS cross-system coupling facility connection is not available, because IMS Connect has terminated, or because IMS Connect has lost communications with TCP/IP).

Related tasks:

“Purging undeliverable commit-then-send output” on page 299

Related reference:

“Output message from message exit to client” on page 257

Chapter 18, “OTMA header fields used by IMS Connect,” on page 265

IMS Connect client call flows

The following examples show flows for IMS Connect client conversational and non-conversational transactions.

All sample flows shown apply to both persistent and transaction TCP/IP sockets, and all flows use this protocol: commit mode 1 (send-then-commit), synch level = confirm, with ACK and NAK.

The following sample flows are illustrated:

- Non-conversational, running to successful completion using ACK
- Conversational, running to successful completion using ACKs
- Non-conversational, where client sends NAK in response to message
- Conversational, where client sends NAK in response to one of the messages
- Non-conversational, terminated by Host application before successful completion of transaction
- Conversation terminated by Host application before successful completion of transaction

Non-conversational, running to successful completion using ACK

The following example shows a non-conversational flow with commit mode=1, synch level=confirm, and ACK (transaction runs to successful completion).

CLIENT REQUEST	FLOW	IMS CONNECT REQUEST
SEND----->	IRM/TRAN/DATA	----->RECEIVE
RECEIVE<-----	DATA/CSM<-----	SEND
SEND----->	IRM/ACK-----	----->RECEIVE
RECEIVE<-----	RSM<-----	SEND DEALLOCATE CONFIRM
	RSM reason code = DEALLOCATE CONFIRM X'61' (97)	
	(97 = IMS Host application has committed the transaction)	

Conversational, running to successful completion using ACKs

The following example shows a conversational flow with commit mode=1, synch level=confirm, and ACK (transaction runs to successful completion).

CLIENT REQUEST	FLOW	IMS CONNECT REQUEST
SEND----->	IRM/TRAN/DATA	----->RECEIVE
RECEIVE<-----	DATA/CSM<-----	SEND
SEND----->	IRM/ACK-----	----->RECEIVE
SEND----->	IRM/DATA-----	----->RECEIVE
RECEIVE<-----	DATA/CSM<-----	SEND
SEND----->	IRM/ACK-----	----->RECEIVE
	.	
	.	
SEND----->	IRM/DATA	----->RECEIVE
RECEIVE<-----	DATA/CSM<-----	SEND
SEND----->	IRM/ACK-----	----->RECEIVE
RECEIVE<-----	RSM<-----	SEND DEALLOCATE CONFIRM
	RSM reason code = DEALLOCATE CONFIRM X'61' (97)	
	(97 = IMS Host application has committed the transaction)	

Non-conversational, where client sends NAK in response to message

The following example shows a non-conversational flow with commit mode=1, synch level=confirm, and NAK (transaction terminates with a NAK from client application).

CLIENT REQUEST	FLOW	IMS CONNECT REQUEST
SEND----->	IRM/TRAN/DATA	----->RECEIVE
RECEIVE<-----	DATA/CSM<-----	SEND
SEND----->	IRM/NAK-----	----->RECEIVE
RECEIVE<-----	IMS MESSAGE "DFS555.."	

Conversational, where client sends NAK in response to one of the messages

The following example shows a conversational flow with commit mode=1, synch level=confirm, and NAK (transaction terminates with a NAK from client application).

CLIENT REQUEST	FLOW	IMS CONNECT REQUEST
SEND----->	IRM/TRAN/DATA	----->RECEIVE
RECEIVE<-----	DATA/CSM<-----	SEND


```

SEND----->IRM/ACK----->RECEIVE
SEND----->IRM/DATA----->RECEIVE
RECEIVE<-----DATA/CSM<-----SEND
SEND----->IRM/ACK----->RECEIVE
.
.
SEND----->IRM/DATA----->RECEIVE
RECEIVE<-----DATA/CSM<-----SEND
SEND----->IRM/NAK----->RECEIVE

RECEIVE<----- IMS MESSAGE "DFS555.."

```

Non-conversational, terminated by Host application before successful completion of transaction

The following example shows a non-conversational flow with commit mode=1, synch level=confirm, and ACK (transaction terminated by host application before successful completion).

```

CLIENT          FLOW          IMS CONNECT
REQUEST

SEND----->IRM/TRAN/DATA ----->RECEIVE
      Host Application abnormally terminates
RECEIVE<----- IMS MESSAGE "DFS555.."

```

Conversation terminated by Host application before successful completion of transaction

The following example shows a conversational flow with commit mode=1, synch level=confirm, and NAK (transaction terminated by host application).

```

CLIENT          FLOW          IMS CONNECT
REQUEST

SEND----->IRM/TRAN/DATA ----->RECEIVE
RECEIVE<-----DATA/CSM<-----SEND
SEND----->IRM/ACK----->RECEIVE

SEND----->IRM/DATA----->RECEIVE
RECEIVE<-----DATA/CSM<-----SEND
SEND----->IRM/ACK----->RECEIVE
.
.
SEND----->IRM/DATA----->RECEIVE
RECEIVE<-----DATA/CSM<-----SEND
SEND----->IRM/ACK----->RECEIVE

      Host Application abnormally terminates

RECEIVE<----- IMS MESSAGE "DFS555.."

```

IMS Connect client message protocol sequence for IMS DFS messages and IMS command output

Table 57 on page 345 and Table 58 on page 345 show the required actions to be taken when different IMS DFSⁿⁿⁿⁿⁿ messages or IMS command output is sent to the IMS Connect client. The two tables illustrate whether or not an ACK is required to be sent when the client receives an IMS DFS message or output from an IMS command, both for synch level Confirm and synch level None and for commit mode 0 (CM0) and commit mode 1 (CM1).

Note: The client code can test the CSM_FLG1 byte for the presence of the CSM_ACK_NAK flag; it can also test the RSM_FLG1 byte for the presence of the RSM_ACK_NAK flag. It performs this test to determine if an ACK or NAK is required. Otherwise, it performs the analysis outlined in Table 57 and Table 58. Table 57 and Table 58 also define whether or not the client requires a READ in order to receive the "Deallocate Abort" response (RSM) from IMS Connect. Notes for both tables immediately follow Table 58.

Table 57. Persistent sockets: client message protocol sequence for IMS DFS messages and IMS command output

Message output to client	CM1, Synch level confirm	CM1, Synch level none	CM0, Synch level confirm	CM0, Synch level none
Invalid transaction code DFS064	DFS064 ¹	DFS064 ¹	N/A	N/A
Transaction stopped DFS065	DFS065 ¹	DFS065 ¹	N/A	N/A
Transaction abended DFS555	DFS555 ⁷	DFS555 ¹	N/A	N/A
Output DFS2082	DFS2082 ²	DFS2082 ¹	N/A	N/A
IMS Command Output	Cmd output ¹	Cmd output ¹	N/A	N/A
Security Failure DFS1292	DFS1292 ¹	DFS1292 ¹	N/A	N/A
Segment greater than 32 K	DFS1294 ⁵	DFS1294 ⁵	N/A	N/A

Table 58. Transaction sockets: client message protocol sequence for IMS DFS messages and IMS command output

Message output to client	CM1, Synch level confirm	CM1, Synch level none	CM0, Synch level confirm	CM0, Synch level none
Invalid transaction code DFS064	DFS064 ¹	DFS064 ¹	DFS064 ¹	N/A
Transaction stopped DFS065	DFS065 ¹	DFS065 ¹	DFS065 ¹	N/A
Transaction abended DFS555	DFS555 ⁷	DFS555 ¹	DFS555 ⁷	N/A
Output DFS2082	DFS2082 ²	DFS2082 ¹	No output ³	N/A
IMS Command Output	Cmd output ¹	Cmd output ¹	Cmd output ⁴	N/A
Security Failure DFS1292	DFS1292 ¹	DFS1292 ¹	DFS1292 ¹	N/A
Segment greater than 32 K	DFS1294 ⁵	DFS1294 ⁵	DFS1297 ⁶	N/A

Notes:

1. Does not require an ACK to DFS messages.
2. Requires both an ACK to DFS messages and a second read to get a deallocate response.

3. The read to receive the transaction output will time out. No data will be received. OTMA treats commit mode=0 and Synch level=Confirm as asynchronous output. If the IMS Host application does not return a message (ISRT to IOPCB), OTMA does not send a deallocate. The TIMEOUT= value specified in the IMS Connect configuration file will have to expire before the disconnect is complete.
4. Requires an ACK to command output. A second read is not required to get a deallocate response. The command output gets treated as asynchronous output.
5. Does not require ACK to DFS1294 output. A second receive is required to receive the DFS555 message.
6. Client will receive DFS1297 rather than DFS1294. The DFS1294 message does not require an ACK. No DFS555 message gets sent, so a second receive is not required. The application is committed, and the application output gets discarded because the segment is larger than 32 K.
7. Does not require an ACK to DFS messages.

Reason codes for commit mode=1, synch level=confirm

For CM1, there are three reason codes associated with a zero (0) return code, and two reason codes associated with an X'04' return code, which provide information to the client application. The sample flows illustrate how each of these codes are used. The code meanings are listed in the following table.

Table 59. Information reason codes for CM1, synch level=confirm

Return code	Reason code	Description
X'00'	94	Response - only output from host from non-conversation
X'00'	95	Conversation - last output from host from on conversation
X'00'	96	Conversation/response - middle of conversation
X'04'	97	Deallocate commit - successful completion of host application
X'04'	98	Deallocate abort - abnormal termination of host application

IMS Connect dead letter queue (HWS\$DLQ)

In certain instances, if OTMA receives a NAK response from IMS Connect, OTMA stores the undelivered message on the IMS Connect dead letter queue. The IMS Connect dead letter queue is identified by the tpipe name HWS\$DLQ.

The instances in which OTMA stores messages on HWS\$DLQ include:

- When IMS Connect returns a NAK response to OTMA for a message that is missing the user data section of the OTMA header. In this case, IMS issues message HWSD0255W.
- When IMS Connect returns a NAK response to OTMA because IMS Connect could not process an asynchronous callout request. In this case, IMS issues message HWSP1510E.

Except in the case where a message has no user data section, you can retrieve messages on HWS\$DLQ by specifying HWS\$DLQ as the alternate client ID on the

RESUME TPIPE call. Messages on HWS\$DLQ that are missing the user data section of the OTMA header must be dequeued before subsequent messages on the queue can be retrieved.

To dequeue a message from HWS\$DLQ:

1. Stop HWS\$DLQ by issuing the command `/STOP TMBER tmembername TPIPE hws$dlq`.
2. Dequeue the message by issuing the command `/DEQUEUE TMBER tmembername TPIPE hws$dlq PURGE1`.
3. Start HWS\$DLQ by issuing the command `/START TMBER tmembername TPIPE hws$dlq`.

To view the queue counts for HWS\$DLQ issue the command `/DISPLAY TMBER tmembername TPIPE hws$dlq`.

Ping support for IMS Connect

To determine whether or not IMS Connect is available, you can send a ping request to IMS Connect. The ping support operates like a transaction and has the appearance of a transaction code and data. When you send the request PING IMS_CONNECT, the response is PING RESPONSE.

The user message exits, HWSJAVA, HWSSMPL1, HWSSMPL0, and HWSSOAP1 provide ping support. User message exits HWSCSLO0 and HWSCSLO1 do not support the ping function. If you write your own user message exit, you can choose to add the ping function support in your exit.

The communication sequence of the ping request consists of four steps. If you are using HWSSOAP1 user message exit, the third and the fourth steps, receive and disconnect, differ from the receive and disconnect steps for the HWSJAVA0, HWSSMPL0, and HWSSMPL1 user message exits.

For the HWSJAVA0, HWSSMPL0, and HWSSMPL1 user message exits, the communication sequence of the ping request is:

1. Connect.
2. Send PING IMS_CONNECT (must be sent in uppercase).
3. Receive:
 - If you are using HWSSMPL0 or HWSSMPL1, the user-written client application receives `WSC0030I *PING RESPONSE* *CSMOKY*`.
 - If you are using HWSJAVA0, the IMS TM Resource Adapter application receives `WSC0030I *PING RESPONSE*` in the application data portion of the OTMA header and return code of 48 and reason code ICONSUCC in the user data portion of the OTMA header.
4. Resolve socket connection.

IMS Connect maintains the socket connection, except in the following circumstances:

- If the socket is not persistent, IMS Connect disconnects the socket.
- If IMS Connect expects an ACK or NAK response from the client application, IMS Connect disconnects the socket.
- If the client application is in IMS conversational mode, IMS Connect disconnects the socket.

For the HWSSOAP1 user message exit, the communication sequence of the ping request is:

1. Connect.
2. Send PING IMS_CONNECT (must be sent in uppercase).
3. Receive PING_RESPONSE.
4. Disconnect.

Related reference:

 [IMS TM Resource Adapter user message exit routine \(HWSJAVA0\) \(Exit Routines\)](#)

 [User message exit routines HWSSMPL0 and HWSSMPL1 \(Exit Routines\)](#)

 [SOAP Gateway exit routine \(HWSSOAP1\) \(Exit Routines\)](#)

Chapter 20. IMS Connect two-phase commit support

IMS Connect supports two phase-commit for both IMS TM transactions and IMS DB database requests.

The IMS TM transactions and IMS DB database requests that are participants in two-phase-commit transactions are coordinated by z/OS Resource Recovery Services or an external coordinator (for example, IBM WebSphere Application Server).

When accessing IMS TM through IMS Connect, the external coordinator must use the IMS TM Resource Adapter as the resource adapter. Together, the IMS TM Resource Adapter and IMS Connect handle the data flow for two-phase-commit processing.

When accessing IMS DB through IMS Connect, the external coordinator can use any of the following resource adapters and APIs:

- IMS Universal Database resource adapter
- IMS Universal JDBC driver
- IMS Universal DL/I driver
- The Distributed Relational Database Architecture (DRDA)

To handle the data flow for two-phase-commit processing for IMS DB support, IMS Connect works with the CSL Open Database Manager (ODBM) component, together with the IMS Universal drivers.

This topic provides an overview of two-phase commit and some key scenarios that IMS Connect supports.

Overview of two-phase commit protocol

Two-phase commit protocol is comprised of a set of actions that ensure a transaction involving multiple databases does not produce unsynchronized updates.

Two-phase commit provides a way for a series of database interactions on multiple different data sources to be grouped together and completed or rolled back as a single transaction. Two-phase commit transactions that represent a series of database interactions on multiple data sources are referred to as *global transactions*.

At the beginning of a global transaction, a global transaction ID (XID) is generated and used by an external transaction manager to drive the two-phase commit processing across all of the resource managers involved. Each database interaction within the scope of the global transaction is executed upon its associated resource manager. The results of the interactions are then sent back to the application for processing. When the database interactions within the scope of the global transaction are finished, a prepare call is sent to each resource manager that was accessed by the global transaction. The prepare call provides each resource manager a chance to determine and report on its ability to commit the work it has done as a part of the global transaction. Upon receiving verification that each resource manager can commit its work, the transaction manager sends a commit call to each resource manager.

At any point in time prior to sending the commit call, the two-phase commit transaction can be rolled back. If the transaction is rolled back, a rollback call is sent to each resource manager involved in the transaction and the temporary changes are removed or discarded.

If any database failures occur during the commit phase, the external transaction manager tries to reestablish a connection with the failed resource manager and resumes its call for the resource manager to commit.

If the transaction manager fails during the commit phase, the transaction manager performs recovery processing upon restart and attempts to reestablish a connection with all of the resource managers involved. When the transaction managers reestablishes the connections, it resumes its call for the resource managers to commit.

Distributed two-phase commit support

Distributed two-phase commit protocol uses TCP/IP to communicate transactions between various platforms (for example, Windows, AIX®, Solaris, Linux).

A distributed TCP/IP transaction normally involves the following components:

- An application component
- An application server, such as WebSphere Application Server
- A resource adapter, such as one of the IMS Universal Database resource adapters, or the IMS TM Resource Adapter.
- A resource manager
- A transaction manager, such as IMS TM
- An enterprise information system (EIS), such as IMS DB

In distributed two-phase commit protocol, a client that is deployed on an application server issues a transaction. The application server acts as an external transaction manager (external coordinator) to manage transactions across one or more resource managers. To access the resource manager of an enterprise information system, the external coordinator must use a resource adapter. The IMS Connect client accesses the resource manager (IMS) through IMS Connect using TCP/IP.

IMS supports the X/Open XA protocol through z/OS Resource Recovery Services (RRS), which IMS uses to participate in two-phase-commit processing on z/OS. As a result, IMS Connect communicates with RRS and passes to RRS the transaction context from the IMS Connect client. In turn, RRS as the syncpoint coordinator coordinates the changes so that all or no updates are made to IMS. In RRS, the set of changes that are made or not made within the transaction scope is called a unit of recovery (UR).

IMS Connect plays dual roles in two-phase-commit processing. IMS Connect, acts as an extension to RRS (the syncpoint manager) and is considered the server distributed syncpoint resource manager (SDSRM). As the SDSRM, IMS Connect allows RRS to communicate with other syncpoint managers as needed to ensure coordination of the distributed resources the application accesses. IMS Connect also is the communication resource manager (CRM). As the CRM, IMS Connect controls access to distributed resources by allowing an application component to communicate with other application components and resource managers that may

be on different systems. Also, as the CRM, IMS Connect assists in processing a syncpoint event and communicates the events to distributed syncpoint managers.

IMS distributed transaction processing can be broken down into two types of transactions: global transactions that use two-phase commit protocol and transactions that use the one-phase commit optimization. The one-phase commit optimization is used when only one resource manager is involved in the transaction.

Support for the IMS Universal drivers

IMS Connect supports two phase-commit and one-phase commit for IMS DB database requests that are received from IMS Universal drivers.

Global (XA) transactions with the IMS Universal drivers

IMS support for global, two-phase commit transactions for the IMS Universal drivers uses z/OS Resource Recovery Services (RRS) multisystem cascaded transactions to coordinate database access requests between IMS Connect, the CSL Open Database Manager (ODBM), and IMS within an IMSplex.

The IMS Universal drivers provide the LocalTransaction and XAResources interfaces for communication with IMS Connect. IMS Connect maintains the XID and associates it with the work context token and the parent unit of recovery (UR), which represents the transaction as submitted by the IMS Universal drivers.

When IMS Connect receives a database access request from any of the IMS Universal drivers that requires two-phase commit, IMS Connect calls RRS and becomes the *coordinator* of the multisystem cascaded transaction. IMS Connect creates a parent unit of recovery (UR) and passes a parent UR token with the database access request to the CSL Open Database Manager (ODBM).

ODBM becomes a *subordinate* to the multisystem cascaded transaction and uses the parent-UR token received from IMS Connect to create a cascaded child UR.

When the client application requests prepare for commit, the IMS Universal drivers send a prepare signal to IMS Connect. IMS Connect then issues a prepare request to RRS, and waits for ODBM to complete phase one of the two-phase commit process. If the IMS resource manager is prepared to commit, ODBM returns the prepare to commit confirmation to RRS and RRS sends the results to IMS Connect. IMS Connect then sends the result of the prepare command back through the IMS Universal drivers to the transaction manager. At this point, the UR maintained by IMS Connect and ODBM is in an indoubt state until a commit or rollback command is sent in from the external transaction manager.

If all the resource managers associated with the global transaction can commit, the transaction manager hardens the commit decision and drives the IMS Universal drivers to commit the change. The IMS Universal drivers send a commit signal to IMS Connect and IMS Connect tells RRS that the overall decision is to commit all resources. RRS tells IMS to commit the changes. After IMS commits the changes, RRS then returns to IMS Connect with the information that the local resources have been committed. IMS Connect tells RRS to delete its log records.

The following figure illustrates the first 17 steps of the flow of a distributed two-phase commit global transaction that is submitted to IMS Connect from the IMS Universal drivers. The transaction involves one IMS system. RRS needs to be

active on each z/OS image; however, IMS Connect does not need to be on the same z/OS image as ODBM and IMS.




Figure 33. Distributed two-phase commit global transaction client flow for IMS Universal drivers (1 of 2)

The following figure illustrates the remaining steps of the flow of the distributed two-phase commit global transaction shown in the preceding figure.




Figure 34. Distributed two-phase commit global transaction client flow for IMS Universal drivers (2 of 2)

One-phase commit global transactions with the IMS Universal drivers

If the external transaction manager detects that only one resource manager is enlisted in a given transaction, the transaction manager can perform one-phase-commit optimization.

| In the one-phase-commit optimization, the first phase of the two-phase commit
| operation, the prepare phase, is omitted and the transaction manager sends only
| the *commit request* to the resource manager to commit the changes without first
| sending *prepare to commit*.

| z/OS Resource Recovery Services (RRS) transactions that use one-phase commit
| are also referred to as an RRS *local transaction*. Local transactions do not contain an
| XID.

| When IMS Connect receives from the IMS Universal drivers a database access
| request that is an RRS local transaction, IMS Connect does not call RRS and passes
| the local transaction to ODBM directly. ODBM in turn calls RRS and coordinates
| the commit or rollback of the database request.

| The following figure illustrates the flow for a distributed one-phase commit global
| transaction. RRS needs to be active on each z/OS image; however, IMS Connect
| does not need to be on the same z/OS image as ODBM and IMS.

| The following figure illustrates the first 17 steps of the flow of a distributed
| one-phase commit global transaction that is submitted to IMS Connect from the
| IMS Universal drivers. The transaction involves one IMS system. RRS needs to be
| active on all z/OS images; however, IMS Connect does not need to be on the same
| z/OS image as ODBM and IMS.




Figure 35. Distributed one-phase commit global transaction client flow for IMS Universal drivers (1 of 2)

The following figure illustrates the remaining steps of the flow of the distributed one-phase commit global transaction shown in the preceding figure.




Figure 36. Distributed one-phase commit global transaction client flow for IMS Universal drivers (2 of 2)

Support for the IMS TM Resource Adapter

IMS Connect supports two phase-commit and one-phase commit for IMS TM transactions that are received from the IMS TM Resource Adapter.

Global (XA) transactions with IMS TM Resource Adapter

A global (XA) transaction is controlled and coordinated by an external transaction manager (external coordinator) to a resource manager.

The transaction normally requires coordination across multiple resource managers that may reside on different platforms. The transaction must be sent to IMS Connect as send-then-commit (commit-mode 1) with a sync level of syncpoint.

To access an enterprise information system, the external coordinator sends an XID, which is defined by the X/Open XA standard, to a resource adapter. In addition to the length and FormatID fields, an XID has two other parts: the global transaction identifier (GTRID) and the branch qualifier (BQUAL). Because IMS does not support X/Open XA protocol, the IMS TM Resource Adapter uses the LocalTransaction and XAResources interfaces to participate in transactions coordinated by the external coordinator to communicate with IMS Connect. IMS Connect maintains the XID and associates it with a work context token and an IMS name. IMS Connect then passes the context token to z/OS Resource Recovery Services (RRS).

IMS Connect sends the transaction output back to the IMS TM Resource Adapter which returns the output data to the client. Upon sending the output message to the IMS TM Resource Adapter successfully, IMS Connect sends an ACK to IMS to acknowledge the message. After making requests to IMS, the application component indicates to the IMS TM Resource Adapter that it is ready to commit the changes. At this point the IMS TM Resource Adapter sends a prepare signal to IMS Connect. IMS Connect, in turn, tells RRS to initiate the prepare phase. If the IMS resource manager is prepared to commit, RRS collects the prepare to commit confirmation from the resource manager and sends the results to IMS Connect. IMS Connect will then send a *request to commit* signal to the IMS TM Resource Adapter to request committing the changes.

When the IMS TM Resource Adapter receives the request to commit signal, it tells the external coordinator that the resources on the IMS system can be committed. The transaction manager determines the overall results. If all the resource managers can commit, the transaction manager hardens the commit decision and will drive the IMS TM Resource Adapter to commit the change. The IMS TM Resource Adapter sends a commit signal to IMS Connect and IMS Connect tells RRS that the overall decision is to commit all resources. RRS tells IMS to commit the changes. After IMS commits the changes, RRS then returns to IMS Connect with the information that the local resources have been committed. IMS Connect tells RRS to delete its log records.

Requirements:

IMS Connect, RRS, and IMS must all be on the same z/OS image (LPAR). Also, all message flows of a global transaction, which can include multiple separate transactions submitted to IMS from the same global transaction, must be processed by the same IMS and IMS Connect pair. Because of this requirement, you might not be able to send global transactions to IMS if you use software that distributes workload across multiple instances of IMS Connect, such as the z/OS Sysplex Distributor. Such workload distribution software is not likely to guarantee that all flows from the same global transaction will be routed to the same IMS and IMS Connect pair.

The following series of figures illustrate the flow of distributed two-phase commit global transactions between IMS TM Resource Adapter and IMS in different configuration scenarios.

The following figure shows the flow of a single two-phase commit transaction between IMS TM Resource Adapter and a single IMS system. In the figure, IMS Connect and IMS are running on the same LPAR.




Figure 37. Flow of a single distributed two-phase commit global transaction

The following figure illustrates the flow of a distributed two-phase commit global transaction. The transaction involves two IMS systems.




Figure 38. Distributed two-phase commit global transaction client flow

One-phase commit global transactions with IMS TM Resource Adapter

If only one resource manager is registered in a transaction that is making changes to shared resources, the transaction manager can perform one-phase-commit optimization. An external coordinator is not required.

The transaction manager can send the phase two *commit request* directly to the resource manager to commit the changes. IMS Connect does not have to go through phase one, *prepare to commit* of the two-phase commit protocol and can go directly to phase two, *commit request*.

Requirement:

IMS Connect, RRS, and IMS must all be on the same z/OS image.

The following figure illustrates the flow for a distributed one-phase commit global transaction.




Figure 39. Distributed one-phase commit optimization client flow

Two-phase commit support for IMS TM Resource Adapter Local Option

IMS Connect also supports two-phase commit through Local Option.

To enable two-phase commit flow, two fields in the OTMA prefix must be set correctly before the transaction is sent to IMS Connect. First, the OMHDRSYN flag must be set to 'X'02' (for example, OMHDRSL2) to indicate SyncLevel Syncpt. Second, the OMHDCID field must be set to a 16-byte z/OS Resource Recovery Services (RRS) context token. Then the ensuing flow is as follows:

1. When IMS Connect detects that the incoming transaction contains the context token and SyncLevel Syncpt, it calls RRS to switch the context off and pass the transaction to IMS.
2. IMS receives the transaction from IMS Connect and schedules the transaction for processing. Because the context token and SyncLevel Syncpt are set, the transaction output will be delivered first before sync point processing is started. When the client receives the transaction output, it is expected to send an acknowledgment to IMS Connect.
3. When the client decides to start two-phase commit processing after it has sent the positive acknowledgment to IMS Connect, it needs to inform RRS of its decision. When RRS is informed of the decision, IMS will be driven for its own

two-phase commit processing. As a result, the transaction will commit or backout based on the final decision from RRS.

Restriction: When IMS Connect and IMS are running on different z/OS images (LPARs), the local option does not support two-phase commit processing. The IBM WebSphere Application Server for z/OS, IMS Connect, RRS, and IMS all must be on the same z/OS image.

The following figure illustrates the two-phase commit flow using the local option provided by IMS Connect.




Figure 40. Two-phase commit flow for local option

Chapter 21. Unicode considerations for IMS Connect

IMS Connect support for Unicode allows Unicode data to be sent to and from an IMS Connect client application. In certain circumstances, IMS Connect translates the Unicode data.

The IMS Connect support for Unicode requires that the client application and the IMS host application both support:

- Unicode data
- The same Unicode encoding schema (either UTF8, UTF16, or UCS-2) as the structure and content of the message being sent and received

IMS Connect supports ASCII and EBCDIC data streams both to and from the client application. If the client application sends ASCII data to IMS Connect, the ASCII is translated to EBCDIC. The subsequent output from IMS Connect to the client application is translated back to ASCII from EBCDIC. If the client application sends EBCDIC data to IMS Connect, no translation is required. With Unicode support, an IMS client application can also send and receive Unicode data to and from IMS Connect, specifically UTF8, UTF16, or UCS-2 data streams.

IMS Connect supports language groups 1, 2, and 3.

The client application uses the IMS Request Message (IRM) to:

- Tell IMS Connect if the data it is sending is Unicode and if the IMS transaction code is being sent as Unicode. IMS Connect transforms the transaction code if it is being sent as Unicode and then sends the transformed code and Unicode data to IMS.
- Tell IMS Connect the Unicode encoding schema that is being used (UTF8, UTF16, or UCS-2).

The transaction code can be sent as Unicode, ASCII, or EBCDIC; however, it must be a valid IMS transaction code of up to 8 bytes that occupies an 8-byte field. In the field, the code must be left justified and, if it is shorter than 8 bytes, padded with blanks. If a blank follows the 8-byte transaction code field, it is considered to be part of the Unicode data.

Message translation

All IMS error messages (for example, DFS555) are sent as either ASCII or EBCDIC. The client application uses the IRM_ID field of the IMS request message (IRM) header to tell IMS Connect which type to send. IMS Connect does not transform messages to Unicode.

For example, if IRM_ID is EBCDIC, the IMS error message (DFSnnnn) is sent as EBCDIC; if IRM_ID is ASCII, the IMS error message (DFSnnnn) is translated from EBCDIC to ASCII.

IRM_ID also identifies the code type of the OTMA header.

An IMS client application can send the IMS transaction code as ASCII, EBCDIC, or Unicode. When the IMS client application sends the transaction code as Unicode, the IMS Connect user message exit (HWSSMPL0 and HWSSMPL1) translates the

transaction code from Unicode TO EBCDIC. When the client application sends the transaction code as ASCII and the remaining data as Unicode, only the transaction code is translated to EBCDIC. A valid 8-byte IMS transaction code can be constructed from the following characters and must begin with an alphabetic character:

- A through Z (uppercase only)
- 0 through 9
- Special characters #, \$, @

An IMS host application that supports Unicode must define an 8-byte field in the input message definition to contain the transaction code. If you pad the 8-byte field with a blank, it is sent as an EBCDIC blank.

If the client application sends Unicode data, the output message is not transformed and is treated as Unicode. For RESUME TPIPE calls, the client application must specify in the IRM if the output should be treated as Unicode or not. During message switching, the IMS host application must ensure that the output message is formatted correctly (using a specific Unicode schema or EBCDIC) for its destination.

Input message format sent by the client

The following table contrasts the message structure for input messages sent by the client and defines the valid ASCII, EBCDIC, and UNICODE formats.

Table 60. Input message structure - message sent by client

EBCDIC IRM	ASCII IRM	If OTMA headers are passed by client	Transaction code	Data
Y	N/A	EBCDIC	EBCDIC	UNICODE
Y	N/A	EBCDIC	UNICODE	UNICODE
N/A	Y	ASCII	ASCII	UNICODE
N/A	Y	ASCII	UNICODE	UNICODE

Output message format received by the client

The following table defines the valid output message elements when the client sends UNICODE data.

Table 61. Output message structure - message received by client

If input message was EBCDIC IRM	If input message was ASCII IRM	RMM	RSM	Output CSM	Output data
Y	N/A	EBCDIC	EBCDIC	EBCDIC	UNICODE
N/A	Y	ASCII	ASCII	ASCII	UNICODE

Related reference:

“Message structures and IMS Connect user message exit routines” on page 250

Chapter 22. TCP/IP settings for IMS Connect

You can choose different TCP/IP values to maximize your environment settings for IMS Connect.

The following TCP/IP parameter values affect IMS Connect:

TCPNODELAY=ENABLE

- Data is transmitted by TCP/IP per client SEND.
- TCP/IP waits one millisecond per transmission.
- Multiple client TCP/IP SENDS can result in multiple TCP/IP transmissions.

TCPNODELAY=DISABLE

- Data is collected by TCP/IP from client TCP/IP SENDS, before transmission.
- TCP/IP waits until the buffer is full before transmission.
- Multiple client SENDS results in 1 to *n* TCP/IP transmissions to IMS Connect.

SO_LINGER=Y, VALUE=0

- Immediate return to client code.
- A client request to close the socket can bypass data sent with a previous client TCP/IP SEND request, but may result in the loss of the client SEND data.

SO_LINGER=N

- Immediate return to client code.
- A client request to close the socket can bypass data sent with a previous client TCP/IP SEND request, but may result in the loss of the client SEND data.

SO_LINGER=Y, VALUE=10

- Return to client code when an ACK is received from the host, or wait for 10 seconds before sending close.
- Socket close will not bypass data sent.

DELAYACK

DELAYACK is used to minimize non-data transmissions from the host. If DELAYACK is used, the z/OS TCP/IP waits 200 milliseconds before sending an ACK to the remote server TCP/IP. However, if the ACK is appended to the data being sent from IMS Connect, there is no delay.

If your client application performs a single SEND followed by a READ, DELAYACK is recommended.

DELAYACK can be set on the TCP/IP "Port Statement" or on the "Gateway Statement."

NODELAYACK

NODELAYACK is used to allow non-data transmissions from the host to flow without data. If NODELAYACK is used, the z/OS TCP/IP immediately sends an ACK to the remote server TCP/IP. The ACK is not appended to the data being sent from IMS Connect.

If the client code sends one SEND followed by a READ to the host with a NODELAYACK setting, an ACK is sent separately.

If the client code sends two or more SENDs followed by a READ to the host, the host TCP/IP will send an ACK immediately to the data received. This will allow the next SEND of data from the client to flow.

NODELAYACK is recommended if your client application sends more than one SEND followed by a READ.

NODELAYACK can be set on the TCP/IP "Port Statement" or on the "Gateway Statement."

Part 6. IMS VTAM network administration

This describes information that will help you to administer your IMS network.

Chapter 23. Introduction to the IMS Transaction Manager network

The following topics introduce the major IMS Transaction Manager (TM) network components, and describes the way communications are established, messages are transmitted, and communications are terminated between IMS and a logical unit.

IMS TM network overview

IMS is a general-purpose database and transaction manager system that provides the necessary support for an advanced telecommunications network. Virtual Telecommunications Access Method (VTAM) controls the physical transmission of data in the network, directing data to IMS from various logical units and from IMS to the appropriate logical units.

An IMS telecommunications network must include the following components:

- IMS
- VTAM (Although IMS uses the network facilities of VTAM, it can also control non-VTAM devices such as Basic Sequential Access Method (BSAM). VTAM is the preferred access method for IMS.)
- Communications hardware (such as control units)
- Terminals

The network can optionally include any of the following components:

- IMS Transaction Manager (IMS TM) services, such as:
 - Extended Terminal Option (ETO)
 - Fast Path
 - Message Format Service (MFS)
 - Intersystem Communication (ISC)
 - Multiple Systems Coupling (MSC)
 - Advanced Program-to-Program Communication (APPC)
- Common Queue Server (CQS) and a coupling facility with any of the following structures:
 - Shared-queues structures
 - Shared-data structures
 - Resource structure
- VTAM generic resource groups
- Open Transaction Manager Access (OTMA)
- Common Service Layer (CSL) including:
 - Operations Manager (OM)
 - Resource Manager (RM)
 - Structured Call Interface (SCI)

In addition, an IMS telecommunications network can operate within one of the following frameworks:

- IBM Systems Network Architecture (SNA), which brings together multiple products in a unified design. SNA formally defines the functional responsibilities of the network components.
- A client-server environment, using Advanced Program to Program Communications (APPC) or OTMA.

Definitions:

- A *logical unit* is an addressable resource that can be an application program, a terminal, or a subsystem such as Customer Information Control System (CICS). A logical unit can also be a component of a general-purpose terminal system that consists of a programmable controller and its attached operator terminals, printers, and auxiliary control units.
- The word *terminal* is used throughout this manual to describe devices and also applies to controllers and remote subsystems. The operator terminals can be keyboard printers, display stations with keyboards, communication terminals, or a mixture of these devices.

A network consisting of IMS and programmable logical units enables users to distribute functions throughout network components. This distribution of function reduces processing requirements that are placed on the central processor (also referred to as the *host*), and it can reduce the impact on the rest of the network when one component encounters a problem.

The following figure illustrates the components of a complete communications network system. The arrows in the figure indicate communications that occur between components. The figure shows the following components:

- IMS and its application programs
- Virtual Telecommunications Access Method (VTAM)
- Tivoli NetView for z/OS
- z/OS operating system (including APPC/MVS if APPC/IMS is used)
- IBM 37x5 Communications Controller and Network Control Program (NCP)
- Terminal




Figure 41. Components of a network

The following list summarizes how each of the components in the preceding figure participates in the network:

IMS

- Checks transaction security
- Schedules the proper application program
- Directs output to the proper terminal
- Provides checkpoint and recovery capabilities

Application program

- Reads data from terminal
- Writes data to processor
- Reads data from processor
- Writes data to terminal

VTAM

- Connects and disconnects terminal from the network
- Sends data from terminal to IMS
- Permits both monitoring and modifying of the network
- Sends data from IMS to the terminal
- Manages physical network (with Tivoli NetView for z/OS)

Communications Controller

- Adds line control characters
- Transmits data
- Receives data
- Removes line control characters
- Checks for transmission errors
- Controls buffering

Network Control Programs (NCP)

- Sends and receives data from communication lines and adapters
- Checks for and records errors

Planning an IMS network requires an understanding of each component and of its relationship to the others.

The following topics provide additional information:

Programmable logical unit (LU)

Definitions:

- A *programmable logical unit (LU)* is an input/output device that is in session with IMS. Application programs in remote logical units can be designed to control more than one terminal.
- When a logical unit informs VTAM that it wants to communicate with IMS, VTAM notifies IMS using the VTAM Logon exit routine. IMS then accepts the request, and VTAM logically connects the logical unit to IMS. This logical connection is called a *session*.

Some of the functions performed by the remote application program include:

- Reading from and writing to associated terminals
- Editing and verifying the data that is received from a terminal
- Reading from and writing to disk storage within the remote LU

- Reading from and writing to the host in which IMS is running
- Editing and verifying data that is received from the host in which IMS is running
- Communicating with other network logical units
- Formatting display and printer devices
- Operating offline when the host, VTAM, IMS, or NCP is unavailable

Depending on the system type, each LU can consist of one or more terminals. An application program that controls an LU consisting of more than one terminal or component must be able to direct output to a specific device. Therefore the application program must be capable of some form of data interrogation in order to make the proper device selection. IMS assists the application program in this process by:

- Allowing LU components to be defined and addressed individually
- Providing, in the header of each output message, the identification of the component to which the message is directed

Communications Controller and Network Control Program (NCP)

VTAM uses the facilities of NCP, which runs in the 37x5 Communications Controllers. VTAM uses NCP to:

- Control lines and devices that are attached to the controllers
- Transmit data between the logical unit and the host CPC
- Perform error recovery
- Collect statistics about the network

Virtual Telecommunications Access Method (VTAM)

VTAM controls the allocation of network resources, and enables these resources to be shared among VTAM users. To VTAM, IMS is a single VTAM user; VTAM is unaware of the IMS application programs.

The IMS application programs use the IMS CALL interface to request IMS services; IMS uses VTAM macros in order to activate VTAM facilities.

If APPC/IMS is active, VTAM regards it as a separate user.

Using VTAM USERVAR with IMS

IMS uses VTAM user variables (USERVARs) in XRF complexes to help maintain and manage sessions when an active IMS subsystem fails. You can, in limited cases, also use VTAM USERVARs to point to a VTAM APPLID of an IMS that is not part of an XRF complex; however, IMS only supports this use of VTAM USERVARs if the IMS being pointed to is the session's primary logical unit (PLU).

For example, you might use a VTAM USERVAR in a non-XRF context if you change the APPLID of an IMS system, but want to temporarily allow LUs to continue connecting to the IMS system using the old APPLID.

Note: In an ISC or MSC environment, a VTAM USERVAR might behave unpredictably when pointing to an IMS subsystem that is not part of an XRF complex. This unpredictable behavior is due to the fact that IMS subsystems can be

either the PLU or the secondary logical unit (SLU) in sessions between two IMS subsystems in ISC and MSC environments.

Related Reading:

- For more information on VTAM and how it is used, see *z/OS Communications Server: SNA Programming*.
- For more information on VTAM USERVARs, see the *z/OS Communications Server: SNA Network Implementation Guide*

IMS product features

IMS comprises two main product features:

- The Database Manager (IMS DB), which can control your databases
- The Transaction Manager (IMS TM), which can control your data communications and application programs

IMS provides:

- Standard functions required by application programs
- An execution environment for concurrently running application programs that serve many online users
- Control of full-function and Fast Path databases

IMS and the application programs that it controls run under z/OS.

Related concepts:

 [Extended Recovery Facility Overview \(System Administration\)](#)

IMS Transaction Manager services

This topic describes several specialized optional Transaction Manager services.

APPC/IMS and LU 6.2 devices

APPC/IMS is a part of IMS TM that allows you to use Common Programming Interface Communications (CPI-C) to build CPI application programs. APPC/IMS supports APPC with facilities provided by APPC/MVS.

IMS supports implicit and explicit application program interfaces (APIs) for APPC support. The implicit API for APPC support allows the IMS application program to communicate with APPC devices using the same techniques employed with other remote devices. This support allows an application program, written by a programmer who has no knowledge of APPC LU 6.2 devices, to:

- Be started from an APPC LU 6.2 device
- Receive input messages from originating LU 6.2 devices
- Send output messages back to originating LU 6.2 devices
- Start transaction programs on remote LU 6.2 devices

The same application program can work with all device types (LU 6.2 and non-LU 6.2) without new or changed coding.

The explicit application programming interface (API) for APPC support is the CPI-C interface, and it is available to any IMS application program. The IMS

application program can issue calls ⁹ to APPC/MVS through this interface.

Fast Path

Fast Path is capable of performing transaction and database processing at high rates. When your system requirements include a high transaction volume, using Fast Path can be advantageous if you do not require all the facilities of full-function processing. Examples of such applications include teller transactions in banking and point-of-sale transactions (inventory update) in retail marketing. Fast Path input and output messages use the expedited message handler (EMH) and bypass message queuing and the priority-scheduling process.

Extended Terminal Option (ETO)

IMS Extended Terminal Option (ETO) provides dynamic terminal and local and remote logical terminal (LTERM) support for IMS TM.

Definition: A *logical terminal (LTERM)* is a user destination. For statically defined terminals, each LTERM is associated with a physical terminal. For ETO terminals, each LTERM is associated with a user, and is associated with a physical terminal only after a user has signed on to a physical terminal.

You can add terminal hardware and LTERMs (users) to the IMS without first defining them. ETO gives you the option of eliminating macro statements in the IMS system definition of VTAM terminals and LTERMs. ETO enhances the availability of your IMS by eliminating the need for you to bring the system down in order to add terminals and LTERMs.

ETO also enhances security for the IMS TM user by associating all output with a specific user, instead of with a device. ETO requires the user to sign on.

ETO reduces the IMS system definition time for those systems where the terminal network is defined dynamically.

Intersystem Communication (ISC)

You can exchange data between your IMS and other external subsystems.

Definition: An *external subsystem* is a subsystem that provides a set of database resources that IMS can use, but does not control. Examples of external subsystems to which your IMS can connect include:

- Customer Information Control System (CICS)
- Other IMS systems
- User-written subsystems

The session that is created between IMS and an external subsystem is called an *application-to-application session*. The IMS that operates in this session uses a function called Intersystem Communication (ISC), which uses SNA protocols.

9. You can also use z/OS ATBxxx call services. For information on these call services, see *z/OS MVS Programming: Writing Transaction Programs for APPC/MVS*.

Message Format Service (MFS)

IMS Message Format Service (MFS) is a facility that formats messages to and from terminals, so that application programs need not deal with device-dependent data in input or output messages. An application program can format messages for different device types using a single set of editing logic, even when device input and output differ from device to device.

IMS MFS formats messages to and from user-written programs in remote controllers and subsystems so that application programs need not deal with the transmission-specific characteristics of the remote controller.

Multiple Systems Coupling (MSC)

Multiple Systems Coupling (MSC) enables you to enter transactions in one IMS and process them in another IMS. The responses from IMS can be returned to the terminals that entered the transactions, or to other terminals.

Related concepts:

“Fast Path expedited message handler” on page 385

“IMS messages and their scheduling” on page 392

“Logical terminals (LTERMs)” on page 391

“Message Format Service” on page 437

Related tasks:

Part 2, “CPI Communications and APPC/IMS,” on page 35

Part 3, “Extended Terminal Option (ETO),” on page 77

Part 7, “Intersystem Communication (ISC),” on page 457

Part 8, “Multiple Systems Coupling (MSC),” on page 675

The Data Communication Control (DCCTL) environment

The Data Communication Control (DCCTL) environment allows you to use the IMS Transaction Manager independently of the IMS Database Manager.

DCCTL provides improved system performance in terms of throughput, system availability, and integrity. DCCTL can coexist with current IMS application programs and installed terminals.

The following scenarios involving DCCTL do not require modifications to your existing environment:

- Application programs that access external database managers can use DCCTL without any modifications. For example, DCCTL can function as a front-end transaction manager for DB2 for z/OS. DB2 for z/OS subsystems do not require modifications in order to run with DCCTL.
- IMS exit routines and IMS application programs that access external subsystem resources in a DCCTL environment do not require modifications.
- Global resource management is the same in a DCCTL environment as it is in a DB/DC environment.

The following procedures might require modifications for a DCCTL environment:

- Operational procedures
- The general procedure that is produced by the system definition

- Application programs that contain a mixture of calls that access both external subsystems and IMS databases do require changes. DL/I calls result in a status code of AD.

Your existing system definition procedures support the generation of a DCCTL system. DCCTL executes with a collection of control blocks that identifies your data processing environment. These control blocks describe the system, data communication, and transaction manager components.

Using a DCCTL environment class system definition, you can generate a TM batch environment. Using TM batch, you can either take advantage of the IMS Batch Terminal Simulator (BTS) or access DB2 for z/OS systems. TM batch support allows only DBB and DL/I regions. It does not provide DL/I database capabilities.

Related reading: For more information on accessing DB2 for z/OS, see *DB2 for z/OS Application Programming and SQL Guide*.

Related concepts:

 DCCTL environment (System Administration)

Operating an IMS network

Operating a basic IMS network involves several tasks.

Operating tasks include:

- Establishing a communication session between a logical unit and IMS
- Sending data between a logical unit and IMS
- Terminating the session between a logical unit and IMS
- Restarting the session after a failure

Operating the network with APPC/IMS

APPC/IMS supports IMS commands for network operation, but the LU 6.2 device handles normal operations such as session startup, transaction initiation, and error handling that does not require master terminal operator (MTO) intervention.

Initiating a session with IMS

A session can be initiated by a logical unit, by the VTAM network operator, by the IMS master terminal operator (MTO), automatically by VTAM, or by IMS itself. After a logical unit connects to IMS, it remains connected until one of the following actions occurs:

- The logical unit itself requests disconnection.
- The IMS MTO requests disconnection.
- Another VTAM application program requests connection to the terminal.
- IMS, VTAM, NCP, the logical unit, or the entire network is stopped.

After the physical connection between the controller and VTAM is established, an LU-to-LU session is initiated. The LU that is requesting the session informs VTAM that it wants to communicate with IMS. VTAM notifies IMS of the request through the VTAM Logon exit routine. IMS indicates that it will accept the request, and VTAM then logically connects the LU to IMS. A session is required before communication between the LU and IMS can be accomplished. To open a session,

all nodes in the communication path (IMS, NCP, line, and station) must be in active status. After a session is established, VTAM directs all data between the logical unit and IMS.

IMS also supports SNA communication links in an Extended Recovery Facility (XRF) complex.

Logging on and signing on to IMS

The following definitions apply to logging onto IMS and signing onto IMS:

Definitions:

- *Logging on* to a terminal establishes a session with IMS for that terminal.
- *Signing on* to a terminal identifies a user to IMS.

Logging off and signing off from IMS

The following definitions apply to logging off from IMS and signing off from IMS:

Definitions:

- *Logging off* from a terminal ends a session with IMS for that terminal.
- *Signing off* from a terminal ends an identification of a user to IMS.

Related concepts:

 [Extended Recovery Facility Overview \(System Administration\)](#)

The shared-queues environment

Operating in a shared-queues environment allows multiple IMS systems in a sysplex environment to share IMS message queues and EMH message queues. The IMS systems work together as an IMSplex providing a single-image view of multiple IMS systems.

The shared-queues environment distributes processing loads between the IMS systems in the IMSplex. Transactions that are entered on one IMS can be made available on the shared queues to any other IMS that can process them. Results of these transactions are then returned to the initiating terminal. End users need not be aware of these activities; they view the processing as if they were operating a single-system.

Definitions:

- A *shared queue* is a collection of messages that are associated by the same queue name. A shared queue is managed by a Common Queue Server (CQS) and can be shared by CQS clients in a IMSplex.
- A *Common Queue Server* receives, maintains, and distributes data objects from a shared queue that resides in a coupling facility list structure for its client.
- A *CQS client* is an IMS DB/DC or DCCTL system that accesses shared queues through its own CQS.
- A *coupling facility* is a special, logical partition that provides high-speed caching, list processing, and locking functions in a sysplex environment.
- A *sysplex environment* is a set of z/OS systems that communicate and cooperate with one another through certain multisystem hardware components and software services in order to process workloads.

- An *IMSplex* is one or more IMS control regions, managers, or servers that work together as a unit. Typically, but not always, IMS systems in an IMSplex:
 - Share either databases or resources or message queues (or any combination)
 - Run in a z/OS Parallel Sysplex® environment
 - Include an IMS Common Service Layer (CSL)

In general, IMS handles messages in the following manner:

1. IMS systems register interest in those queues for which they are able to process messages.
2. When an IMS receives a message and places it on the shared queue, all IMS systems that have registered interest in that queue are notified.
3. One IMS retrieves the message and processes it.
4. The IMS that processes the message places a response on the queue.
5. The IMS that submitted the original message is notified that the response message was placed on the queue.
6. The IMS that submitted the original message sends the response message to the originating terminal.




Figure 42. Basic shared-queues environment

Related reference:

 Queue types (System Administration)

Benefits of using shared queues

The major benefits of operating in a shared-queues environment are automatic workload balancing, incremental growth, and increased reliability.

The following list explains how shared queues provides these benefits:

Automatic workload balancing

A message that is placed on a shared queue can be processed by any participating IMS that is available to process work.

Incremental growth

You can add new IMS systems as workload increases.

Increased reliability

If one IMS fails, work that is placed on a shared queue can still be processed by other IMS systems.

Recommendations:

- Use generic resource groups with shared queues.
- Share all data in an IMSplex across the IMSplex.

Related concepts:

“Balancing sessions with generic resources” on page 381

 Data sharing in IMS environments (System Administration)

Required components of a shared-queues environment

Shared-queues processing requires a number of different components.

Although you can operate many different configurations of a shared-queues environment, the required components of shared-queues processing, shown in Figure 43 on page 380, include:

Common Queue Server (CQS)

One CQS is required for each client, though multiple IMS systems can share a CQS. Each CQS accesses the shared queues, which reside on coupling facility list structures.

CQS client

One or more IMS DB/DC or DCCTL subsystems that can access the shared queues using CQS client requests.

z/OS coupling facility list structures

A type of coupling facility structure that maintains the shared queues.

Definitions:

- A *list structure* is an area of storage in a coupling facility that enables multisystem applications in a sysplex environment to share information that is organized as a set of lists or queues. The list structure consists of a set of lists and an optional lock table.
- CQS maintains list structures in pairs, called *structure pairs*, with a primary list structure and an overflow list structure.
- The *primary list structure* contains the shared queues.
- The *overflow list structure*, if defined, contains shared queues that overflow after the primary list structure reaches a predefined threshold.

z/OS system log

One z/OS system log is used for each structure pair. CQS places recovery information about the work it has processed and about the list structure pair in the z/OS log streams. These log streams are then shared by all CQs that access the list structure pair.

CQS checkpoint data set

One CQS checkpoint data set is maintained for each structure pair of each CQS. The CQS checkpoint data set contains CQS system checkpoint information.

CQS structure recovery data sets (SRDSs)

CQS maintains two SRDSs for each structure pair so that shared queues on the structures can be recovered. The SRDSs maintain structure checkpoint information for the shared queues.




Figure 43. Components of a shared-queues environment

Related concepts:

- ➡ Overview of the Common Queue Server (Communications and Connections)
- ➡ Writing a CQS client (System Programming APIs)

Overview of the Common Queue Server

The Common Queue Server (CQS) is an internal interface by which IMS communicates with shared message queues. You can use IMS commands to initiate CQS requests.

The CQS address space is started by the IMS.

CQS provides the following services:

- Notifies registered clients when work exists on the shared queues.
- Provides clients with an interface for accessing shared queues and CQS.
- Writes CQS system checkpoint information to a CQS checkpoint data set.
- Writes structure checkpoint information to an SRDS for recovery of a shared-queues list structure.
- Provides structure recovery and overflow processing for the shared-queues list structure.

- Drives the following CQS client exit routines:
 - The Client CQS Event exit routine notifies the client of system events, such as CQS abnormal termination and CQS restart completion.
 - The Client Structure Event exit routine notifies the client of structure events, such as structure copy, structure recovery, structure overflow, structure checkpoint, and structure resynchronization.
 - The Client Structure Inform exit routine notifies the client when work exists on the shared queues.
- Drives the following CQS user-supplied exit routines:
 - The CQS Queue Overflow user-supplied exit routine allows the exit to approve or reject a queue that CQS selects for overflow.
 - The CQS Initialization/Termination user-supplied exit routine is notified when CQS initializes and when CQS terminates after disconnecting from all structures (under normal termination conditions).
 - The CQS Client Connection user-supplied exit routine participates in connecting clients to structures and in disconnecting clients from structures.
 - The CQS Structure Statistics user-supplied exit routine gathers structure statistics during CQS system checkpoints.
 - The CQS Structure Event user-supplied exit routine tracks structure activity, such as structure checkpoint, structure rebuild, structure overflow, and structure status change. It also tracks when CQS connects to a structure or disconnects from a structure.
- Provides the Log Print utility with sample JCL. By using the sample JCL, you can print log records from the z/OS log.

Related concepts:

 CQS administration (System Administration)

 Enabling shared queues (System Administration)

Related reference:

 CQS client exit routines (Exit Routines)

Balancing sessions with generic resources

If you are operating in a sysplex environment and have multiple IMS systems, you can initiate a session by using the name of a generic resource group. VTAM balances the sessions among generic resource members in a generic resource group.

If you do not require the services of a specific IMS, initiate the session using a generic resource name, rather than the APPLID name of a specific IMS. VTAM Generic Resources is intended to run in a sysplex environment, but a sysplex environment is not required.

This topic gives an overview of the benefits and terminology of VTAM Generic Resources.

Benefits of generic resource groups

The benefits of using generic resource groups include:

Automatic session workload balancing

Using generic resources is complementary to using shared queues; generic resources distributes network traffic among multiple IMS systems, while shared queues distributes back-end application workload.

Single-image resources

You can use a single generic resource name to access multiple IMS systems, offering a single-system image while you are using the resources of many IMS systems.

Enhanced IMS system availability

In general, if one IMS fails, you can log on to another IMS in that generic resource group.

Exception: You might not be able to logon if the terminal has LOCAL status recovery mode.

Share global messages

You can obtain messages on shared queues from any IMS in the generic resource group.

Recommendation: Before attempting to obtain messages from a terminal that is in either conversation or fast-path response mode, you may need to re-logon to the system from which the input was originally submitted. If the status recovery mode is LOCAL, you must re-logon to the system from which the input was originally submitted. If the status recovery mode is GLOBAL or NONE, you can logon to any IMS within the generic resource group.

Generic resource group definitions

Definitions:

- A *generic resource group* is a set of IMS subsystems that have the same generic resource name, enabling VTAM to distribute terminal sessions among them.
- A *generic resource member* is an IMS subsystem that belongs to a generic resource group.
- A *generic resource name* is the common name by which VTAM knows all the IMS subsystems that belong to a generic resource group.
- An *APPLID name* is a unique application program name by which VTAM knows an IMS subsystem. In a generic resource group, VTAM uses the APPLID name to differentiate each IMS subsystem.
- In the context of generic resources, an *affinity* is an association between a VTAM logical unit and a specific IMS subsystem in a generic resource group.

Generic resource affinity

VTAM establishes an affinity between a terminal and a specific IMS in the following circumstances:

- For all sessions initiated by IMS
- When logging onto IMS using a generic resource name

Until the affinity is reset, subsequent logons resolve to the same IMS.

Affinity is automatically reset in one of two ways:

- By VTAM at session termination if VTAM-managed affinity is established for the session.
- By IMS at session termination if IMS-managed affinity is established for the session, unless the terminal has end-user significant status in the local IMS (status recovery mode of LOCAL).

Related concepts:

IMSpdex terminal management

Using Resource Manager (RM) and a resource structure enhances your ability to manage IMSplex-wide TM resources and to share terminal-related information in a DB/DC or DCCTL environment.

When operating with the Common Service Layer (CSL), specifically RM, the resource structure provides a way to consistently define and maintain resources across the IMSplex. By sharing resource information throughout the IMSplex, you also gain transparency and state recovery for terminals and users.

This topic gives an overview of the benefits of using RM and a resource structure to manage your TM resources.

Note: If you have defined your IMSplex without a resource structure, terminal management is the same as it is for IMS systems running in local mode only.

Related concepts:

➡ CSL RM administration (System Administration)

➡ CSL administration (System Administration)

➡ Planning for Transaction Manager resources in an IMSplex (System Administration)

Related reference:

➡ DFSDCxxx member of the IMS PROCLIB data set (System Definition)

Benefits of managing resources with a resource structure

Using Resource Manager (RM) and a resource structure to manage IMS TM resources provides a number of benefits.

The benefits of using RM and a resource structure include:

Enforcement of resource type consistency

Prevents the same name being used for more than one resource type

Enforcement of resource name uniqueness

Ensures that certain resources can only be active one at a time within the IMSplex

Global callable services

Allows exit routines to obtain LTERM, node, and user resource information across the IMSplex

Terminal and user status recovery

Allows terminals and users to resume work without having affinity to any IMS

VTAM can manage Generic Resource affinities

Specifying GLOBAL as the status recovery mode allows VTAM to manage Generic Resource affinities instead of IMS. If an IMS system fails, VTAM can reassign terminals to another IMS system even if the terminals have affinity to the failed IMS system.

Related concepts:

➡ How RM and the resource structure impact IMS activities (System Administration)

Related reference:

➡ DFSDCxxx member of the IMS PROCLIB data set (System Definition)

Shared TM resources

TM resource sharing is automatic when a resource structure is present in an IMSplex with Resource Manager (RM).

When sharing TM resources the following points apply:

- Each IMS connected to the resource structure can use the TM resources of all the other IMS systems connected to the resource structure.
- RM enforces resource name uniqueness among IMS systems connected to the resource structure.
- RM enforces resource type consistency among IMS systems connected to the resource structure.

If you need to disable TM resource sharing while maintaining all of the connections between the resource structure and the IMS systems in the IMSplex, specify STM=NO in the DFSDCxxx PROCLIB member.

Related reference:

➡ DFSDCxxx member of the IMS PROCLIB data set (System Definition)

Resource name uniqueness

When Resource Manager (RM) is active and a resource structure is defined, IMS automatically enforces name uniqueness for certain TM resources.

The TM resources for which IMS enforces name uniqueness include:

- VTAM LTERMs
- VTAM single-session nodes
- User IDs
- Users

You can disable the enforcement of resource name uniqueness, and TM resource sharing in general, by specifying STM=NO in the DFSDCxxx PROCLIB member.

Related reference:

➡ DFSDCxxx member of the IMS PROCLIB data set (System Definition)

Resource type consistency

IMS automatically enforces type consistency for TM resources when RM is active and a resource structure is defined in the coupling facility.

IMS validates the type of resources for the following message destinations:

- LTERMs
- LTERMs defined as APPC descriptors
- MSNAMEs
- CPI-C transactions
- non-CPI-C transactions

You can disable the enforcement of resource type consistency when a resource structure is present by specifying STM=NO in the DFSDCxxx PROCLIB member. After specifying STM=NO, resource type consistency is enforced only for non-CPI-C and CPI-C transactions.

Related reference:

 DFSDCxxx member of the IMS PROCLIB data set (System Definition)

Fast Path expedited message handler

This topic briefly describes the Fast Path expedited message handler (EMH) and how it processes Fast Path messages.

Fast Path message scheduling

Fast Path schedules input messages by associating them with a load balancing group.

Definition: A *load balancing group* is a group of Fast Path input messages that are ready for balanced processing by one or more copies of a Fast Path program. One load balancing group exists for each unique Fast Path message-driven application program.

The messages for each load balancing group are processed by the same Fast Path program. The EMH controls Fast Path messages by:

- Managing the complete execution of a message on a first-in-first-out basis
- Retaining the messages that are received in the control program's storage without using auxiliary storage or I/O operations
- Supporting multiple copies of programs for parallel scheduling
- Requiring that programs operate in a wait-for-input mode

Routing codes and balancing groups

IMS places input messages on message queues by using the transaction code, and attempts to process one queue until that queue is exhausted, within limits specified by system definition. Message processing is grouped for load balancing and synchronized for database integrity and recovery.

Definitions:

- A *message queue* is a data set on which messages are placed before being either processed by an application program or sent to a terminal.
- A *transaction code* is a one- to eight-character alphanumeric code that invokes an IMS message processing program.
- A *routing code* is a user-defined code that the EMH uses to enable transactions to be routed to programs within a load balancing group. A message is assigned to a balancing group by a routing code. The routing code can be the same as the transaction code, or it can be any other name you choose.

Incoming messages have preassigned routing codes, or a routing code can be dynamically set for the input message. You declare routing codes during system definition, and they are then associated with individual application programs.

Fast Path input edit/routing exit routine (DBFHAGU0)

All Fast Path-exclusive or Fast Path-potential messages are sent to the Fast Path Input Edit/Routing exit routine so that the application program can examine, change, or edit the input message in the input area.

Message buffering with a Fast Path-capable system

IMS buffer management uses a single EMH buffer when the input message has all of the following characteristics:

- Source terminal is VTAM
- Terminal is not enabled for front-end switch (FES)
- Input is single segment
- Input is non-MFS edited
- Input is not an IMS command
- Execution of at least one transaction is through Fast Path

An expedited message handler (EMH) buffer is dynamically allocated to the terminal when IMS receives the first message that has all of the characteristics listed. The EMH buffer remains allocated to the terminal for use with subsequent input messages (that have the same characteristics) until the user signs off, or until the session terminates. The buffer is then released to the EMH pool.

Fast Path EMH and shared queues

In a shared-queues environment that includes Fast Path, you have the option of sharing Fast Path EMH messages using an EMH queue (EMHQ). The EMHQ distributes the processing of EMH messages among multiple IMS systems.

If you do not intend to share EMH messages, you can delete the EMHQ statement from the DFSSQxxx PROCLIB member to prevent IMS from allocating the EMHQ structures and data sets.

Criteria for Fast Path application programs using EMH

Application programs that use EMH must conform to all of the following criteria:

- Each transaction must be initiated by a single-segment, response-type message.
- Each input message must require only a single-segment response message or no response.
- The input and output message length cannot exceed the EMH buffer size.
- ETO and LU 6.2 users cannot access main storage databases (MSDBs) with terminal-related keys.
- No IMS conversational processing is performed.
- Fast Path programs cannot act as automated operator programs or issue IMS commands.
- Fast Path transactions cannot be sent over any of the four types of multiple system coupling (MSC) physical links for processing in another IMS.

Related concepts:

 Enabling shared queues (System Administration)

 Fast Path (System Administration)

Related reference:

✚ Fast Path Input Edit/Routing exit routine (DBFHAGU0) (Exit Routines)

Chapter 24. Planning the network

The following topics describe information that will help you plan your IMS network.

Planning for network administration

As an IMS network administrator, you should plan for each of the activities listed in the following table, which also lists where you can find additional information on each activity.

Table 62. Network administration activities

Activity	Related reading
Identifying the terminals and other devices required in the online system, verifying that IMS supports these devices, and identifying any incompatibilities.	For information on the terminals that IMS supports, the communication modes that IMS supports for each terminal, and the features for each terminal, control unit, or CPC, see <i>Terminals and equipment supported by IMS Version 11 (Release Planning)</i> .
Establishing appropriate security features in the network	For information on establishing security, see “Planning for security” on page 410.
Determining message editing requirements	For information on editing and formatting messages, see Chapter 26, “Editing and formatting IMS messages,” on page 437.
Determining requirements for exit routine	For information on using exit routines to customize your IMS network, see <i>Exit routines (Exit Routines)</i> .
Enforcing suitable naming conventions	For information on establishing naming conventions, see <i>Establishing naming conventions (System Administration)</i> .
Defining the network at IMS system definition, and coordinating the system definitions for the host	For information on IMS network system definition, see <i>Defining terminals with data communication macros (System Definition)</i> .
Defining ETO descriptors and MFS device characteristics table entries	For information on ETO and the MFS device characteristics table, see Chapter 6, “Administering the Extended Terminal Option,” on page 87.
Defining LU 6.2 descriptors to APPC/IMS	For information on APPC/IMS and LU 6.2, see Chapter 4, “Administering APPC/IMS and LU 6.2 devices,” on page 47.
Testing the readiness of the network	For information on testing the network, see <i>Ensuring network readiness (System Administration)</i> .
Monitoring the performance of the network, and analyzing message loading, such as transaction routing and message queue size	For information on performance monitoring, see <i>Performing capacity planning (System Administration)</i> .

Documenting network and terminal requirements

The size of your IMS installation and your terminal profile determine how you document your network requirements.

The terminal profile

Creating a terminal profile is the best approach for gathering the information you need for your system definition. Your terminal profiles also assist you in coordinating the installation arrangements. They help you describe, to the system programming staff, how each terminal is to be used.

Begin by examining the application program specifications or the hardware plans to identify the terminals that are to be used.

Terminals attached using VTAM

Terminals attached through VTAM are defined according to terminal type. Communication and attachment modes include:

- Binary synchronous communications (BSC)
- Synchronous data link control (SDLC)
- Local attachment (directly to a channel)

When you define your IMS or select an ETO logon descriptor, your choice of terminal type determines the communication protocol and the MFS formatting that IMS uses for that terminal. You can use more than one terminal type for a particular physical terminal; however, your choice must be compatible with the application programs that execute on the IMS.

IMS terminal network

An IMS network is characterized by physical terminals and other devices. IMS supports many terminals, including display devices, printers, and remote intelligent controllers.

IMS uses an access method, such as VTAM, to communicate with a physical terminal. Be sure to understand the difference between:

- VTAM's view of terminals that are connected to a host
- IMS's supervision of those connections

VTAM's responsibilities include:

- Receiving the input message and processing it
- Transmitting data
- Managing line and terminal communication status

IMS's responsibilities include:

- Monitoring the message traffic
- Routing messages to application programs
- Handling message storage, scheduling, and recording

Although IMS uses the network facilities of VTAM, it can also control devices that use BSAM.

VTAM is the preferred access method for IMS. You can use OTMA for non-VTAM networks.

To describe a physical terminal configuration to IMS, you code system definition macros, such as the TYPE macro, that describe the terminal attachments and their characteristics (such as communication type, features, or options). You also define ETO and LU 6.2 descriptors.

Terminal connections to IMS

IMS supports various modes of communication or attachment, usually using VTAM. Remote attachments are either switched or nonswitched.

Terminals that are attached by using VTAM are assigned a terminal type. The addresses of the lines and terminals that IMS uses are transparent to IMS, because they are not specified during IMS system definition. Instead, all VTAM resources have names assigned to them.

Definition: Each logical unit is assigned a *node name* that your installation defines.

The node name that is chosen for the VTAM definition is also used in IMS system definition and commands. Using VTAM enables IMS users to share network resources with other VTAM applications.

XRF terminal support

XRF supports all terminals that IMS supports. XRF terminal support is divided into three categories: class 1, class 2, and class 3.

Related concepts:

“Terminals in an XRF complex” on page 416

 Extended Recovery Facility Overview (System Administration)

Logical terminals (LTERMs)

A *logical terminal (LTERM)* is a user destination. For statically defined terminals, each LTERM is associated with a physical terminal. For ETO terminals, each LTERM is associated with a user, and is associated with a physical terminal only after a user has signed on to a physical terminal.

Each logical terminal has an installation-defined name, called the *LTERM name*.

Reserve one LTERM and identify it as the IMS master terminal. This logical terminal is the control point for the online IMS. Commands associated with this LTERM start and stop the system, control the system resources, and display the status of those resources. The master terminal operator (MTO) can use the /ASSIGN command to alter the physical device that is associated with the LTERM. When ETO users sign on, altering the physical device that is associated with the LTERM is automatic. The device is identified by either the node name or the line number with the physical terminal identifier.

When a user enters a transaction, the logical terminal name is associated with the input message, and input messages are queued by transaction code. The output message queue designation is actually the LTERM name itself. Although this name is often the same for output as for input, an application program can cause the output to be directed to an alternate LTERM. An input terminal can also specify an LTERM name as the destination of a message. A message switching (input) edit routine can append the current input LTERM name to the message.

Related concepts:

“Designing logical terminal networks” on page 397

Related reference:

 Transaction Manager exit routines (Exit Routines)

APPC/IMS and LU 6.2 terminal support

APPC/IMS does not use an LTERM for input. APPC/IMS provides two facilities (that are similar to LTERMs) for determining output destinations: LU 6.2 descriptors and side information.

LU 6.2 descriptors allow an IMS LTERM name to define an LU 6.2 destination (specifying LUNAME, TPNAME, and other LU 6.2 destination characteristics). These LU 6.2 LTERMs can be used the same way as regular LTERMs in IMS application programs.

Side information contains system-defined values provided by CPI Communications (CPI-C), and supplies LU 6.2 destination information. For APPC/IMS to establish a conversation with a partner program, CPI-C requires initialization information, such as the name of the partner program and the name of the LU at the partner's node. CPI-C provides a way to use system-defined values for these required fields. These system-defined values are specified in the side information, and are accessed by a symbolic destination name. The symbolic destination corresponds to an entry in the side information containing the partner_LU_name, the mode_name, and the TP_name.

Related reading: For more information on side information, symbolic destination names, and the entries in the side information, see *z/OS MVS Planning: APPC/MVS Management*.

IMS messages and their scheduling

An *IMS message* can be one of four types of data communication for which IMS controls processing: transactions, messages sent to LTERMs, IMS Commands, and DFSAPPC.

Transactions

Transactions are input messages that are destined for processing by application programs. Transactions are identified by a one- to eight-character *transaction code*. Transactions can be entered at terminals or generated by application programs. In a multiple-systems environment, a transaction can originate in a remote IMS, or in another subsystem (like DB2 for z/OS).

Messages sent to LTERMs

Messages that are sent to LTERMs are identified by LTERM names. These messages can be either of the following types:

- Output messages sent from application programs to communicate with a logical terminal. They usually acknowledge or give results of work, and are sent to the originating terminal. Application programs can also send output to logical terminals other than the one that originated the input message.
- Input messages whose destination is a logical terminal. This message type is known as a *terminal-to-terminal message switch*. The text of the message becomes the output message to the destination LTERM.

IMS commands

IMS commands are entered by terminal users who request that IMS display or alter the status of one or more IMS resources. Most commands originate from the master terminal operator (MTO), but other terminals, as well as application programs, can also enter them.

DFSAPPC

DFSAPPC is an IMS service that is used for exchanging messages between LU 6.2 devices, or between any combination of LU 6.2 and non-LU 6.2 devices. Messages are delivered by allocating a new conversation with the designated LU 6.2 destination device. If the allocated conversation fails, the output is stored for future delivery. Messages are held on the IMS message queue until they can be successfully delivered.

Message queuing

In IMS systems that do not use shared queues, all input and output messages are queued in IMS-controlled virtual storage. Most command output is queued in virtual storage. Application programs handle messages serially if SERIAL=YES is specified on the TRANSACT macro.

If virtual storage is exceeded, messages are saved on direct-access storage. In this way, messages can be received as input or saved for output, although the resources that are necessary to process them might not be immediately available.

The maximum number of concurrent terminal I/O requests and the amount of virtual storage to hold the messages are specified during IMS system definition. Both variables can be altered when IMS is started.

In a shared queues environment, input and output messages are queued to a queue structure in a coupling facility. The shared queues environment is documented in *IMS Version 11 System Administration*.

Message segments

When a terminal is used to enter data into the IMS online system, the input can consist of either single segments or multiple segments. A sequence of several segments might be required to specify a message that can be interpreted by an application program. The device operation determines the end-of-segment.

For input, detection of the end-of-message condition by IMS indicates that a complete message has been received.

When the first end-of-segment condition is detected, IMS examines the beginning of the segment to find a destination. If a destination is declared to be a single-segment message, the end-of-segment signals the end-of-message.

After editing, the first (or only) segment of a message has the following structure:

LL	ZZ	DEST-CODE	b	MMMM
2 bytes	2 bytes	1 to 8 bytes	1 byte	message text

These parameters are explained in the following table.

Table 63. Message segment format

Message segment part	Bytes	Explanation
LL	2	Total length of segment including LL and ZZ parts
ZZ	2	Reserved halfword
DEST-CODE	1 to 8	Destination code (usually a transaction code)
MMMM	Varies	Message text

Each input message is uniquely identified by its destination code. The destination is normally a transaction code, but it can also be an IMS command or LTERM for message switching.

Invalid destinations

Even when the destination code meets normal resource naming conventions, the input destination can be incorrect:

- If ETO is not active, destinations that cannot be identified in IMS are considered invalid, and the input message is rejected.
- If ETO is active, IMS assumes these destinations are LTERM names and creates dynamic user structures. If the destination is invalid, a *dead-letter queue* is created. Several commands are provided, including the /DISPLAY command, for the MTO to process the dead-letter queues.

Transaction codes

The transaction code has optional attributes that affect the processing program's scheduling eligibility by the IMS control program.

Application programs are defined in separate but related macro instructions. The application program that is designated to process a particular transaction code is considered by IMS to be another transaction attribute. An application program might process several transaction codes, but a transaction code can be associated with only one program.

Message scheduling

Definition: *Message scheduling* is the process by which a completely received input transaction is united with its associated application program for processing.

Some of the factors that affect the message scheduling process are:

- The variable attributes associated with the transaction code
- The number and relative importance of other transaction codes
- The number of received messages that are not yet processed
- The intent of associated application programs toward the data to be processed
- The amount of currently available space in control-block storage pools and buffers

In addition, each of the following activities affect the message scheduling process:

- Selecting system definition options
- Designing and using databases
- Specifying buffer sizes
- Declaring and selecting transaction code priorities

By properly combining these activities, you can manipulate the sequence in which messages are processed and enhance system performance.

CPI-C transactions

IMS resources are made available to the CPI Communications driven application program when the application program issues the APSB (Allocate_PSB) call. The CPI Communications driven application program can use the SAA Resource

Recovery Commit (SRRCMIT) and Backout (SRRBACK) calls to initiate an IMS synchronization point or backout. CPI Communications driven application programs should use the SAA Resource Recovery calls to initiate an IMS sync point prior to program termination.

Definition: A *synchronization point* (also referred to as a *sync point* throughout this manual) is an occurrence within a program or subsystem wherein resources are committed and a reference point is established for subsequent restart, if necessary.

Recommendation: Define CPI transactions with a different message class from that used for non-CPI transactions.

Related concepts:

➡ DFSAPPC message switch (Application Programming)

➡ Performing capacity planning (System Administration)

Related tasks:

“Asynchronous output to an invalid destination” on page 126

Part 2, “CPI Communications and APPC/IMS,” on page 35

Message flow within the IMS online system

This topic describes the events that take place as a message that is entered at a terminal flows through an IMS online system. Although some exceptions to the order of events and processing exist, the majority of transactions are processed this way.

IMS provides two ways to customize processing: basic edit (the default) and MFS.

Message flow from terminal to program

The initial entry from a terminal is a message that is processed asynchronously by IMS. After the message is received from the terminal, the message is formatted using basic edit, MFS, or exit routines.

The IMS message scheduler then uses the status of the message queues and a prioritizing algorithm in order to select the next transaction to schedule. When the program is selected for scheduling in the dependent region, the first segment of the message is made available to the program.

Message flow from program to output terminal

While the application program is executing, it can use DL/I calls to access databases. The data communication support of DL/I is used by the program to request messages, using the GU function for the first message segment and the GN function to retrieve subsequent message segments. The program can then send a response to the entering terminal by inserting the reply to the I/O PCB. The entering terminal's LTERM is the I/O PCB after a successful GU call.

If the message is handled by MFS, it can be transformed from the program output format to the device output format. Otherwise, the message is passed unchanged to the device.

When the output message has been completely received by a terminal or by a program, it is dequeued. Any failure in message delivery causes IMS to keep the recoverable message queued for a later delivery.

Message flow from program to alternate destination

The application program can also send output to an alternate destination. Using an alternate PCB, the program can insert a message to another LTERM. In this case, the message handling is the same, but the message queue on which it is placed has a different LTERM name.

Message flow from program to program

The application program can also use the alternate PCB to send a message to a program; that is, it can generate a secondary transaction that is placed on a message queue. The processing of the secondary transaction is dependent on the selection of the transaction by the scheduling algorithm.

Related tasks:

Chapter 26, "Editing and formatting IMS messages," on page 437

Conversational transactions

Conversational transaction processing allows you to retain message continuity from a given terminal, even when the program that processes the conversation is not retained in storage throughout that conversation.

When the transaction is defined as conversational, the application program can use a scratchpad area (SPA) in order to interrelate messages from a given terminal.

The *scratch pad area (SPA)* is a work area used in conversational processing in order to retain information from an application program across executions of the program. A unique SPA is created for each conversation.

Contents of the SPA

Typical contents of the SPA are data items entered at the terminal and data that is retrieved from databases to be saved between iterations of the conversation. To simplify operator procedure and application design, database updates are retained in the SPA until the last iteration of the conversation.

Any subsequent data entry from a terminal that is already operating in conversational mode causes a message processing program (MPP) to receive both the contents of the SPA and the input terminal data. Each input message is considered an individual unit of work. Each interaction can be with a different program, although the same program is typically used.

Message processing for conversational transactions

When the message is a conversational transaction, the following sequence of events occurs:

- IMS removes the transaction code and places it at the beginning of a message segment. The message segment is equal in length to the SPA that was defined for this transaction during system definition. This is the first segment of the input message that is made available to the program. The second through the *n*th segments from the terminal, minus the transaction code, become the remainder of the message that is presented to the application program.
- When the conversational program has prepared its reply, it inserts the SPA to IMS. The program then inserts the actual text of the reply as segments of an output message.

- IMS saves the SPA and routes the message to the input LTERM.
- If the SPA insert specified that another program is to continue the same conversation, the total reply (including the SPA) is retained on the message queue as input to the next program. This program then receives the message in a similar form.
- A conversational program must be scheduled for each input exchange. The other processing continues while the operator at the input terminal examines the reply and prepares new input messages.
- To terminate a conversation, the program places blanks in the transaction code field of the SPA, and inserts the SPA to IMS.
- The conversation can also be terminated if the transaction code in the SPA is replaced by any nonconversational program's transaction code, and the SPA is inserted to IMS. After the next terminal input, IMS routes that message to the other program's queue in the normal way.

ETO conversations

ETO conversations maintain associations with users, rather than with terminals. The conversation is associated with the terminal only while the user is signed on. The conversation can be restarted on any other ETO terminal.

Program switches for conversational programs

A conversational program can use a deferred program switch or an immediate program switch.

Definitions:

- During a *deferred program switch*, the program responds to the originating terminal but causes the next input from the terminal to go to another conversational program.
- During an *immediate program switch*, the program passes the SPA (and, optionally, a message) to another conversational program without responding to the originating terminal. In this case, it is the next program's responsibility to respond to the originating terminal.

Related tasks:

Part 3, "Extended Terminal Option (ETO)," on page 77

Message switches

A *message switch* occurs when an input message specifies an LTERM name instead of a transaction code. IMS formats the message and inserts it into the output message queue for that LTERM.

Designing logical terminal networks

The IMS system definition describes the characteristics and relationship of communication lines, static terminals, and logical terminals (LTERMs).

For a nonswitched terminal, the relationship between a physical terminal and a logical terminal within IMS is a static relationship defined during system definition.

When only one user operates a physical terminal, only one logical terminal is associated with that physical terminal. If multiple users operate a physical terminal, the terminal is associated with many logical terminals.

You can structure the IMS system definition so that a separate logical terminal is assigned for each user of a particular static terminal.

Definitions: The information in this topic uses *physical terminal* to represent anode, for VTAM devices.

Logical terminals can be assigned to physical terminals for both input and output. When a logical terminal is assigned to a physical terminal for output, all messages that are sent to that logical terminal are transmitted to its associated physical terminal. More than one logical terminal can be assigned to a given physical terminal for output. Only one physical terminal can receive the output for a given logical terminal. The relationship between physical and logical terminals for output is shown in the following figure.




Figure 44. Relationship between physical and logical terminals for output from IMS

Logical-terminal chains

When a logical terminal is assigned to a physical terminal for input, any message that is entered from the physical terminal is considered to have originated at the logical terminal. When more than one logical terminal is assigned to a physical terminal for input, a chain of input logical terminals is formed. Any input from the physical terminal is considered to have originated at the first logical terminal on the chain.

The following figure shows the relationship between physical and logical terminals for input.




Figure 45. Relationship between physical and logical terminals for input to IMS

If the first logical terminal is not allowed to transmit a message (for example if it is not authorized or it is a stopped logical terminal), all logical terminals on the input chain are interrogated in chain sequence for their ability to transmit messages. If the physical terminal is a Finance, SLU 1, SLU P, or LU 6.1, only the logical terminals associated with the input component are scanned. The first appropriate logical terminal found is considered the originator of the message. If no appropriate logical terminal is found, the message is rejected with an error message.

Logical-terminal queues

Using a message queue for received input messages or for pending output messages enables an application program to be independent of the arrival time and message transmission. Because this message queue is associated with a logical terminal, rather than with a physical terminal, the message queue can be moved from device to device, independent of the application program. The message queue can even be moved among device classes.

Logical terminals provide stability for application programs. Regardless of how the physical terminal network changes, the application programs remain insensitive to these changes.


The application program interface to a logical terminal is conceptually the same as that for the database system:

- GU calls retrieve message segments from a queue.
- ISRT calls insert message segments to a queue.

Separating input and output devices

In certain application programs, it might be necessary to associate a different physical device for output than the one that is used for input. If the physical terminal type is an input-only device and output is required, a different device must be associated for output.

IMS system definition and commands support assignment of output devices that are different from the input device. For example, the application program that is processing a message from a display might need to send some of the output to a printer. The possible physical-to-logical relationships are shown in the following figure.



Notes to figure regarding applications

1. Normal assignment of one or more logical terminals or physical terminals. Output is sent to input terminal. Application is insensitive.
2. Application uses specific logical terminal for output. Application is insensitive to input.

Figure 46. Possible physical-to-logical terminal relationships

Logical and physical terminal relationships

The following figure shows the communication flow between a terminal user, a physical terminal, a communication line, and a logical terminal in a nonswitched communications network.




Figure 47. A nonswitched communications network

IMS system definition describes the characteristics and relationship of physical terminals, communication lines, and logical terminals. On a nonswitched communication line, the relationship between a physical terminal at one end and a logical terminal within IMS at the other end is a stable relationship defined during system definition or during signon for ETO terminals.

Except for LU 6.1, terminals using VTAM have physical-to-logical terminal (PTERM-to-LTERM) relationship as a terminal on nonswitched line.

The relationship that is established between a physical terminal and one or more logical terminals at system definition can be changed using commands or by creating a new system definition. The /ASSIGN command dynamically changes logical and physical relationships, and it is usually executable only from the master terminal.

Master terminal

The master terminal is the central point of control for IMS. If a VTAM master terminal is defined, IMS establishes a connection with it during startup.

Recommendation: Although several device types are supported, a VTAM SLU 2 device is the preferred type for the master terminal.

Restriction: Neither an ETO terminal nor an LU 6.2 terminal can be defined as the primary or secondary master terminal.

The master terminal operator (MTO) must be familiar with all aspects of operating the system. The physical location of the master terminal in relation to the computer console is important. If, for security reasons, the MTO and master terminal are not close together, telephone communication should be provided.

Master terminals in an XRF complex

In an XRF complex, each IMS must have its own master terminal, and it can have a secondary master terminal.

Typically, the secondary master terminal is a printer (output-only device type). System messages for the active and alternate subsystems go to their respective master terminals.

During a takeover, the console operators must know the status of the takeover on both systems. The operator at the alternate subsystem must know when I/O prevention in the failed active subsystem is complete; terminal switching of class-3 terminals and failed class-2 switch attempts must be handled.

Related concepts:

 Choosing the IMS master terminal (System Definition)

NTO terminals

When using the NTO licensed product, VTAM can support some start-stop and teletype devices, known as *NTO devices*.

The receipt of a message can be confirmed by using transaction-response mode to “lock” the keyboard following the input entry. When IMS responds with available output to a particular terminal for a specific transaction, it “unlocks” the keyboard and sends the data.

Even if the NTO devices are used primarily for inquiry transactions, the transactions can be made recoverable. Because input is not acknowledged, the application program should indicate in the output response the inquiry transaction that caused it.

Both input and output editing routines (which are provided by your installation and are entered after MFS processing) can be made available for transactions using NTO devices.

Because NTO uses VTAM, IMS is not aware of whether a terminal user is logging on to a session or is logging off. Therefore, existing session status might cause problems when one terminal user ends a session and another terminal user resumes the session. This problem can be minimized by defining all NTO devices as Response mode.

To avoid the problem of continuing session status, a terminal user can clear any preset transaction codes that were previously set using /SET commands by entering the /RESET command either:

- Before issuing the /RCLSDST command when ending an NTO session
- Before issuing the /SIGN ON command when starting an NTO session

A terminal user who enters /RCLSDST or the MTO who enters /CLSDST terminates a session but does not reset conditions that were set during that session.

Related concepts:

 How to communicate with IMS from an NTO device (Operations and Automation)

Resource modes and states

IMS keeps three types of status information for the terminals in the network.

The types of status information that IMS keeps are:

- The mode of operation
- The state of an inoperable terminal (such as temporarily unusable)
- The recovery status

Terminal and user operating modes

A terminal can be in more than one mode at the same time.

Possible modes for a terminal or user include:

Response Mode

Response mode is established through static terminals, ETO users, or transaction specifications. Response mode can be used with full function or Fast Path processing. In response mode, after the entry of an input message, the terminal is locked until a reply is received, and no additional incoming data is accepted.

Definitions:

- For operator-driven terminals like 3270s, *locked* means that the keyboard is locked.
- For programmable terminals like LUPs, *locked* means that IMS delays the acknowledgment of the input message until the output message is ready to be sent.

For LU 6.2, the originator of the transaction must issue a Receive command to get the response message (synchronous output). Failure to do so is treated as a protocol violation. Depending on the system, the application program could fail and IMS could send an error message.

Terminal-response mode is terminated when the response has been sent and dequeued. If Fast Path is used, terminal-response mode is automatically continued following an IMS failure, when a static terminal logs off, or when an ETO user signs off. If full-function operation is used,

terminal-response mode is terminated when IMS is restarted, when a static terminal logs off, or when an ETO user signs off, unless the terminal or user is defined with both SRMDEF=LOCAL and RCVYRESP=YES, in which case the full-function response mode is recoverable after either a terminal logoff or a user signoff.

The ETO user, not the terminal, is in response mode. If the user signs off and Fast Path is used, response mode is recovered for the user, but the terminal is no longer in response mode. The user cannot enter another transaction until the response to the first transaction is received.

Conversation Mode

After you enter a transaction that is defined using the SPA= parameter on the TRANSACT macro, the terminal is in conversation mode. While the terminal is in conversation mode, other transactions cannot be entered from that terminal. However, the terminal is not locked, in the same sense that response mode locks the terminal. The terminal remains in conversation mode until the conversation is terminated, when the message has been sent and dequeued, and the application program has placed blanks in the transaction code field in the scratch pad area (SPA).

A conversation can abnormally terminate under the following conditions:

- When an application program abends
- When the IMS MTO issues an /EXIT command, a /START NODE command, or a /START USER command
- When an inconsistent definition exists in the application program between IMS and MSC

For LU 6.2, all iterations of the IMS conversation must use the same LU 6.2 conversation; each iteration of the IMS conversation is demarcated by using the LU 6.2 CMPTR (Prepare_To_Receive) call. If the LU 6.2 conversation ends prior to the end of the IMS conversation, the IMS conversation is abnormally terminated.

The ETO user, not the terminal, is in conversation mode.

Exclusive Mode

A terminal is placed in exclusive mode when the /EXCLUSIVE command is issued. The Exclusive mode:

- Restricts the output received by the terminal.
- Remains with ETO users after they sign off. ETO users that sign off while they are in exclusive mode remain in exclusive mode when they sign on the next time.
- Terminates with an /END or /START NODE command.

Lock Mode

Lock mode prevents a terminal from sending and receiving messages. A terminal, node, or logical terminal (LTERM) is placed in lock mode when the operator issues the /LOCK command. Lock mode is reset by issuing the /UNLOCK command or the /IAM command.

Test Mode

Test mode ensures that any input message entered into a terminal is transmitted back to the terminal. A node is placed in test mode by the /TEST command. Test mode is reset by an /END command or a /START command. Test mode is not significant and is not carried across restart and ETO signoff and signon.

Related tasks:

Related reference:

 [IMS commands \(Commands\)](#)

Terminal and user states

Several terminal and user states render the terminal or user inoperable.

The states in which a terminal is inoperable include:

Stop State

The stopped state prevents the delivery of any output queued on a logical terminal that is associated with a physical terminal.

The /STOP NODE command results in the termination of the session between IMS and the node. This termination occurs immediately for most devices but only at the end of the message for 3270 devices and SLU 2 devices.

The /STOP NODE command also prevents a new session until a /START command or /RSTART command is issued.

The /STOP command, the /PSTOP command, and the /MONITOR command also cause a terminal to enter the stopped state. This state is reset by the /START command or /RSTART command. The /STOP NODE or USER command allows ETO users to retain their status after signoff or logoff.

QERROR State

A logical terminal is placed in a stopped state if an I/O error is encountered while attempting to read from or write to a message queue. This condition is reset when the MTO operator issues a /START command.

QLOCK State

A logical terminal is prevented from sending any additional output until the locked state is reset by a specific request received in the LU 6.1 session. The LTERM is also prevented from receiving input that might create additional output. This condition is reset when the MTO issues a /START command. A /PSTOP or /PURGE command is ignored for LTERMs that are in QLOCK state.

INOP State

The physical terminal is considered inoperable by IMS device support whenever an error is detected and put in the INOP state. All logical terminals associated with this physical terminal are also considered inoperable. The /START or /RSTART command resets the inoperable condition.

COMPINOP State

Component inoperative can be set in one of two ways:

- When an error is detected that is isolated to a component of the terminal
- When the /COMPT or /RCOMPT command is issued for terminals that are defined to VTAM

All logical terminals associated with this component are ineligible for message output. The component inoperative state is reset when the operator issues a /START LINE PTERM command, a /START NODE command, another /COMPT command, or a /RCOMPT command. Special signals from the device, such as device end from a 3270 device or the “component available” status from a SLU 1 can also cause a reset.

PAGE, SCREEN, and COMPONENT PROTECTION State

This is a state supported for video terminals, SLU P, Finance, and LU 6.1 devices. Logical terminals associated with these physical terminals are not eligible for output selection.

SNA QUIESCE State

When IMS sends output messages to a VTAM programmable node and the node wants to stop receiving, the node signals IMS to stop transmissions after an end-of-chain has been sent. IMS does not send any more output to the terminal until the terminal sends the SNA release-quiesce (RELQ) command.

Related concepts:

 3270 terminal screen protection (Operations and Automation)

Related tasks:

“Suspending output from IMS” on page 934

Part 3, “Extended Terminal Option (ETO),” on page 77

Resource status recovery

Resource status recovery defines how a terminal or user is recovered.

Recovery information can be shared using RM; therefore a resource can be recovered without using the IMS log records. In an IMS failure, a terminal or user can resume work without having affinity to the failed system or having to wait for IMS restart. You can also choose to recover resources only on a local system, or to delete the resource status.

This topic defines the resource status classifications and recovery modes. It then describes how recovery mode status is used with Fast Path and XRF.

Resource status classification

IMS classifies resource status to determine how much information needs to be stored in the resource structure or in local log records.

Nonrecoverable Status

Status only exists when the resource is active. Status is deleted when the resource becomes inactive and is not recovered at terminal logoff, user signoff, or IMS restart.

Recoverable Status

Status is recovered, but does not prevent the resource from being deleted across signoff, logoff, or IMS restart.

Significant Status

Status is recovered and the resource is not deleted across signoff, logoff, or IMS restart. Where the status is maintained depends on whether the status is command or end-user.

Command Significant Status

Status relates to the resource command, such as STOP, TRACE, and MFSTEST. Status recovery is always maintained globally by RM in the resource structure, if defined. Status is unaffected by status recovery mode. Not all statuses set by commands are significant.

End-User Significant Status

Status relates to resource work: conversation, STSN, and Fast Path.

The status frequently changes, which can affect performance. Therefore, you can specify the status recovery mode as either GLOBAL, LOCAL, or NONE.

Related reference:

“Status recovery of TM resources” on page 409

Status recovery mode for end-user significant status

The status recovery mode defines the scope and location of recovery for resources with End-User Significant Status.

You can set the default mode for each IMS, which is used for all resources unless overridden during a logon or signon. You can then set what End-User Significant Status to recover.

The following list describes how to set the status recovery modes for TM resources:

- Specify IMS defaults in DFSDCxxx.
- Modify user descriptors.
- Set the DFSLGNX0 override for static terminal logon and dynamic STSN terminal logon.
- Set DFSSGNX0 override for dynamic non-STSN user signon.

The following list describes the three recovery modes for End-User Significant Status.

GLOBAL Status Recovery Mode

All recoverable status is saved locally in the IMS record logs, but uses RM to recover the status instead of the logs. Status is restored at the next logon or signon and is available to any IMS in the IMSplex. When the resource becomes active, status is copied to the local system. When the resource becomes inactive, status is deleted from the local system.

RM, a coupling facility resource structure, and shared queues are required. If you are sharing queues and have a resource structure, GLOBAL is the default. The default is overridden by DFSDCxxx, user descriptors, or logon and sign-on exit routines.

LOCAL Status Recovery Mode

All recoverable status is saved locally in the local control blocks and log records. Status is restored at the next logon or signon and is only available to the same IMS that the user or node was accessing. Additionally, the user or node can only access the IMS where the local status is. This affinity is called *RM affinity*, which is enforced when End-User Significant Status exists. With RM affinity, IMS does not allow a terminal or user to log on to or sign on to an IMS if RM indicates that the user or terminal has RM affinity to another IMS. This affinity occurs because end-user significant status (conversation, STSN, or Fast Path) is being recovered on another IMS.

When the IMS with the RM affinity fails, the RM affinity still exists. The user or node can access another IMS immediately if the Logon exit routine (DFSLGNX0) or the Sign-on exit routine (DFSSGNX0) allows it, but resource status is not recovered and local status is deleted on the failed IMS at restart.

LOCAL is the default if a resource structure or shared queues are not used.

RM is not required. If RM is not active, RM affinity is not enforced.

NONE Status Recovery Mode

No status is saved in RM or local log records. At logon or signon, significant status is cold. STSN, conversation, and fast path status is automatically nonrecoverable.

RM and a resource structure are not required.

Recoverability of specific resource types

One way to specify the recoverability of specific resource statuses (conversation status, STSN status, and Fast Path status) is to use the parameters in the DFSDCxxx PROCLIB member.

The following parameters apply to the recoverability of conversation status, STSN status, and Fast Path status.

RCVYCONV= YES | NO

Specifies whether conversation status is recovered. This parameter does not affect output messages. If conversation status is not recovered, the output is still recovered and delivered asynchronously. YES is the default when SRMDEF=GLOBAL or LOCAL. NO is the default and YES is invalid when SRMDEF=NONE.

RCVYFP= YES | NO

Specifies whether Fast Path status and Fast Path output are recovered. YES is the default when SRMDEF=GLOBAL or LOCAL. NO is the default and YES is invalid when SRMDEF=NONE.

You must specify YES if you specify that STSN status is recovered.

RCVYRESP= YES | NO

Specifies whether full-function response mode is recovered after either a terminal logoff or a user signoff. RCVYRESP=NO is the default. RCVYRESP=YES is invalid when SRMDEF=GLOBAL or NONE or when a CSL Resource Manager (RM) is active.

RCVYSTSN= YES | NO

Specifies whether STSN status is recovered for STSN terminals (SLU P, Finance, and ISC). This parameter affects only the recovery of STSN sequence numbers and does not affect output messages. YES is the default when SRMDEF=GLOBAL or LOCAL. NO is the default and YES is invalid when SRMDEF=NONE.

Related reference:

“Fast Path recovery”

Fast Path recovery

Recovering Fast Path transactions depends on the status recovery mode, the Fast Path recovery, the STSN recovery, and where the Fast Path transaction runs.

The following table lists the recovery and status of Fast Path transactions based on these criteria. If the transaction runs locally without going through the EMH queue, IMS cannot recover Fast Path status globally in RM. In the situation where the status recovery mode is GLOBAL but the transaction is running locally, the status recovery mode is temporarily changed to LOCAL with the terminal or user having RM affinity.

Table 64. Determining Fast Path recoverability

Status recovery mode	Fast Path recovery (RCVYFP)	STSN recovery (RCVYSTSN)	Local Fast Path (DBFHAGU0)	Shared EMH queues (DBFHAGU0)
LOCAL or GLOBAL	NO	NO	Status and messages discarded	Status and messages discarded
LOCAL or GLOBAL	NO	YES	INVALID	INVALID
LOCAL	YES	NO	Status and output recovered locally. STSN cold started.	Status and output recovered locally. STSN cold started.
GLOBAL	YES	NO	Status and output recovered locally. STSN cold started.	Status and output recovered globally. STSN cold started.
LOCAL	YES	YES	Status and output recovered locally. STSN recoverable.	Status and output recovered locally. STSN recoverable.
GLOBAL	YES	YES	Status and output recovered locally. STSN recoverable.	Status and output recovered globally. STSN recoverable.

Recovery modes for XRF terminals

Status recovery modes affect each Extended Recovery Facility (XRF) terminal class differently during a takeover in the XRF environment.

Operating in an XRF environment does not affect status recovery modes. However, recovery of class-1 terminals is slightly different when using status recovery modes in a sysplex environment compared to a an environment not using RM.

Class-1 terminals

IMS relies solely on local log records for an XRF takeover, even if the status recovery mode is GLOBAL. When session takeover is complete, IMS reconciles any differences in status between RM and the new active system. The GLOBAL mode provides the benefit of a user or terminal being able to log on to another IMS in the IMSplex following normal termination, but for class-1 terminals, LOCAL mode may be more appropriate.

If the status recovery mode is LOCAL, RCVYRESP=YES, and RM is not used, full-function response mode is recovered for class-1 terminals.

Class-2 terminals

IMS uses local log records during a takeover if status recovery mode is LOCAL and RM if status recovery mode is GLOBAL. The new active system logs active terminals on, but user signon is required. IMS uses local log records to initiate logon, but determines recovery based on the status recovery mode of the user or terminal.

Class-3 terminals

IMS does not automatically take over these terminals. Recovery is based on the status recovery mode.

Related concepts:

“Planning for the Extended Recovery Facility (XRF)” on page 414

Status recovery of TM resources

IMS saves recoverable statuses for LTERMs, nodes, users, and user IDs. This topic lists the statuses saved for each.

LTERM recovery status

IMS saves the following LTERM recoverable statuses:

- LTERM name
- EDIT=UC (upper case translation specification)
- Node or user owner

IMS saves the following LTERM command significant statuses:

- /ASSIGN SAVE
- STOP

Node recovery status

IMS saves the following recoverable statuses:

- Node name
- Device type
- Allocated LTERM name
- Allocated user name

IMS saves the following command significant statuses:

- EXCLUSIVE
- MFSTEST
- STOP
- TRACE

IMS saves the following end-user significant statuses:

- Conversation
- Fast Path
- STSN
- Full-function response mode after a terminal logoff or a user signoff, but only if RM is not used, SRMDEF=LOCAL, and RCVYRESP=YES. Full-function response mode is not recovered after a termination or failure of IMS.

User recovery status

IMS saves the following user recoverable statuses:

- User name
- User ID
- Allocated LTERM names
- Allocated node names
- Autologon parameters

IMS saves the following user command significant statuses:

- EXCLUSIVE
- MFSTEST
- STOP
- /CHANGE AUTOLOGON SAVE

IMS saves the following user end-user significant statuses:

- Conversation
- Fast Path
- Full-function response mode after a terminal logoff or a user signoff, but only if RM is not used, SRMDEF=LOCAL, and RCVYRESP=YES. Full-function response mode is not recovered after a termination or failure of IMS.

User ID recoverable status

IMS saves the following LTERM recoverable statuses:

- User ID
- Terminal Name

Planning for security

To prevent unauthorized use of a terminal in the IMS network, you can use RACF (or an equivalent product).

RACF is a licensed program available under the z/OS operating system.

RACF allows you to control access to:

- Physical terminals
- Logical terminals
- Transactions
- Commands

If you do not use RACF security, IMS allows only certain commands to be entered at user terminals (excluding the master terminal). This is called *default terminal security*.

Using RACF, you can design security profiles based on user ID and you can define two levels of security for your network:

- You can control the use of the terminals connected to your network.
- You can control the resources that can be accessed from the terminal.

You control use of a terminal by signon verification security. For example, a terminal user enters an identifier as a parameter on a /SIGN command or in response to a DFS3649 message. You can use RACF, an exit routine, or both to validate the signon. The user ID is logged with each input and output message and with each database change.

Related concepts:

 Potential use of IMS commands (Operations and Automation)

Related reference:

 Terminal security defaults for IMS type-1 commands (Commands)

Authorizing transactions in a TM network

During transaction authorization, RACF checks to ensure that the terminal user is authorized to enter the transaction.

Related concepts:

 [IMS security \(System Administration\)](#)

Using RACF to secure transactions

To control which users can issue which transactions, define the controlled transactions to RACF as TIMS class, and grant authorization to RACF-defined users or groups of users.

You can also use System Monitoring Facility (SMF) logging to track the successes and failures of transaction authorization. Using RACF, request auditing capabilities for your transaction security profiles.

Using the Transaction Authorization exit routine (DFSCTRN0)

You can use the Transaction Authorization exit routine (DFSCTRN0) to check the validity of a transaction that is entered by a user who signed on using the /SIGN command.

You can use the Transaction Authorization exit routine in conjunction with RACF.

Related reading: For more information on using the Transaction Authorization exit routine (DFSCTRN0), see *IMS Version 11 Exit Routines*.

Authorizing commands in a TM network

For command authorization, you can use RACF and the Command Authorization exit routine (DFSCCMD0).

Related concepts:

 [IMS security \(System Administration\)](#)

Using RACF to secure commands

To control which users can issue controlled commands, define controlled commands to RACF as CIMS class, and grant authorization to RACF-defined users or groups of users.

You can use SMF logging to track the successes and failures of command authorization. Using RACF, request auditing capabilities for your command security profiles.

Using the command authorization exit routine (DFSCCMD0)

You can use the Command Authorization exit routine (DFSCCMD0) to restrict keywords and parameters by editing the input command buffer.

The Command Authorization exit routine checks the validity of commands. You can use this exit routine with RACF.

Related reference:

 [Command Authorization exit routine \(DFSCCMD0\) \(Exit Routines\)](#)

Transaction command security

Transaction command security can provide another level of access control.

Transaction command security is indirectly related to a terminal. If the terminal user enters a transaction to start a program that issues IMS commands, both the transaction code and the set of commands that program can issue must be authorized.

Password security

Password security requires that a transaction or a command that is entered at a terminal have a password. This allows you to implement secondary security to verify that the user who is issuing a specific transaction or command is authorized to do so.

Using RACF and a security profile based on user IDs

RACF provides the REVERIFY option, which requires your password to be entered when a command or a transaction is entered. You need to use RACF for signon authorization and include the REVERIFY parameter as APPLDATA when defining the transaction or command to RACF. You must specify RVFY=Y as an execution parameter in order to use REVERIFY.

Security for APPC/IMS

APPC/IMS requires security using the System Authorization Facility (SAF) interface to RACF, or an equivalent security environment. RACF is optional for remote transactions from LU 6.2 application programs.

APPC/IMS supports both the Transaction Authorization exit routine (DFSCTRN0) and the Command Authorization exit routine (DFSCCMD0).

Restrictions:

- APPC/IMS does not support the /SIGN command, because it is not required in order to validate the user ID. z/OS validates user IDs when using RACF; therefore, each APPC/IMS message has a validated user ID.
- For IMS commands entered from remote LU 6.2 application programs: if you do not use RACF or the Command Authorization exit routine (DFSCCMD0), the default command security allows only the following four commands:
 - /BROADCAST
 - /LOG
 - /RDISPLAY
 - /RMLIST

To allow other commands, use DFSCCMD0 or RACF.

Related reference:

 DL/I calls for transaction management (Application Programming APIs)

Security for ETO

You can use RACF security for command and transaction authorization on both static and ETO terminals.

Related concepts:

“Planning a high-security environment with ETO” on page 98

Planning for Fast Path terminals

To have a Fast Path-capable system, specify Fast Path support on the FPCTRL macro.

IMS systems with a very high transaction rate use Fast Path's expedited message handler (EMH) facility. EMH is a performance option that expedites message processing by imposing restrictions on message lengths and segmentation. LU 6.2 terminals can use EMH; no definitions or specifications are required.

With Fast Path, an EMH buffer is acquired from the EMH buffer pool when the first eligible input message is received for a Fast Path transaction. The buffer remains allocated to the terminal for future use until the session terminates or the user signs off. If an EMH buffer is allocated to the terminal, it is reused if it meets the requirements of the next input call. If the message requires a larger buffer, Fast Path swaps it for a larger one.

If, at entry to IMS, an input message satisfies all of the following criteria, it is edited in a Fast Path EMH buffer instead of in a full-function message queue buffer.

- Terminal is not FES capable
- Terminal is not the MTO
- Input is single segment
- Input is non-MFS edited
- Input is not an IMS command
- Execution of at least one Fast Path transaction is scheduled using Fast Path

Exception: If the message is not a Fast Path transaction, the message is moved to a full-function message queue buffer.

When a message with all of these characteristics is received, an EMH buffer is allocated to the terminal. The EMH buffer remains allocated to the terminal until either the user signs off or the session terminates. The buffer is used repeatedly for Fast Path transactions until a larger EMH buffer is required for special applications. When a larger buffer is needed, the current EMH buffer is swapped for the larger-size EMH buffer. The smaller EMH buffer is released and returned to the EMH pool.

The EMHL parameter in the IMS control region initialization (also known as the startup parameter) specifies the default EMH buffer size for all Fast Path transactions for all Fast Path-eligible terminals. A Fast Path transaction can specify an EMH buffer that is larger than the default. Specify larger application EMH buffer sizes on the APPLCTN or TRANSACT stage_1 macros by using the FPATH=*size* parameter when you generate the system. If the EMH buffer pool is exhausted, message DFS3971 is sent to the input terminal.

Buffers are extracted from the EMH buffer pool, which expands and contracts dynamically. The size of the pool depends on the number of terminals that are concurrently entering Fast Path transactions and the buffer sizes that are required to satisfy each request.

You can say FPATH=No or Yes on the APPLCTN and TRANSACT macros, or you can specify FPATH=*size*, where *size* is the EMH buffer size that is required to run

the transaction. `FPATH=size` implies `FPATH=Yes`. The minimum EMH buffer size is 12 bytes, and the maximum is 30720 bytes.

MSDBs are available to ETO terminals, unless they have terminal-related keys.

The `OPTIONS` keyword on the `TERMINAL` macro or ETO descriptor is used to:

- Declare `FORCRESP` or `TRANRESP` so that Fast Path-eligible terminals operate in response mode
- Specify `PAGDEL` for automatic page deletion where appropriate

You can use non-VTAM terminals to enter Fast Path transactions. For non-VTAM terminals, the `LINE` macro must indicate response mode by using the `RESP=TERM` parameter.

Related reference:

 [Terminals and equipment supported by IMS Version 11 \(Release Planning\)](#)

Planning for the Extended Recovery Facility (XRF)

An *XRF* complex enables you to rapidly resume end-user services when a failure occurs.

You can also use XRF to switch all transaction and database processing from one IMS system to another for scheduled maintenance or during off-peak periods.

When you do not use XRF, a high level of availability with short interruptions in end-user service can be difficult to achieve when IMS fails. Termination processing must complete, and the cause of the failure determined before restart can take place. A restart can be time-consuming if IMS regions must be restarted and terminal sessions reestablished.

An XRF configuration consists of a primary system, called the *active IMS*, and a secondary system, called the *alternate IMS*. The alternate IMS tracks the activities of the active IMS in the same or in a different central processing complex (CPC). The active IMS processes the IMS workload in the same way that it would without XRF. It does not do any additional processing or logging. Using the log, the active IMS informs the alternate IMS of its activities. Surveillance mechanisms alert the alternate IMS to problems in the active IMS. The active IMS is not aware of the activities of the alternate IMS.

The alternate IMS does not process any transactions; it maintains the status of the active IMS, and monitors it for any indications of trouble. The alternate IMS tracks resource changes in the active IMS, and updates its control blocks and buffers while continuously changing its status of readiness. In this manner, work can be quickly shifted from the active IMS to the alternate IMS without interruption. This shift of work from the active IMS to the alternate IMS is known as the takeover process. The following figure shows an example of a simple XRF complex.




Figure 48. Sample XRF complex

Related concepts:

[Extended Recovery Facility Overview \(System Administration\)](#)

Using MNPS or USERVAR in XRF complexes

You can use IMS multinode persistent sessions (MNPS) support for XRF. Due to performance limitations, this is not a recommended solution.

You can implement one of two types of XRF complexes:

- An XRF complex that uses VTAM MNPS with an MNPS ACB.
- An XRF complex that uses a user variable (USERVAR) with USERVAR tables.

Both types of XRF complex provide the same increased availability for IMS, but differ in how they manage terminal sessions and the takeover process. Other differences include:

- XRF complexes that use a USERVAR have additional hardware requirements.
- VTAM and the network terminals use different names to communicate with IMS in each type of complex:
 - In an XRF complex that uses an MNPS, IMS is known to VTAM and the network terminals by the name of the MNPS ACB (the MNPS name).
 - In an XRF complex that uses a USERVAR, IMS is known to VTAM by the name of the APPLID ACB (the APPLID name) and to the network terminals by the USERVAR.

XRF complexes that use MNPS still have an APPLID ACB, but it is not used for the purposes of XRF.

Related concepts:

Terminals in an XRF complex

XRF supports all terminals that your IMS users are currently using. XRF provides emergency support with a minimum of effort from you.

XRF requires that you specify keywords on several system definition macros or ETO descriptors. Unless you want to change the switching priority of a terminal from the defaults, or to change the terminal type sessions, you need not change any terminal definition macros.

XRF automatically gives the terminals at your installation the best service that the terminal's characteristics allow. This topic describes these characteristics.

XRF provides support to the following three types of terminals:

Class-1

At takeover, the sessions persist. The sessions are switched transparently to the alternate IMS without losing the session.

Class-2

At takeover, the sessions end. The alternate IMS automatically tries to reestablish the sessions. The users are briefly without service.

Class-3

At takeover, the sessions end. The operator or the end users can establish new sessions with the new active IMS if the terminals are physically connected to it.

Because APPC/IMS devices are Class 3, LU 6.2 conversations that are in progress at the time of failure are not preserved during a takeover. A new LU 6.2 conversation must be allocated from the new active IMS after a takeover occurs. The new active IMS processes APPC/IMS activity exactly the same as the old active IMS did prior to takeover. The new active IMS has the same LU name as the old active IMS.

Class-1 terminals

To take full advantage of XRF, define most of the terminals in your complex as Class-1 terminals. Class-1 terminals are those terminals defined on the UNITYPE keyword of the TYPE macro as either:

- SLUTYPE1
- SLUTYPE2
- SLUTYPEP
- FINANCE

Class-1 terminals are handled differently depending on whether you use VTAM Multinode Persistent Session (MNPS) support.

Class-1 terminal recovery in an XRF system with MNPS

When you use XRF with MNPS, class-1 terminals do not have backup sessions on the alternate IMS. Instead, when a class-1 session is initiated in the active IMS, IMS tells VTAM to track both the session state and data traffic information.

During an XRF takeover, IMS requests that VTAM verify the class-1 terminal sessions and uses the data traffic information to transparently restore sessions. Higher priority terminals are restored before lower priority terminals.

The terminal types listed in the table in “Specifying terminal support” on page 418 are eligible for class-1 service if they are owned by an XRF-capable VTAM that is MNPS enabled. To enable VTAM for MNPS, specify PERSIST=MULTI on the APPL definition.

Class-1 terminal recovery in an XRF system with USERVAR

In an XRF system with USERVAR, the alternate IMS maintains backup sessions for all class-1 terminals with open sessions on the active IMS. These sessions are established during the tracking phase when a session on the active IMS is initiated. In the event of a takeover, the alternate IMS switches all class-1 terminals to the backup sessions.

From the terminal user's point of view, only one session with IMS exists. From the NCP point of view, two host sessions exist. The actual messages are only passed between the node and the active IMS. The alternate IMS tracks the status of the message activity by monitoring the IMS log.

When the active IMS establishes or terminates a terminal session with class-1 terminals, the alternate IMS establishes or terminates its backup session with the terminal. When the active IMS fails, the alternate IMS takes over the terminal session (without losing control from the viewpoint of the node) by changing the mode of the established backup session from BACKUP to ACTIVE. When the session is switched, the NCP sends the new active IMS its view of the terminal's status. IMS compares this with its own status record to decide what, if any, recovery action to take. When the REVERIFY operand is used with RACF, the user must sign on again.

The terminal types listed in the table in “Specifying terminal support” on page 418 are eligible for class-1 service if they are both:

- Owned by an XRF-capable VTAM
- Connected to an XRF-capable NCP in a 37x5 Communication Controller

Class-2 terminals

Class-2 terminals do not have backup-session support on the alternate IMS, but they are restarted automatically by IMS after takeover.

When a Class-2 terminal session is established on the active IMS, the alternate IMS tracks the session initiation and termination by using log records. When the active IMS terminates abnormally, the alternate IMS tries to establish a new session with the network resources that were active before the failure.

When IMS tries to reestablish a session with a terminal that has no available network path, the attempt fails. Before IMS can reestablish a session on a locally attached terminal, the operator must switch the terminals to the new active IMS. If the complex has many of these kinds of terminals, you might consider giving the operator control before the takeover actually proceeds. At this time, the operator could perform the switch, allowing a faster recovery for sessions on these terminals. To give the operator this control, use the AUTO parameter in member DFSHSBxx in IMS.PROCLIB.

See the table in “Specifying terminal support” for terminals that qualify as class-2.

If you have an application terminal that communicates with a programmable controller in an XRF complex, you might need to add code to the communications portion in order to achieve the highest level of transparency available. Additionally, for the ISC terminal, you must define a session for the link that connects the ISC terminal to an XRF IMS.

Class-3 terminals

Class-3 terminals do not have backup session support on the alternate IMS. Their terminal sessions must be restarted manually on the new active IMS after takeover, either by the MTO or by user logon.

All IMS terminals are eligible for class-3 service. Class-1 and class-2 terminals that have indicated no switching should be done are eligible for class-3. This is done by specifying BACKUP=NO on the system definition macro or ETO logon descriptor.

Class -3 terminals use the existing procedures to connect with the new active IMS after takeover. If no path to the new IMS systems exists, reestablishing the sessions might not be possible. All LU 6.2 conversations on the failing IMS must be reallocated after a takeover.

Related concepts:

➡ Terminals in an XRF complex (System Administration)

Related reference:

“Recovery modes for XRF terminals” on page 408

“Coding ETO descriptors” on page 100

➡ Macros used in IMS environments (System Definition)

Specifying terminal support

The following table summarizes the keywords you can use to change the priority or terminal type of various devices.

All terminals that IMS supports can be used in an XRF complex. The terminal characteristics determine the type of support a terminal can have in an XRF complex, with class-1 having the best support and class-3 having no support.

Table 65. Specifying terminal support

Device	Class	BACKUP keyword on TERMINAL macro or ETO logon descriptor
SLUTYPE1 (3286) SLUTYPE2 (3278,3279,...) SLUTYPEP (4700,8100, SERIES 1,...) FINANCE 3601	Eligible for class 1,2,3	For class-1 service, BACKUP=(1-7,YES). For class-2 service, BACKUP=(1-7,NO). For class-3 service, BACKUP=NO.
LUTYPE6 (ISC) MSC (VTAM) NTO (leased line) 3270 (VTAM) SPOOL line groups Locally attached devices	Eligible for class 2,3	For class-2 service, BACKUP=(1-7,NO). For class-3 service, BACKUP=NO.

Table 65. Specifying terminal support (continued)

Device	Class	BACKUP keyword on TERMINAL macro or ETO logon descriptor
LU 6.2	Eligible for class 3	Not Applicable

The primary and the secondary master terminals are special types of terminals that exist simultaneously on both the active and alternate IMS systems. They are not treated as class-1 or class-2 terminals. The master terminals on the alternate IMS are not disrupted by a takeover. After takeover, the master terminal on the alternate IMS processes the undelivered messages from the active IMS's master terminal.

The BACKUP keyword on several macros in the system definition statement or ETO logon descriptor specifies the actual support a terminal is to receive. Your installation determines the switching priority of the terminals, with level 1 being the lowest priority and level 7 the highest. Although IMS orders VTAM requests by priority, the requests might be completed in any order because of unpredictable factors in the network, such as pacing, line speeds, and congestion.

If BACKUP=(1-7,YES), the terminal is automatically taken over without losing session control. If the terminal is connected to an XRF-capable VTAM, this constitutes class-1 support.

For both class-1 and class-2 terminals, the default switching priority is 4.

If BACKUP=NO, this terminal must be restarted manually after takeover. This forces class-3 support.

If options on the BACKUP parameter are set incorrectly, IMS resets the incorrect option at session initialization. For example, if you specify BACKUP=(1-7,YES) for an ISC terminal, IMS will reset YES to NO because the highest class of support that an ISC terminal can receive is class-2.

XRF with USERVAR and the BACKUP parameter

For XRF systems with USERVAR, BACKUP=(1-7,YES) also tells the alternate IMS to maintain backup sessions for all class-1 terminals. The maximum number of terminal backup sessions that can be established on the alternate IMS is specified in the NCP BUILD statement with the BACKUP=n option. If the (n+1) class-1 terminal user tries to log onto the active IMS, the logon is denied.

If BACKUP=(1-7,NO), this terminal is automatically restarted after takeover, but no BACKUP session is established. This forces class-2 support.

Planning for Rapid Network Reconnect (RNR)

Rapid Network Reconnect (RNR) automatically reconnects IMS VTAM terminal sessions across outages (IMS, z/OS, CPC, or VTAM) and subsequent IMS restarts on the same or another CPC within an IMSplex. The use of RNR can provide greater availability of VTAM sessions and eliminate the need to clean up sessions and restart IMS after an outage.

Note: The following topics apply only to RNR systems and should not be confused with XRF systems, which also use VTAM multinode persistent sessions (MNPS).

Specifying levels of support

RNR support can be specified at three levels: IMS execution parameters; or ETO logon descriptor `OPTIONS=` parameter; or by session using the Logon exit routine (DFSLGNX0).

RNR support is specified in the DFSDCxxx PROCLIB member by indicating `RNR=ARNR`, `RNR=NRNR`, or `RNR=NONE`. `ARNR` activates RNR and specifies that a basic session reconnect will be performed automatically by IMS during the /START DC process following an IMS or VTAM restart. `NRNR` specifies that automatic session reconnect will be activated, but not used unless overridden by the logon descriptor or the logon user exit routine. Unless overridden, the session will be terminated following restart of IMS or VTAM. `NONE` will not activate RNR at all.

If RNR is specified without a parameter (`RNR=` without `ARNR` or `NRNR`), `NRNR` is the default. If RNR is not specified at all (no `RNR=`), RNR will not be activated just as if `RNR=NONE` was specified.

RNR support can also be indicated on the ETO logon descriptor using the `OPTIONS= ARNR|NRNR` parameter. The RNR support specified at this level overrides that set at the system level using IMS execution parameters. If IMS was initialized with RNR unspecified, then specification of RNR support on the ETO logon descriptor is ignored.

RNR support can be indicated on the Logon exit routine (DFSLGNX0) using the `LGOPT=LGOARNR` parameter. The RNR support specified at this level overrides that set using the ETO logon descriptor or the system level execution parameters. If IMS was initialized with RNR unspecified, then specification of RNR support on DFSLGNX0 is ignored.

Using the DFSDCxxx PROCLIB member, you can also specify the maximum time for session persistence following an IMS or VTAM failure. The parameter for specifying this time is `PSTIMER`. `PSTIMER` can be set from 1 to 86400 seconds. The default setting is 3600 (1 hour). If 0 is indicated, then no timer is used and session persistence will continue indefinitely.

Note: If the VTAM START option for MNPS, `HPRPST`, is set to a lower value than `PSTIMER`, then VTAM START overrides the `PSTIMER` setting.

Changing levels of support

RNR support can be turned on or off between sessions or between IMS restarts, and it can be changed at any level: in the IMS PROCLIB member, in the ETO logon descriptors, and in the Logon exit.

After RNR has been set on or off, it is not applied until the next session cold start. For IMS terminals (non-ISC, Finance and SLUP), the result is that the last RNR option selected becomes active at each logon. For ISC, Finance, and SLUP terminals, the result is that the last RNR option selected spans multiple logons and IMS restarts and becomes active only at the next session cold start.

Persistent Session Tracking

VTAM Persistent Session Tracking is provided for both single-node persistent sessions (SNPS), and multinode persistent sessions (MNPS).

The level of VTAM persistent session support desired for IMS is specified on the APPL definition statement using the `PERSIST=MULTI|SINGLE` parameter.

VTAM SNPSs have the following characteristics:

- Reconnect must be on the same CPC as the original IMS.
- Only IMS failures and reconnects are supported.

VTAM MNPSs have the following characteristics:

- Reconnect may be on another CPC in an IMSplex.
- IMS, VTAM, z/OS, and CPC failures and reconnects are supported.

If you use VTAM MNPS, keep in mind that VTAM end nodes must be running in an APPN/HPR (High Performance Routing) network environment. In addition, all VTAM instances must be connected to a Parallel Sysplex coupling facility using the ISTMNPS structure.

APPC Persistent session support is provided by APPC/MVS. However, APPC conversations are not automatically restarted following an outage.

Termination of Persistent Session Tracking

If a session terminates before IMS has closed the VTAM access control block (ACB), VTAM persistent session tracking, and the ability of IMS to reconnect, are terminated on all terminals.

An IMS-initiated shutdown could be executed, for example, by issuing the `/CHECKPOINT FREEZE` or `/STOP DC` or `/CLSDST` commands.

A VTAM session is also disconnected if VTAM persistent session tracking is prematurely terminated due to a non-IMS-initiated session close if that session close occurred before an IMS restart completion, which can happen for one of the following reasons:

- Remote operator-initiated session termination; executing, for example, the `/RCL` command
- Network operator-initiated session termination; executing, for example, the `VARY NET` or `INACT` commands
- Timeout of persistent session tracking indicated by the `PSTIMER` value

If IMS is defined for RNR on MNPS, the VTAM Network Operator can determine whether persistence is currently active by using the `VTAM DISPLAY ID` command.

When shutting down VTAM, the `HALT NET` or `HALT NET, QUICK` commands terminate session persistence, and the `HALT NET, CANCEL` command maintains session persistence.

IMS shutdown and RNR

When IMS is shut down using the MVS `MODIFY` command for IMS, RNR quickly terminates the IMS network, and VTAM persistence is retained.

After the /START DC command is issued, IMS then automatically reconnects all sessions specified with ARNR that were still active at ABEND; sessions that were specified with NRNR are terminated.

When IMS is shut down using the /STOP DC command followed by the /CHECKPOINT command, VTAM terminals are disconnected for each session, resulting in no automatic session reconnect following an IMS restart.

Using RNR with XRF or VGR

RNR is not supported on a system configured with XRF. Attempting to activate RNR on an XRF-configured system produces a warning message indicating that execution will continue with RNR disabled.

RNR can be used in conjunction with VGR; however, IMS support for RNR takes precedence over VGR support. If a /START DC command is entered during an IMS restart, and RNR=NRNR, then the session is terminated and VGR, if active, performs the appropriate level of affinity and terminal status management. If a /START DC command is entered during an IMS restart, and RNR=ARNR, then the session is scheduled for reconnection and a new session with another IMS cannot be established by invoking VGR. If IMS is cold-started, then all active VTAM sessions are terminated and all active VGR affinities and statuses are deleted.

Restriction: RNR does not support OTMA connections.

Terminal reconnect protocols

The reconnect protocol used for terminals with RNR activated depends on the kind of work in progress and the types of terminals in use at the time of the outage.

The following table shows the reconnect protocol used for each type of terminal following a session outage. All terminals are assumed to have RNR activated. If a session cannot be reconnected, error message DFS2050 or DFS2055 is sent to the Master Terminal Operator.

Table 66. Rapid Network Reconnect protocols

Terminal type	Connect protocol	Terminal subtype	Message generated
SLU1	SNA CLEAR	• Static sessions	DFS3650 ¹
	SNA SDT	• ETO printers	
		ETO non-printers	DFS3649
SLU2	SNA CLEAR	Static sessions	DFS3650
	SNA SDT	Static sessions	DFS3649
NTO	SNA CLEAR	Static sessions	DFS3650
	SNA SDT	ETO sessions	DFS3649
SLU0	SNA UNBIND	Static sessions	DFS3650
(Non-SNA 3284 and 3286)	SNA BIND	ETO sessions	DFS3649
SLU0	SNA CLEAR	ETO sessions	No message
(Finance and SLUP) ²	SNA STSN		
	SNA SDT		

Table 66. *Rapid Network Reconnect protocols (continued)*

Terminal type	Connect protocol	Terminal subtype	Message generated
ISC	SNA UNBIND		No message
(Primary Half-Session)	SNA BIND		
	SNA STSN		
	SNA SDT		
ISC	SNA UNBIND		No message
ETO w/ ALOT=0 ³	UNBIND		No message
	UNBIND		

Notes:

1. Message DFS3649 indicates that signon is required. Message DFS3650 indicates that no signon is required.
2. Reconnect is available only for Single-Node Persistent Sessions.
3. Signon data must be supplied as logon user data or by the DFS LGNX0 exit.

Signon security

Depending on whether RNR is activated or not and the types of terminals you are using, signons or logons of VTAM sessions may be required following IMS restarts.

If RNR is not activated, then both logon and signon of sessions are required. If RNR is activated, then the following types of terminals require signon after a session reconnect:

- SLU0 (non-printer, 3270, non-SNA)
- SLU1 (non-printer only)
- SLU2
- NTO

If RNR is activated, then the following types of terminals sign on automatically after a session reconnect:

- SLU0 (Finance/3600)
- SLU0 (SLUP)
- SLU1 (3284, 3286, non-SNA)
- ISC (LU 6.1)

Chapter 25. Defining the network

Performing a system definition for an IMS terminal network includes several tasks.

The tasks include:

- Defining the names of the terminals
- Defining the terminal device types
- Specifying optional parameters, such as buffer sizes

The IMS network consists of static and dynamic terminals (ETO and LU 6.2 devices). You do not need to define your ETO or LU 6.2 devices to IMS.

Although only static terminals in the IMS network are defined by system definition macros, both static and dynamic terminals must be defined to VTAM. You must coordinate several aspects of the IMS requirements with parameters that are specified during VTAM network generation. The following topics highlight those coordination activities, but do not attempt to relate the separate IMS system definition in a detailed way with the series of VTAM and NCP definitions.

Related concepts:

 Defining terminals with data communication macros (System Definition)

Preparing for the operational network

In establishing the network, you must incorporate the IMS system definition requirements into the corresponding VTAM generations. You should also track the overall progress and the hardware installation schedules.

The principal activities involved in establishing the IMS network consist of:

- Gathering the IMS requirements for physical terminals.
In addition, gather the following information regarding terminal requirements:
 - Organize input data for stage-1 IMS system definition for static terminals.
 - Define ETO descriptors and display screen characteristics for dynamically allocated terminals.
 - Define LU 6.2 descriptors.
 - Use the terminal profiles as part of the documentation for the IMS system design.
- Matching the terminals that are specified for IMS with those available for wider use by the installation.

Multiple users can exist on high-function terminals, and the terminals might not be dedicated to the IMS online system. You need to be able to identify these terminals to the network generation personnel and to identify expected usage by the end users.

- Matching IMS system definition parameters to their counterparts in network generation.

Correct device function might require the matching of IMS system definition parameters with VTAM and NCP network parameters. Be especially careful when defining buffer sizes. IMS and network buffer sizes must be compatible. Although system definition in IMS for ETO and LU 6.2 devices are not necessary, you must define them to VTAM.

- Monitoring the hardware installation plan.
Create a plan for tracking the progress of hardware installation and network generation to ensure that the change to production mode is not jeopardized. This tracking might involve interacting with hardware specialists who install and maintain terminal hardware.
- Anticipating terminal installations.
You might want to predefine terminals to IMS in anticipation of their installation.

Related tasks:

Chapter 6, “Administering the Extended Terminal Option,” on page 87

Part 2, “CPI Communications and APPC/IMS,” on page 35

Coordinating IMS definition and network definition

Provide input to VTAM and NCP regarding the network requirements described in the following topics.

Using IMS as a host subsystem

You must define the IMS control region to VTAM as an application. Your installation needs to decide on a name. This name is used in the definition of application nodes that are added to the SYS1.VTAMLST system data set.

Example: The following is an example of a control region definition:

```
IMS  APPL AUTH=(ACQ),PRTCT=password
```

This example shows that a password is required in order to start communication with the subsystem. Use the PASSWD keyword on the COMM macro to specify, for IMS system definition, the corresponding password. If the VTAM definition does not use IMS as the name for the IMS subsystem, code the chosen subsystem name in the COMM macro for IMS, using the APPLID keyword. The default APPLID name is IMS.

Defining VTAM nodes

The local node names are also added into SYS1.VTAMLST.

You need to take special care that the exact terminal names used in the IMS system definition are repeated in the local list members of that data set. Other characteristics of the terminal are specified in the VTAM entries, and you must ensure that no inconsistencies exist between IMS and VTAM parameters. Associated with the node name are:

- The use of program function keys or selector pen
- Keyboard characteristics
- Screen sizes
- Model number
- LU type
- Transmission service level

IMS verifies screen size, model, LU type, and transmission service for dynamic terminals. If these are not correct, a session is not established.

Recommendation: If you are a DC administrator, assure that the VTAM definitions are correct. These definitions are used to select ETO descriptors and MFS format characteristics. IMS system definition sets these definitions for static terminals and ignores VTAM definitions.

Related concepts:

 Defining VTAM terminals (System Definition)

Estimating VTAM storage requirements

Examine the full set of terminal options that are planned for the IMS network and select those communication options that might have an effect on the VTAM storage requirements.

For example, choose the appropriate:

- Protocol (BID or NOBID)
- Type of response patterns
- Number of active nodes

Using your terminal profiles, you can respond to the need for system programming information to estimate storage requirements for IMS support.

Determining VTAM buffer pool values

Provide the VTAM system programmer with IMS input and output message sizes. The values you select for the number and size of receive-any buffers, as specified on the RECANY keyword of the COMM macro, are of special importance.

You must specify the largest number of receive-any buffers that are required in order to support the VTAM network. You can override the size with the execution parameter, RECASZ, and the number with the execution parameter, RECANY. Session initiation fails for terminals that are added to the system if they require a larger receive-any buffer size than the size that is currently specified. You must restart the IMS system, specifying the increased buffer size, before the terminal can establish a session with IMS.

Determining the NCP buffer pool values

The NCP buffer pool sizes and thresholds are based on the VTAM buffer pool information.

Work with the system programmer to assess the volume of predicted traffic for application programs. No specific IMS requirements exist for NCP system definition parameters.

Determining static and dynamic terminal signon requirements

All VTAM-defined terminals can sign on by providing user data with the session initiation request, regardless of whether signon is required for the terminal.

You cannot sign on a user to more than one terminal simultaneously, unless you specify SGN=G, M, or Z on the startup procedure to allow multiple signons. Otherwise, IMS issues message DFS3649A or DFS2467I.

A user cannot enter signon data from an output-only device through logon or a /SIGN command. When signon data is omitted at session initiation for an output-only device that requires a signon, message DFS2085I is sent to the MTO.

Multiple signons

You cannot use the same name for dynamic LTERMs as for static LTERMs that are defined during system definition.

You can allow users to sign on concurrently to one or more terminals, static or dynamic, by specifying SGN=M in either the IMS or DCC procedures.

The IMS master terminal receives a security violation message, DFS286, for rejected signon attempts. If you do not want to receive the message, set SECCNT to 0 on the COMM macro, the MSGEN macro, the SECURITY macro, or with the SECCNT initialization EXEC parameter.

Session status message—DFS3650I

When signon is achieved or if signon is not required, message DFS3650I is sent to the user indicating the status of the session with IMS.

Exception: The following terminals do not receive these messages:

- Autologon terminals
- SLU-1 terminals running in unattended mode
- ISC
- SLU P
- 3600/Finance

In addition, if NOTERM is specified on the TERMINAL macro or on the ETO user descriptor, no DFS3650I message is received.

The information presented on the DFS3650I message panel shows whether the user is in conversation mode. This status panel also shows whether user output security exists for the terminal.

Related tasks:

Part 3, “Extended Terminal Option (ETO),” on page 77

Related reference:

-  COMM macro (System Definition)
-  MSGEN macro (System Definition)
-  SECURITY macro (System Definition)

Related information:

-  DFS3650I (Messages and Codes)

Checking requirements for LOGON MODE tables

When a static VTAM terminal is connected to IMS and a session is initiated, VTAM needs to know the LOGON MODE identifier.

A VTAM LOGON MODE table contains the SNA bind parameters for the session and the identifier (the LOGMODE parameter value, which is used at logon) for that terminal. Also, a default identifier for each SDLC-connected terminal exists.

These considerations apply to static VTAM terminals.

Coordinating LOGON MODE identifiers

When you are deciding on the LOGON MODE identifiers, you need to:

- Match the naming conventions that are being used for the VTAM network in your installation. Your choice of a name might need to follow this convention.
- Arrange for automatic starting of IMS-remote terminals, or have them controlled by the VTAM operator. Otherwise, choose names that are meaningful to a remote terminal operator, if in order to initiate a session the remote terminal operator (RTO) enters the VTAM LOGON command.
- Choose names that are helpful to the master terminal operator if sessions are to be started with the /OPNDST command. If a terminal can be operated with different SNA protocols, the name should help distinguish between the operating modes.
- Use the MODETAB keyword on the TERMINAL macro or ETO logon descriptor to specify the default identifier that is to be used in the following situations:
 - The default identifier is not used.
 - The terminal operator does not specify the identifier in the LOGON command.
 - The master terminal operator does not specify the identifier in the /OPNDST command.

Examples

The code below illustrates the nature of the LOGON MODE table entries that are described to VTAM for a 3270 SDLC connection or a 4730 device. IMS defines these as SLU 2 and SLU P respectively.

```
SE3270  MODETAB
IBMS3270 MODEENT LOGMODE=S3270,FMPROF=X'02',TSPROF=X'02',      X
          PRIPROT=X'71',SECPROT=X'40',COMPROT=X'2000'
MT4730  MODEENT LOGMODE=MT4730,FMPROF=X'04',TSPROF=X'04',      X
          PRIPROT=X'B1',SECPROT=X'B1',COMPROT=X'6080',          X
          RUSIZES=X'0000'
MODEEND
```

The operator can enter either the VTAM LOGON command or use the VTAM USSTAB command for installation-determined logon syntax.

The following example specifies a KEYWORD=*value* format:

```
U3270  USSTAB
      USSCMD      CMD=IMS,REP=LOGON
      USSPARM     PARM=APPLID,DEF=IMS
      USSEND
      END
```

Logon requirements for the master terminal

Because IMS automatically logs on the master terminal after the IMS START command is issued, you must code the MODETAB keyword on the LU statement for the physical device that is used as the MTO if the VTAM default is not to be used.

When the MTO is using the /OPNDST command, a MODE operand applies to all node names presented in the command.

Related tasks:

Part 3, “Extended Terminal Option (ETO),” on page 77

Specifying initial VTAM configurations

Based on your IMS requirements, start lists (ATCSTRyy) and configuration members (ATCCONxx) are placed in the SYS1.VTAMLST library. This information should reflect those terminals that are to be activated at VTAM startup, and those VTAM nodes that are known to VTAM at startup.

For ISC (LU 6) devices and MSC VTAM links, if the session is started from another subsystem, the session is not automatically recovered. Session failure messages are sent to the MTO indicating that you must manually restart the session.

You should emphasize flexibility in the IMS requirements, and try to reduce the requirements for a large part of the network to be immediately available. This might entail more elaborate startup instructions for start lists and configuration members. This might also be reflected in your MTO instructions and the extent to which the /OPNDST command is used. The delayed start up of network sections can overlap with IMS processing in order to allow IMS to start up more quickly. Also, the network startup sequence can be executed by an automated operator program executing in IMS.

Using SON/COS support in IMS

Session outage notification (SON) and class of service (COS) are facilities of VTAM and SNA that allow IMS to recognize a session outage.

IMS attempts to restart the session for recoverable session failures. Both SON and COS must be specified in VTAM to be available for IMS use. In addition, the ASR option must be specified for the IMS node using the system definition process or the /CHANGE command.

The IMS SON and COS support are available for all the VTAM SNA terminals and MSC VTAM links. When multiple virtual routes exist between IMS and the other LU, use an alternate route to avoid a network outage.

The SON facility enables VTAM to inform IMS that a session failure has occurred. The SON facility gives installations the option of allowing IMS to automatically restart sessions (without end-user intervention) if a recoverable session failure occurs in the network.

You need to code SONSCIP=YES on the VTAM APPL definition statement for the IMS application in order to activate automatic session restart after a session outage. If you decide not to include the SONSCIP definition for IMS, VTAM and IMS proceed with session termination without automatic session restart.

The COS facility allows VTAM to control selection of actual routes by designating a set of virtual communication routes based on speed of traffic, data type, and security considerations. These sets of routes are defined in the VTAM COS table. The mode table entry that is used for session establishment can specify a COS entry name.

When the set of session BIND parameters associated with the terminal contains a COS name, these specified virtual routes are scanned to determine which one can be used for the session. If the selected virtual route fails, restart of the session causes the COS set to be scanned again for use of a different virtual route for session reestablishment.

For program LUs (as well as Finance, LU P, MSC, and ISC), in-flight messages can be recovered and retransmitted as necessary, thereby preserving the integrity of the message exchange across the restart.

For device LUs, in-flight messages might be lost or duplicated.

SON and COS support preserves the session setup options (BIND parameters), as well as the class of service associated with the failing session by using the same mode table entry name in order to reestablish a new session.

Restrictions: Because the boundary node NCP cannot distinguish an upstream route failure from a failure of a host CPC, SON and COS support cannot be used for terminals in XRF configurations that have backup sessions.

Related reference:

 /CHANGE commands (Commands)

Starting an IMS network

To make IMS ready to receive VTAM logon (session initialization) requests, use the IMS /START command with the DC keyword.

Before a session with IMS can be established, VTAM and NCP must be active.

The /START DC command tells VTAM to pass to IMS any queued VTAM logon requests to IMS, as well as logon requests for any logical unit that is defined to VTAM as belonging to IMS.

The /START DC command activates the following processes:

- Initiates IMS data communication processing
- Opens the VTAM access method control block
- Enables the IMS VTAM Logon exit routine

Any logon requests VTAM receives before the IMS /START DC command but after the IMS VTAM access method control block (ACB) has been opened are queued in VTAM until the /START DC command is completed. If VTAM is active when IMS is initialized, the IMS VTAM ACB is opened. If VACBOPN=DELAY has been specified, then the VTAM ACB open is delayed until the /START DC command is entered.

Use the /START APPC command to start APPC/IMS.

Session initiation

A *session* is the logical connection between a logical unit (such as a terminal) and a VTAM application program (such as IMS). A session must be established before data can be transmitted between a logical unit and IMS.

Session initiation is requested in one of five ways:

- The terminal operator enters the LOGON sequence. VTAM verifies the command and passes the request to IMS.

The terminal operator can identify IMS in the logon sequence with the following names:

- The APPLID name, if the terminal operator is requesting a session with a specific IMS.

- The MNPS ACB name, if the terminal operator is requesting a session with an XRF with MNPS system.
- The USERVAR, if the terminal operator is requesting a session with an XRF with USERVAR system.
- The generic resource name, if the terminal operator is requesting a session with a generic resource group.
- The z/OS VTAM network operator requests session initiation on behalf of the logical unit by using the VTAM VARY command with the LOGON option. VTAM processes the request and passes it to IMS.
- VTAM passes a logon request to IMS for each logical unit that is defined to VTAM as belonging to IMS.
- The IMS master terminal operator requests session initiation for a logical unit by entering the IMS /OPNDST command.
- ETO autologon initiates the session and supplies the appropriate user data, based on user descriptors and exit routines.

For LU 6.2, session initiation occurs automatically as conversations are allocated. The IMS MTO can request delivery of queued LU 6.2 output by issuing the /ALLOCATE command. When using LU 6.2, remember that the IMS application name (which matches the LU name) for LU 6.2 conversations is different than the LU name IMS has for non-LU 6.2 types, and that the /START DC and /START APPC commands work independently.

Regardless of how session initiation is requested, identical processing occurs when IMS receives the request.

Related concepts:

“Sharing printers between systems” on page 455

“Sharing printers using ETO” on page 120

IMS transaction types and transaction states

Transactions are the most common type of data that is sent from a logical unit to IMS.

IMS supports two kinds of transactions—update and inquiry.

Definitions:

- An *update* transaction can modify a database.
- An *inquiry* transaction can look at data in a database, but it cannot change or update it.

You define transactions as update or inquiry during IMS system definition.

An additional attribute is defined for inquiry transactions—recoverable or irrecoverable.

Definitions:

- *Recoverable-inquiry transactions* are always recoverable, regardless of which element in the network fails.
- *Irrecoverable-inquiry transactions* are not recovered after an I/O error condition occurs or at IMS system restart.

All update transactions are recoverable. All Fast Path transactions must be defined as recoverable, but they can be either inquiry or update.

An LU 6.2 transaction is recoverable only if the asynchronous protocol is used to initiate the transaction. IMS conversational transactions from LU 6.2 programs are not recoverable in the local system, but they are recoverable across the MSC link.

IMS treats an irrecoverable transaction in the same manner as a recoverable transaction, except that all processing that required to achieve recoverability is eliminated. As a result, irrecoverable transactions require less processing time, but they could be lost in the event of a network failure (for example, line failure, CPC failure, or queue failure).

Transactions created for queuing only by the Destination Creation exit routine (DFSINSX0) have a status of DYN. The only purpose of a queue-only transaction is to queue a message to the shared queues. Queue-only transactions are not recovered at restart, unless they are stopped, or were not yet checkpointed.

Determining transaction states

A transaction can be in any one of a number of states.

The possible transaction states are:

LOCK A transaction in LOCK state does not receive messages. Scheduled messages containing this transaction code are stopped.

Restriction: The /LOCK TRANSACTION command cannot be used with Fast Path-exclusive transactions, but it can be used with Fast Path-potential transactions.

MODSTOPPED

A transaction in MODSTOPPED state cannot receive input, because online change processing is in progress. A /MODIFY COMMIT command sets this status. A CPI Communications driven transaction cannot be marked MODSTOPPED.

PSTOP

A transaction in PSTOP state cannot be scheduled; however, the transaction continues to be processed until the limit count is reached. If the limit count is large, the processing interval is long. To ascertain the status of the transaction, use the /DISPLAY command; to alter the status of the transaction, use the /ASSIGN command.

PURGE

A transaction in a PURGE state has had its input messages stopped.

QSTOP

A transaction in QSTOP state cannot be entered during the time between the completion of a /MODIFY PREPARE command and the completion of the corresponding /MODIFY COMMIT or /MODIFY ABORT command. Transactions that are known to be affected by the content of an online change are rejected; that is, the transactions are to be changed or deleted, or they could access databases or programs that are to be changed or deleted. By using the /DISPLAY MODIFY command, you can cause a list of transactions that are affected by a current online change to be displayed. At the terminal, such a transaction is rejected with message DFS3470. If the

transaction uses a Fast Path routing code that is changed or deleted, the rejection message is DFS3471. A CPI Communications driven transaction is never in a QSTOP state.

STOP A transaction in STOP state is stopped. The queuing and scheduling of messages that are destined for a transaction or class of transactions are stopped. However, output can still be queued if it originates from an application program.

USTOPPED

USTOPPED is a status value that is set when an application program attempts to use an IMS DL/I database resource that is not available, and the program terminates abnormally with abend U3033. The USTOPPED condition is not set for a CPI Communications driven program. The CPI Communications driven application program is abnormally terminated with an abend U0125.

Restarting the network after remote takeover

Unlike XRF, which automatically reestablishes data communications with class-1 and class-2 terminals, RSR remote takeover requires that you be involved in restarting the terminal network at the new active site.

Your new active site (the old tracking site) should already have a proper system definition, including all ETO descriptors, terminal definitions specified in the TERMINAL macro, and all necessary VTAM definitions. Then, after remote takeover, you must restart the terminal network.

Before starting non-VTAM terminals, you need to physically switch all telecommunication lines and terminals to the new active subsystem. Because all non-STSN terminals have no knowledge of their work that was done immediately prior to the remote takeover, you receive no indications at the terminal or from IMS whether the remote takeover causes message data to be lost.

VTAM terminals must be disconnected from the old active site before they can log on to the new active subsystems at the new active site.

Finance terminals, SLU-P terminals, ISC terminals, and MSC links can cause IMS to recognize a lost-data condition after a remote takeover. If IMS recognizes this condition, the session resynchronization attempt fails, the MTO is notified, and IMS automatically forces the session or the MSC link to be cold started. Data loss is not recognized if information about the terminal's initial (cold start) logon is also lost. If IMS does not know the terminal was active at the old active site, it cold starts the session at the next logon. The terminal itself can recognize this action and issue messages (DFS2948 in a non-MSC environment, or DFS3311 and DFS3212 in an MSC environment), which you should save for possible manual resynchronization (recovery). From the messages and an associated CVCT log record, you can determine which input and output messages were lost, resulting in the resynchronization error.

An IMS ISC remote primary half session can detect an STSN problem from the RSR secondary half session after remote takeover, but it cannot perform the automatic forced session cold start. The remote primary half session terminates with a DFS2050 message, and the MTO must manually force a session cold start using the following commands:

- /STO NODE *node* USER *user*

- /ASSIGN USER *user* TO VTAMPOOL
- /STA NODE *node*

VTAM MSC links can also be out of synchronization after a remote takeover if link cold start is not automated. The links and their sessions can be resynchronized in the same way as for unrecoverable data transmission errors when using VTAM in a non-RSR environment:

- Stop the link using the /PSTOP command.
- Change the link using the /CHANGE command (SYNCESS to FORCESS).
- Restart the link using the /RSTART command, thus forcing a session cold start.

MTM and CTC links cannot be restarted if significant data is lost, if the sessions are out of synchronization, and if the new active subsystem requires a cold start. If automatic restart fails for SLU-P or Finance devices, you must cold start IMS.

Related concepts:

 Remote Site Recovery overview (System Administration)

VTAM definition for Transport Manager Subsystem

To use the IMS Transport Manager Subsystem (TMS), some VTAM definitions are required, such as APPL definition and mode table definition.

VTAM buffering and pacing

When VTAM runs out of buffers, it returns errors to SEND operations, which result in TMS deallocation. In other words, if pacing causes VTAM to stop accepting data and the TMS buffers for the conversation fill up, log data transport is temporarily suspended. To avoid this TMS deallocation, use pacing (APPL VPACING=) and allocate enough VTAM IOBUFs to support a number of bytes of data equal to:

$VPACING \times \log\_block\_size$

If session pacing is not used, a high-volume component (such as the IMS logger or IMS isolated log sender) will probably consume buffers faster than VTAM can empty them.

VTAM missing interrupt handler

You must be sure to include VTAM Missing Interrupt Handler (MIH) coding in the appropriate VTAM and z/OS definitions. If you do not include MIH in these definitions, and if the CTC link becomes inoperative, MIH continues to appear active but is likely to be hung or unusable. And if the CTC link is inoperative, the TMS is likewise inoperative.

Defining VTAM for Rapid Network Reconnect (RNR)

To use Rapid Network Reconnect (RNR) for VTAM sessions, you must define VTAM for the level of persistent session support desired using the APPL statement for IMS, and assess user security requirements and exposures and define an appropriate level of RNR support for terminals.

Related concepts:

“Planning for Rapid Network Reconnect (RNR)” on page 419

Defining the level of persistent support

VTAM Persistent Session Tracking is provided for both single-node persistent sessions (SNPS), and multinode persistent sessions (MNPS). You must indicate which level of persistent support is needed for RNR.

MNPS allows VTAM sessions to be reconnected to another CPC in a sysplex, if necessary. SNPS requires that VTAM sessions be reconnected to the same CPC they were connected to when the outage occurred.

The level of VTAM persistent session support desired for IMS is specified on the APPL definition statement using the `PERSIST=MULTI|SINGLE` parameter. The default setting is `PERSIST=SINGLE`. If no `PERSIST` specification is entered, `PERSIST=SINGLE` is assumed.

Defining the level of RNR support

RNR support can be defined at three levels: the system level, the terminal level, and the session level.

To use all three levels of control over RNR support, you must:

- Update the IMS execution parameters to activate RNR.
- Update the appropriate IMS ETO Logon descriptors for dynamic terminal support.
- Update the DFSLGNX0 Logon exit support for dynamic override of the RNR option on a session by session basis.

To activate RNR at the system (and default) level, you must specify in the DFSDCxxx IMS.PROCLIB member whether to activate RNR by indicating either `RNR=ARNR` (Activate RNR) or `RNR=NRNR` (No RNR).

Using the PSTIMER parameter in the DFSDCxxx IMS.PROCLIB member, you can also specify the maximum time for session persistence following an IMS or VTAM failure. PSTIMER can be set from 1 to 86400 seconds. The default setting is 3600 (1 hour). If 0 is indicated, then no timer is used and session persistence continues indefinitely.

Note: If the VTAM START option for MNPS, HPRPST, is set to a lower value than PSTIMER, it overrides the PSTIMER setting.

To control RNR support for ETO dynamic terminals, use the `OPTIONS= ARNR|NRNR` parameter in the ETO logon descriptor. The RNR support specified at this level overrides that set at the system level using IMS execution parameters.

To control RNR support on a session by session basis, use the `LGOPT=LGOARNR|LGONRNR` parameter for the Logon exit routine (DFSLGNX0). The RNR support specified at this level overrides that set using the ETO logon descriptor or the system level execution parameters.

The parameters of DFSLGNX0 are documented in its source code.

Chapter 26. Editing and formatting IMS messages

IMS uses two methods to edit and format messages to and from terminals: Message Format Service (MFS) and basic edit routines.

IMS provides samples of user-written exit routines that can be designed to edit:

- Input and output from a terminal
- Transaction codes
- Input message fields
- Input message segments
- Message switching

This topic presents an overview of the advantages of MFS, introduces MFS control blocks for message formatting, and summarizes the characteristics of the different devices MFS supports and the responsibilities of the MFS administrator.

Restriction: MFS does not support LU 6.2 devices or OTMA. The LU 6.2 Edit exit routine (DFSLUEE0) is provided for both input and output messages from LU 6.2 devices when the implicit API support is used. The OTMA Input/Output Edit exit routine (DFSYIOE0) is provided for both input and output messages from OTMA.

Related reference:

-  OTMA Input/Output Edit exit routine (DFSYIOE0) (Exit Routines)
-  LU 6.2 Edit exit routine (DFSLUEE0) (Exit Routines)

Message Format Service

Message Format Service (MFS) is an IMS facility that formats messages to and from terminals, so that IMS application programs need not deal with device-specific characteristics in input or output messages.

MFS formats messages to and from user-written programs in remote controllers and subsystems, so that host application programs need not deal with terminal-specific characteristics of the remote controller.

MFS uses control blocks that the user specifies to indicate to IMS how input and output messages are arranged.

- For input messages, MFS control blocks define how the message that is sent by the device to the application program is arranged in the program's I/O area.
- For output messages, MFS control blocks define how the message that is sent by the application program to the device is arranged on the screen or at the printer. Data, such as literals that appear on the screen but not in the program's I/O area, can also be defined.

In IMS systems, data that is passed between the application program and terminals or remote programs can be edited by MFS or basic edit. The facilities provided by MFS depend on the type of terminals or secondary logical units (SLUs) your network uses.

MFS allows application programmers to deal with logical messages instead of device-dependent data; this simplifies application development. The same

application program can deal with different device types using a single set of logic, whereas device input and output are varied for a specific device type. The presentation of data on the device or operator input can be changed without changing the application program. Full paging capability is provided by IMS for display devices. Input messages are created from multiple screens of data.

A program using MFS need not be designed for the physical characteristics of the device that is used for input and output messages unless it uses very specific device features. Even when these features are used, the program can request that MFS assist in their presentation to the program or the device.

MFS supports SLU-type devices SLU-1, SLU-2, SLU-P, Finance, and LU 6.1. MFS also supports older devices, including IBM 3270 and 3600.

For IBM 3270 or SLU-2 devices, device control characters or orders can be sent directly from or received by the program using the MFS bypass function. This gives the application program more direct control of the data stream. The program uses reserved format names that cause MFS to bypass the edit of:

- the output message
- the next input message that is received from the display terminal

Both logical- and physical-paging facilities are provided for the IBM 3270 and 3604 display stations; these facilities allow the application program to write large quantities of data that MFS can divide into multiple display screens on the terminal. The terminal operator has the capability to page forward and backward to different screens within the message.

MFS components

MFS has several components.

The MFS components include:

- MFS Language utility
- Message editor
- MFS pool manager
- MFS Service utility
- MFSTEST pool manager
- Message Format Service Device Characteristics Table (MFSDCT) utility (DFSUTB00)

The MFS Language utility is executed offline to generate control blocks and place them in a format control block data set named IMS.FORMAT. The control blocks describe the message formatting that is to take place during message input or output operations. They are generated according to a set of utility control statements.




Figure 49. MFS utilities and their output

Note: In MFS test mode, the MFS Language utility can run at same time as the online IMS control region. However, you must use the online change procedure to modify MFS formats when not in IMS test mode.

The following figure shows the MFS online environment. The steps listed following the figure correspond to the numbers in the figure.




Figure 50. Overview of the MFS online environment

1. Input message is sent to the MFS message editor from the terminal.
2. MFS message editor requests pointer to MFS blocks from the MFS pool manager.
3. MFS pool manager checks the message format buffer pool to see if the blocks exist in the pool. If the blocks do not exist, the MFS pool manager reads the blocks from IMS.FORMAT into the buffer pool.
4. MFS pool manager sends the address of the MFS blocks to the MFS message editor.
5. MFS message editor formats the input message for the application program.
6. MFS message editor sends the formatted input message to the message queue to be processed.
7. Application program processes the message and sends the output message to the message queue.
8. Output message is sent from the message queue to the MFS message editor.
9. MFS processes the output message for the terminal just as it processed the input message (steps 2 through 6).
10. The formatted output message is sent to the terminal.

The following figure show the MFS test environment. The steps listed following the figure correspond the numbers in the figure.




Figure 51. Overview of the MFS test environment

Note: You must use the /TEST MFS commands to begin MFS test mode.

1. Input message is sent to the MFS message editor from the terminal.
2. MFS message editor requests pointer to MFS blocks from the MFSTEST pool manager.
3. MFSTEST pool manager checks the communication-line buffer pool to see if the blocks exist in the pool. If the blocks do not exist, the MFS pool manager reads the blocks from IMS.TFORMAT into the buffer pool.
4. MFSTEST pool manager sends the address of the MFS blocks to the MFS message editor.
5. MFS message editor formats the input message for the application program.
6. MFS message editor sends the formatted input message to the message queue to be processed.
7. Application program processes the message and sends the output message to the message queue.
8. Output message is sent from the message queue to the MFS message editor.
9. MFS processes the output message for the terminal just as it processed the input message (steps 2 through 6).
10. The formatted output message is sent to the terminal.

The message editor and MFS pool manager operate online during the normal production mode of operation. The message editor performs the actual message formatting operations using the control block specifications.

Two other MFS components, an MFS Service utility and an MFSTEST pool manager, are available to support optional MFS operations.

The MFS Service utility provides a method for additional control of the format control block data sets. It executes offline and is able to create and maintain an index of control blocks for online use by the MFS pool manager.

The MFSTEST pool manager replaces the MFS pool manager in order to support the optional MFSTEST mode of operation. The /TEST command with the MFS keyword places a logical terminal into MFSTEST mode. For each terminal in MFSTEST mode, combining temporary format blocks with the use of other blocks that are already in production mode allows new applications and modifications to existing applications to be tested without disrupting production activity.

Related reference:

- ➡ MFS Service utility (DFSUTSA0) (System Utilities)
- ➡ MFS Language utility (DFSUPAA0) (System Utilities)
- ➡ MFS Device Characteristics Table utility (DFSUTB00) (System Utilities)

Administering MFS

To take full advantage of the flexible message formatting options offered by MFS and to ensure efficient MFS operation, an MFS administrator should be appointed.

The MFS administrator should be responsible for MFS implementation and administration and should coordinate MFS application design and programming for the installation.

The responsibilities of an MFS administrator include:

- Establishing procedures for the submission of MFS control block requests by application development personnel.
- Establishing procedures and controls for the application of changes to the IMS.TFORMAT library.
- Defining MFS control blocks most efficiently in keeping with the requirements of the specific application and the overall system.
- Minimizing device character transmission, sharing MFS control blocks, and ensuring the most efficient use of MFS without jeopardizing application requirements or operator considerations.
- Establishing and documenting operator guidelines and system standards for MFS. The many options that MFS offers can result in confusing practices, unless you establish and follow standard procedures. Be sure to standardize certain aspects of format design in order to minimize terminal operator training and error rates.
- Deciding if and how the optional index directory should be used and determining buffer pool requirements.
- Monitoring the use of the MFS control blocks and of the MFS buffer pool with the IMS /DISPLAY command and IMS Monitor report output, and modifying MFS parameters as required.
- Making end users aware of the operating characteristics of the different device types and terminal subsystems.
- Informing others about the differences between the various partition formats.
- Establishing and informing others about naming conventions and guidelines. In particular, the MFS administrator should be able to discuss naming conventions for the partition descriptor blocks and the sizes of the display screen, the viewports, and the display characters.
- Communicating information on conventions for and restrictions on MFS formats.
- Defining screen sizes and feature combinations that are not included in the IMS stage-1 system definition.

- Creating the MFS device characteristics table control statements for processing by the MFSDCT utility (DFSUTB00). The MFS device characteristics table entries and default format control blocks are used for ETO terminals.
- Defining input message field edit routines and segment edit routines. MFS and all MFS-supported devices are able to use message edit routines. You can use these edit routines for such common editing functions as numeric validation or conversion of blanks to zeros.

IMS provides a sample of both a field edit and a segment edit routine.

The MFS administrator should be technically competent in all aspects of IMS relative to MFS:

- Online transaction processing
- IMS API for message processing
- Operation with remote controllers
- MFS implementation, device characteristics, and capabilities
- Interpretation of MFS statistics and related IMS Monitor report output

The administrator should also be familiar with the hardware and remote programs for SLU-P, Finance remote programs, or ISC subsystems if such programs are going to operate with MFS by using distributed presentation management.

In addition, because one administrative responsibility is minimizing device character transmission, the administrator should be familiar with the terminal hardware characteristics.

An MFS administrator must communicate with IMS system administrators and application developers, as well as programmable workstation developers and end users. The administrator must be able to enforce installation standards and to modify application specifications for MFS control blocks when necessary to benefit overall system performance. The procedures of related programming groups should recognize this authority of the MFS administrator.

Related concepts:

 Input message field and segment edit routines (Application Programming APIs)

Advantages to using MFS

Two primary advantages to using MFS are that it simplifies the development and maintenance of terminal-oriented application systems and improves online performance.

To simplify IMS application development and maintenance, MFS performs many common application program functions and gives application programs a high degree of independence from specific devices or remote programs.

With the device independence offered by MFS, one application program can process data to and from multiple device types while still taking advantage of their different capabilities. Thus, MFS can eliminate or minimize the changes in application programs when new terminal types are added.

MFS makes it possible for an application program to communicate with different types of terminals without having to change the way it reads and builds messages. When the application program receives a message from a terminal, how the message appears in the program's I/O area is independent of the kind of terminal

that sent it; the appearance depends on the MFS options specified for that program. If the next message that the application program receives is from a different type of terminal, the user does not need to do anything to the application program. MFS shields the application program from the physical device that is sending the message in the same way that a database program communication block (PCB) shields a program from the data in the database and how it is stored.

Other common functions MFS performs include left or right justification of data, padding, exit routines for validity checking, time and date stamping, page and message numbering, and data sequencing and segmenting. When MFS performs these functions, the application program handles only the actual processing of the message data.

The following figure shows how MFS can make an application program device-independent by formatting input data from the device or remote program for presentation to IMS, and by formatting the application program data for presentation to the output device or remote program.




Figure 52. Message formatting using MFS

MFS also improves online performance of a terminal-oriented IMS by using control blocks that are designed for online processing. The MFS control blocks are compiled offline, when IMS is not being executed, from source language definitions. MFS can check their validity and make many decisions offline to reduce online processing. In addition, during online processing, MFS uses look-aside buffering of the MFS control blocks in order to reduce CPU usage and the channel costs of input/output activity.

Because MFS control blocks are reentrant and can be used for multiple applications, online storage requirements are reduced. Optional main-storage indexing and anticipatory fetching of the control blocks can also reduce response time. IMS gains additional performance improvements, because multiple I/O operations can execute concurrently in loading the format blocks from the MFS format library.

In addition, MFS uses z/OS paging services to reduce page faults by the IMS control region task.

Finally, MFS can reduce use of communication lines. Compressing and transmitting only the required data reduces line load and improves both response time and device performance.

Related reference:

➡ Transaction Manager exit routines (Exit Routines)

MFS control blocks

Users can specify four types of MFS control blocks to format input and output for the application program and the terminal or remote program.

Definitions:

- *Message Output Descriptors (MODs)* define the layout of messages that MFS receives from the application program.
- *Device Output Formats (DOFs)* describe how MFS formats messages for each of the devices with which the program communicates.
- *Device Input Formats (DIFs)* describe the formats of messages MFS receives from each of the devices with which the program communicates.
- *Message Input Descriptors (MIDs)* describe how MFS formats messages so that the application program can process them.
- *Message descriptors* are both MIDs and MODs.
- *Device formats* are both DIFs and DOFs.

Because each MOD, DOF, DIF, and MID deals with a specific message, both a MOD and DOF must exist for each unique message a program sends, and both a DIF and MID must exist for each unique message a program receives.

Overview of MFS components and operation

MFS has the following components:

- The MFS language utility, which generates control blocks from user-written control statements and places them in a library called IMS.FORMAT
- The MFS service utility, which is used for maintenance of the control blocks in IMS.FORMAT
- The MFS message editor, which formats messages according to the control block specifications generated by the language utility
- The MFS pool manager, which keeps the MFS control blocks that are required by the message editor in the main-storage MFS buffer pool
- The MFSTEST Pool Manager, which replaces the MFS pool manager when the language utility is being used in test mode

IMS online change also plays an important part in updating the MFS libraries, even though it is not part of MFS. Briefly, online change allows the control block libraries to be modified while the IMS control region is executing.

Related concepts:

 Making online changes (System Administration)

Creating MFS formats with SDF II

SDF II is an interactive tool for designing and generating MFS formats.

SDF II does not replace MFS, but it does make developing and maintaining MFS formats easier. Because SDF II uses a panel editor for designing and testing formats, it frees the MFS programmer from some of the tasks associated with coding MFS source statements.

With SDF II, application programmers and analysts who might not know the special requirements of MFS can perform a part of the programming job that

would otherwise call for specialized knowledge. SDF II uses a panel editor, such as the one shown in the following screen shot, to define and test a panel.

```

                                DEFINE FORMAT

Format . . . . .Positions 1-75 or 80, Lines 1-24 of 24
Marks: V - C . L , S +                               Contents: FORMAT
001
002                                *****
003                                ** E M P L O Y E E   P A Y R O L L **
004                                *****
005
006  LAST NAME:                                FIRST NAME:
007
008  EMPL NO:
009
010  SOC SEC NO:
011
012  RATE OF PAY:
013
014
015
016  INPUT:
017
018
019
020
021
022
023
024
PF1=HELP      2=SPLIT      3=END      4=RETURN      5=RFIND      6=CHANGE
PF7=UP        8=DOWN      9=SWAP     10=LEFT      11=RIGHT     12=CURSOR

```

After testing the panel, the SDF II user can automatically generate the MFS source code, shown in the example code below, that is needed for the format definition. Using MFS only, the programmer would need to code these statements manually.

```

DOF
PAYF      FMT
          DEV      TYPE=(3270,2),FEAT=IGNORE,DSCA=X'00A0'
          DIV      TYPE=INOUT
          DPAGE    CURSOR=((5,15))
          DFLD     '*****',POS=(1,21)
          DFLD     '*  EMPLOYEE PAYROLL  *',POS=(2,21)
          DFLD     '*****',POS=(3,21)
          DFLD     'LAST NAME:',POS=(5,2)
LNAME     DFLD     POS=(5,15),LTH=16
          DFLD     'FIRST NAME:',POS=(5,36)
FNAME     DFLD     POS=(5,48),LTH=16
          DFLD     'EMPL. NO:',POS=(7,2)
EMPNO     DFLD     POS=(7,11),LTH=6
          DFLD     'SOC SEC NO:',POS=(9,2)
SSN       DFLD     POS=(9,14),LTH=11
          DFLD     'RATE OF PAY: $',POS=(11,2)
RATE      DFLD     POS=(11,16),LTH=9
          DFLD     'INPUT:',POS=(16,2)
INPUT     DFLD     POS=(16,9),LTH=30
          FMTEND

MID
PAYIN     MSG      TYPE=INPUT,SOR=(PAYF,IGNORE)
          SEG
          MFLD     'PAYUP '      SUPPLIES TRANCODE
          MFLD     LNAME,LTH=16
          MFLD     FNAME,LTH=16
          MFLD     EMPNO,LTH=6
          MFLD     SSN,LTH=11
          MFLD     RATE,LTH=9

```

	MFLD	INPUT,LTH=30,JUST=R,FILL=C'0'
	MSGEND	
MOD		
PAYDAY	MSG	TYPE=OUTPUT,SOR=(PAYF,IGNORE)
	SEG	
	MFLD	LNAME,LTH=16
	MFLD	FNAME,LTH=16
	MFLD	EMPNO,LTH=6
	MFLD	SSN,LTH=11
	MFLD	RATE,LTH=9
	MFLD	INPUT,LTH=30,JUST=R,FILL=C'0'
	MSGEND	

SDF II is designed for use on a 3270–display device; however, SDF II can create formats for all devices supported by MFS.

Related reading: For more information on SDF II, see *SDF II General Introduction*.

Basic edit

If you do not use MFS, an IMS function called basic edit performs message editing.

For input messages, basic edit:

- Translates messages to uppercase, if specified by the EDIT=UC parameter on the system definition TRANSACT macro.
- Removes leading control characters from the first segment of each message. Leading blanks are also removed from the first segment if the message is not the continuation of a conversation or a message from a terminal in preset mode.
- Removes leading control characters from all subsequent message segments, if the message type is a transaction or a command (except the /BROADCAST command).
- Removes line control characters from all segments.
- Removes trailing carriage return and end-of-block characters from all segments of a transaction.
- Eliminates backspaces, on a one-for-one basis, from all segments when the entering or transmission of backspaces is a normal correction procedure on the entering terminal.
- Removes the password and replaces it with a blank when necessary to provide separation between the transaction code, logical terminal, or command verb, and data that follows.
- Inserts, in front of data that is entered in the first segment of a message, the transaction code or logical terminal name defined by the prior /SET command. A blank is inserted following the transaction code, if it is necessary to obtain separation between the inserted transaction code and the entered data.
- Adds a nonconversational transaction code to the first segment of the next input message, if a terminal is in conversation mode and the application terminates the conversation by inserting a nonconversational transaction code into the SPA.
- Removes the function management header (FMH), if any, that appears at the beginning of the first transmission of a chain for VTAM-supported devices.
- Deblocks message segments at each interrecord separator (IRS) control character, and discards the IRS control character for input from a SLU–1 card reader, a transmit data set (TDS), or a user data set (UDS).
- Deblocks message segments at each new line or forms-feed control character if the optional MFS editing is not selected for SLU–1 consoles. This character can be discarded, depending on the criteria previously described.

- Treats the presence of a TRN (X'35') character immediately followed by a length in the next byte as transparent data.

For output messages, basic edit:

- Changes nongraphic characters in the output message before the data goes to the device.
- Inserts any necessary idle characters after new-line, line-feed, and tab characters.
- Adds line control characters for the operation of the communication line.

For basic edit support of SLU-1 transparent data, basic edit does not alter or delete characters following the destination and password fields if transparent processing has been requested. Indicate transparent processing by specifying:

- BINPDSB1=BINTRNDS on the bind image for VTAM-type SLU-1 terminals
- Edit option BASIC=TRN on the COMPT1 parameter of the IMS TERMINAL macro or ETO descriptors for SLU-1 terminals

Related concepts:

“Conversational transactions” on page 396

IMS editing for Intersystem Communication (ISC)

IMS recognizes several options on a message-by-message basis when communicating with ISC (LU 6.1).

IMS recognizes the following options:

- Use basic edit.

The destination process name (DPN) is BASICEDT, which is a reserved IMS name.

- Treat the data as transparent data.

The IMS preset destination mode, which is established by the /SET command, is used. This option is used if a /SET command is in effect for a given session and a primary resource name (PRN) is not specified on the ATTACH or SCHEDULER FMH. While this option is in effect, IMS cannot recognize input commands within this input data stream.

- Do not edit data that follows the transaction code and input password fields.

The input destination and security checking is the same as in basic edit. A / as the first character indicates an IMS command and causes the IMS preset destination mode to be bypassed or terminated. No additional editing is done on the input message following the input transaction code or LTERM name and optional password field. This is the default input edit for ISC under each of the following conditions:

ATTACH or SCHEDULER FMH is not included in input

ATTACH or SCHEDULER FMH does not specify a DPN

ATTACH or SCHEDULER FMH specifies a DPN but not a PRN

- Do nothing.

IMS uses the input PRN ATTACH parameter as the transaction code or LTERM name for the duration of the single input message. Subsequently, the original IMS preset value can be used again. Basic edit treats the input as transparent with password security already checked.

The chosen option depends on the destination process name (DPN) and the data received by IMS. The DPN is specified on the ATTACH or SCHEDULER FMH.

If a DPN of ISCEDT or its alias is specified the second through fourth options are chosen. The alias for ISCEDT is defined using the EDTNAME keyword on the COMM macro. ISCEDT is also a reserved IMS name.

Related concepts:

“Editing messages” on page 473

Transparency option

Transparent data includes programmed symbols, extended field attributes, and commands.

These unprintable characters might appear on the 3270 screen as hyphens or another representation for unprintable characters.

If IMS attempts to send transparent data to an IBM 3270 terminal that is unable to receive the data, an error is returned to IMS, and the terminal is taken out of service. To restart the terminal, the master terminal operator must issue the /START command.

For LU 6.1, input is processed using the transparency option of basic edit if the primary resource name (PRN) in the function management header (FMH) contains a transaction. Otherwise, the data stream is only scanned for the transaction code and optional password. The remaining data is treated as transparent data.

Related concepts:

 Extended field attributes for output devices (Application Programming APIs)

Unprotected screen option

When the screen is in unprotected status, IMS can send output to the terminal at any time without requiring input from the terminal. This option can be used on a terminal-by-terminal or message-by-message basis.

Bypassing MFS editing

IMS enables an IMS application program to bypass MFS formatting of an output message destined to an IBM 3270 or to SLU-2 devices by specifying reserved format names DFS.EDT or DFS.EDTN. This bypass is intended for subsystems or installation support programs that execute under IMS.

Recommendation: Do not use this bypass for IMS application programs, because the application program loses the productivity, device independence, and migration benefits associated with MFS.

When MFS is bypassed on output, the application program must construct the entire 3270 data stream, beginning with the command code and ending with the last data byte. The user might want to bypass the MFS output edit function to allow application programs to receive the 3270 input data stream from the terminal without MFS or basic editing.

Locking and unlocking the terminal keyboard

When the application uses MFS bypass and the reserved format name DFS.EDTN, IMS enables the application program to lock or unlock the terminal keyboard.

If the lock option is specified, the program is responsible for unlocking the keyboard after processing is completed. If the unlock option is specified, IMS controls the locking and unlocking of the keyboard.

Related concepts:

 MFS message formats (Application Programming APIs)

IMS sensitivity to nongraphic message data

The following topics describe the sensitivity IMS has to specific characters when users attempt to send and receive nongraphic data in IMS messages.

Output message segment editing

For output message segments that MFS edits, only graphic data (X'40' through X'FE') is contained in the output message that is presented to the device. Nongraphic characters, if present in the output message, are changed by MFS before the data is presented to the device.

Device control characters HT, CR, LF, NL, and BS are changed to X'00' for 3270 data streams. For all other device types, all nongraphic characters are also changed to blanks.

If the Distributed Presentation Management (DPM) option of MFS is used for SLU-P or ISC, the user can specify GRAPHIC=NO in the SEG statement. Nongraphic characters, if present in the output segment with GRAPHIC=NO specified, are presented unchanged to the remote program.

For programmable workstations that are supported through VTAM, IMS can insert function management headers (FMHs) and can perform additional editing for device control sequences when splitting a single IMS segment into multiple transmissions.

Editing of input message segments by MFS

If MFS is defined for a device, you should be aware of certain considerations.

The considerations include:

- Specify GRAPHIC=NO in the SEG statement to prevent uppercase translation on a segment if the destination requests it with the EDIT=UC specification on the system definition TRANSACT macro.
- If the first input record is from a 3600, SCS1, or SCS2 device, or from DPM-An, the segment is discarded if the final characters of the segment are:
 - Two asterisks (**)
 - Two asterisks followed by NL (**X'15')
 - Two asterisks followed by IRS (**X'1E')
- The presence of two slashes (//) at the beginning of a message segment is considered an escape sequence.
- If the card feature is defined for an SCS1 device (with the CARD=operand in the DEV statement), all control characters are removed from magnetic card input before the data is presented to the input MFLD.

The definition of the MFS delete characters (LDEL=operand in the DEV statement) and field tab character (FTAB= operand in the DEV statement) for MFS-supported devices, excluding IBM 3270, can direct the editing of input message segments.

If the input is processed by MFS, the editing performed is dependent on the descriptions provided through the Message Format Service language utility. As input segments from the device might have no relationship to input message segments after MFS editing, the input segment from the device is not available to user-written edit routines. Input message segments after MFS editing are available to user-written edit routines.

Related concepts:

 Input message formatting (Application Programming APIs)

Editing of input message segments by basic edit

The following editing is done if basic edit is used for nongraphic message data.

- For the first segment of an input message when a terminal is not in conversation mode, leading characters less than X'41' are removed. For other than the first segment or when a terminal is in conversation mode, leading characters less than X'40' are removed.
- If a terminal is not in conversation or preset mode, a left parenthesis within the first nine positions of the first segment indicates the presence of a password. The left and right parentheses and the password are removed, and the segment is compressed.
- For non-SNA devices, the X'26' character that appears as the final character in a segment is removed.
- Two asterisks (**) or two asterisks followed by NL (X'15') that appear as the final characters of a segment cause the entire segment to be discarded.
- For unbuffered keyboard devices, backspace (X'16') characters are treated as character-delete indicators. Each backspace character and the preceding input character are removed from the segment.
- If the destination of the input message is a transaction, an NL (X'15') character appearing at the end of a segment is removed.
- If a device is in preset mode, the transaction code is added to the first segment.
- For input from IBM 3270 devices, the attention identifier and cursor address are removed, and all start buffer address sequences are changed to blanks.
- If the first character of any segment is a slash (/), the entire input message is treated as a command.
- If an input message is received from an NDS device, or if it is using Intersystem Communication, the data stream is handled wholly as transparent data or transparent except for edit of the transaction code and password.

Controlling output devices

You can control an output device by using control characters in the second byte of the ZZ field of the message prefix (Z2).

Follow these rules:

- SLU-1:
 - The Z2 field, bit X'80', should be set if the segment contains structured field data.
 - The Z2 field, bit X'40', should be set if the segment is the first segment of a new LPAGE series.

- Switched devices:

Exception: IBM 3270

You can use the Z2 field in the message format to request that the line be disconnected after the present message is sent. This field is ignored if the output is physically sent to a device without this capability. The disconnect request is indicated using X'80' as the Z2 field value, this request is recognized if present in any segment.

- IBM 3270 printer (3270P device type):

You can use the Z2 field, bit value X'80', to specify the presence of command code and WCC character in the data stream, which is to be used when MFS editing is bypassed.

- Printer components:

The program can embed, in the text portion of the message, carriage return characters or new line symbols. If the output is going to a local printer (SYSOUT), the first two characters of the message can be carrier control characters.

Small buffer devices

Some terminal devices have hardware limitations of the maximum buffer size that they can process. Additional limits can be imposed by software, either in the device or in the network.

Exceeding these limits can result in an error and failure to deliver a message that is too large for the device to process. Be aware of what these limits are, and of the alternatives available to your application program.

Different device types specify the limits in different ways. For input, the SEGSIZE or BUFSIZE specification in IMS is relevant, and, for output, the OUTBUF specification. For VTAM terminals, the BIND RU size is based on these values. IMS supports message chaining for both input and output, but some terminals do not support chaining. In this case, message length is limited by sizes.

Your application program has control over the length of the output message, either directly or through the MFS definition being used to format the message. If you need to send a message that is longer than the maximum length supported by the target device, you must break the message into multiple messages, each of which is shorter than the maximum length for the device. The DL/I PURG call must be issued by your application program to do this. IMS delivers messages in the order that your application program inserts them; this is a good solution for devices that are not sensitive to message boundaries.

Related reference:

➞ PURG call (Application Programming APIs)

➞ TERMINAL macro (Communications and Connections)

➞ LINE macro (System Definition)

Controlling output

Most output of an IMS online system goes to the input terminal. You have several choices for creating printer output.

- Use local printing through hardware-specific device support.
- Code the application program so that it directs response to the printer component of the input terminal.

- Use the /ASSIGN command to assign the response LTERM to a printer component.
- Code the application program to send output (SYSOUT) to an LTERM representing an IMS system printer.
- Code the application program to send output to an LTERM representing an alternative device for offline printing.

The choice is dependent on the application design. The use of the /ASSIGN command and the control of spool data sets are choices that might involve master terminal operator intervention.

With ETO, the output is associated with the user. The user can move from terminal to terminal. The output follows the user. When the output is available, use an autologon facility to generate a session for output delivery to a specified terminal.

Using a printer component

The printer component for a terminal is often not in continuous operation. The printer might need a remote operator to supervise the paper supplies or special forms.

The /COMPT command is used to make the device *ready* or *not ready* if the remote operator is a VTAM component. The operator can use the /ASSIGN command in order to direct output to a particular component.

Related tasks:

 Using a printer component (Communications and Connections)

Spooled output control

You need to be aware of requirements for IMS spooled output. This is because the output might be needed immediately by the end user or additional output can be impeded if allocated space is exhausted.

Given a spool line group, you need to be aware of how many data sets are used for output. A recommended definition is at least two data sets, one data set printing while the other receives additional input. The line number and PTERM references are also required.

At any time, you can use the /STOP LINE n PTERM nn command to call for the spooled output to be scheduled to a printer.

The action of this command is to close the data set and direct additional output to the next spool data set. If this is the last in the set, output goes to the first. If only one output data set exists at the time the print utility is scheduled, all messages for the LTERM are queued. This could rapidly add to the contents of the message queues. Message queues can fill very quickly, specially when the spool is the secondary master console. The IMSWTnnn procedure to execute the print utility is scheduled automatically by the data-set-full condition. A write error for the data set also starts the procedure.

The printing procedure is tailored to the line group data sets by system definition. You can invoke this procedure at any time by using a /START REGION IMSWTnnn command. If the data sets are not closed, current line output is queued.

Restart can affect spooled output handling. If output data sets were not printed from a prior execution, the /NRESTART command prevents additional output messages. The action of a /START LINE automatically schedules the IMSWTnnn procedure and frees the data set for output from the current system execution.

Using printer components of the IBM 3270 Information Display System

You can define printer components to be part of an IBM 3270 Information Display System. This allows for printed copy of the video output (or input) to be sent to the printer's component.

This support covers the following printer component devices: IBM 3284, 3286, 3287, 3288, or 3289.

Your system definition input indicates whether these are to be connected through a polled-BSC or an SDLC line. The printed copy can be automatically produced or be operator controlled. The input message definition or the control prefix of the output message can cause this automatic printing. The component must be attached to the same 3270 control unit as the display workstation containing the information to be copied—an IBM 3271/3274 or a 3275/3276.

Restriction: A locally attached 3270 does not support the copy function.

Specifying candidate printers

With BSC 3270, using VTAM, requests for copying the screen content to a printer are directed to the first available printer on the same control unit as the screen.

Definition: *First available printer* refers to the sequence of TERMINAL macros for the devices. A predetermined sequence controls the order in which the printer components are selected for message output. The following figure shows two display workstation printer groups.




Figure 53. Sequence of IMS TERMINAL macros

The first printer defined after a display is a candidate printer for that display. Any subsequently defined printer is also a candidate if no display is defined between it and the first candidate. In the preceding figure, if a request is made for a printed copy of Display 1 or Display 2, IMS sends a 3270 copy command to Printer 3, if it is available. If Printer 3 is not available, IMS considers Printer 4. Printer 6 is not a candidate for Display 1 or Display 2, but it is a candidate for Display 5.

The copy function is not permitted to cross 3270 control-unit boundaries. For IBM 3274 and 3276 with SNA protocols, IMS sends the print request to the display, and the printer is selected by the controller. The printer need not be defined to IMS.

Restriction: The 3270R (non-SNA) copy function is not supported for ETO terminals.

Operational considerations

When a request for a copy operation is sent by the operator or an application program, the first available candidate printer is used for output.

The search order corresponds to the system definition sequence. If a printer is stopped, already printing a message, in exclusive status, or not ready, the next candidate is chosen. If a copy request by an operator finds all printers busy, the keyboard is locked until a printer becomes available. Output messages that require the copy function are not sent if a candidate printer is unavailable. The display workstation receives an error message. A retry attempt is made for the message when the error message is cleared from the screen using the message advance function (PA2 key). The format description for a message can also specify a copy action. MFS uses the DSCA operand in the DEV statement to specify a copy action.

Related concepts:

➡ System control area (SCA) and default SCA (DSCA) (Application Programming APIs)

Related reference:

➡ DEV statement (System Utilities)

Sharing printers between systems

If you plan to share the use of a printer other than the master, IMS allows an online IMS system and another subsystem (possibly another IMS system) to alternate their use of the printer.

The IMS support is for VTAM 3270 printers and, in the SNA environment, for LU-1 and LU-4 printers. The /OPNDST NODE Q command simulates a logon to IMS for the requested printer. The printer is automatically acquired when the owning subsystem frees the printer. You can queue output at any time independently of printer availability.

The system definition requirements are for the OPTIONS keyword of the TERMINAL macro. Specify:

OPNDST so that the /OPNDST command is valid for this terminal

SHARE so that the VTAM macro SIMLOGON can be issued by IMS in order to acquire the printer when data is queued for output

RELRQ so that IMS releases the terminal to other VTAM subsystems upon their request

Specify RELRQ on the ETO logon descriptors to have the same features available for ETO terminals.

Each output message that is queued to a shared printer attempts a session.

As an alternative to immediate logon simulation and repeated failures while trying to acquire messages, you can use the Shared Printer exit routine (DFSSIML0). This exit routine monitors terminal (or transaction) status, and takes one of three actions:

- Ignores the request

- Simulates the /OPNDST command
- Queues a transaction for an automated operator program

The Shared Printer exit routine is not called if the printer is connected. You specify the SIMEXIT keyword on the COMM macro to call the Shared Printer exit routine. Your installation must write the Shared Printer exit routine and add it to the library of user routines prior to stage-2 of system definition.

Recommendation: Do not share the secondary master terminal.

Related concepts:

“Designing logical terminal networks” on page 397

Related reference:

 Shared Printer exit routine (DFSSIML0) (Exit Routines)

Part 7. Intersystem Communication (ISC)

These topics introduce Intersystem Communication (ISC) and provide all the information you need to use ISC to connect IMS subsystems with other types of subsystems that support the ISC protocol.

Chapter 27. Overview of Intersystem Communication

ISC is a part of the IMS Transaction Manager. It is one of the ways to connect multiple subsystems. The other means of connection is Multiple Systems Coupling (MSC).

As defined under SNA, ISC is an LU 6.1 session that:

- Connects different subsystems to communicate at the application level.
- Provides distributed transaction processing permitting a terminal user or application in one subsystem to communicate with a terminal or application in a different subsystem and, optionally, to receive a reply. In some cases, the application is user written; in other cases, the subsystem itself acts as an application.
- Provides distributed services. Thus, an application in one subsystem can use a service (such as a message queue or database) in a different subsystem.

SNA makes communication possible between unlike subsystems and includes SNA-defined session control protocols, data flow control protocols, and routing parameters. The functions provided by each of these protocols and parameters are summarized below, and the IMS support for them is described in this topic and its subtopics.

Session control (SC) comprises:

- Initiating sessions between subsystems
- Recovering and resynchronizing sessions, maintaining the integrity of session states and recoverable resources across both session and subsystem failures
- Terminating any or all session paths between IMS and another subsystem

Data flow control (DFC) includes:

- Controlling send and receive protocols within a session.
- Resolving contention for transaction initiation within a session.
- Monitoring error recovery processing.
- Monitoring symmetrical shutdown accomplished using the SNA commands stop bracket initiation (SBI) and bracket initiation stopped (BIS).
- Controlling synchronization of resources. Sync point control ensures that all resources are committed or backed out synchronously.

ISC VTAM routing comprises:

- Using parameters in the SNA-defined function management headers to connect the process required for incoming messages and to route reply messages.

APPC/IMS only supports message switching to LU 6.2 destinations through the DFSAPPC service.

Related concepts:

“DFSAPPC system service” on page 70

“ETO and LU 6.1 (ISC) terminals” on page 129

Related tasks:

Chapter 31, “ISC protocols for VTAM connections,” on page 503

Comparison of ISC and MSC

Both Multiple Systems Coupling (MSC) and Intersystem Communication (ISC) can be used to couple multiple IMS subsystems. Both MSC and ISC enable you to route transactions, distribute transaction processing, and expand beyond the capacity of one IMS system.

MSC is an IMS protocol that enables coupling of IMS systems to other IMS systems only. ISC, however, allows you to connect IMS subsystems with any other subsystem that supports the ISC protocol. This other system can be another IMS, CICS, or a user-written system.

MSC supports three types of links between IMS systems: channel-to-channel (CTC), memory-to-memory (MTM), and VTAM LU 6.1. ISC uses only VTAM LU 6.1.

The following figure compares MSC to ISC.




Figure 54. IMS Multiple Systems Coupling and Intersystem Communication

Both ISC and MSC use VTAM parallel session support. Some key differences exist, however. The following table highlights the major functions of MSC and ISC, and shows the differences in support.

Table 67. Comparing MSC and ISC functions

MSC functions	ISC functions
MSC connects multiple IMS systems only to each other. These IMS systems can all reside in one processor, or they can reside in different processors.	ISC can connect either like or unlike subsystems, as long as the connected subsystems both implement ISC. ISC can couple an IMS subsystem to: • Another IMS subsystem • A CICS subsystem • A user-written subsystem
Communication in the MSC environment is subsystem-to-subsystem.	Communication is between application programs in the two subsystems. The subsystems themselves are session partners, supporting logical flows between the applications.

Table 67. Comparing MSC and ISC functions (continued)

MSC functions	ISC functions
Processing is transparent to the user. That is, to the user, MSC processing appears as if it is occurring in a single system.	Because ISC supports coupling of unlike subsystems, message routing requires involvement by the terminal user or the application to determine the message destination. Specified routing parameters can be overridden, modified, or deleted by Message Format Service (MFS).
Unless MSC-directed routing is used, routing is automatic based on system definition parameters. The terminal operator or application program does not need to know routing information.	ISC provides a unique message-switching capability that permits message routing to occur without involvement of a user application.
MSC supports the steps of a conversation to be distributed over multiple IMS subsystems, transparent to both the source terminal operator and to each conversational step (application).	ISC supports the use of MFS in an IMS subsystem to assist in the routing and formatting of messages between subsystems.
MSC does not support the use of the Fast Path Expedited Message Handler (EMH).	ISC supports the use of Fast Path Expedited Message Handler (EMH) between IMS subsystems.

Related tasks:

Chapter 40, “Administering Multiple Systems Coupling,” on page 695

IMS facilities available to ISC

A variety of IMS facilities are available to Intersystem Communication (ISC).

Distributed transaction processing

When supporting distributed transaction processing, IMS can be either the front-end or back-end processor. A front-end subsystem originates communications traffic (transactions, commands, or message switches) that are to be processed by the back-end subsystem. A message is typically the result of data entered at a terminal. However, the message can also result from processing that occurred in an application within the front-end subsystem. The back-end subsystem processes input messages from another subsystem.

In the most typical configuration, IMS is the back-end system and processes IMS transactions that another subsystem enters.

In addition, special support exists for front-end or back-end system utilization is provided by the Front-End Switch exit routine.

Sometimes, transactions entered into an IMS front end are sent with an ISC message switch to the other subsystem. In this case, the terminal operator or MFS provides the message routing information. The transaction can also be sent as alternate program communication block (PCB) output, in which a user-written application program or MFS format definition provides the routing information.

The fact that IMS is a queued subsystem is a key factor in determining the ISC functions that it supports, particularly when IMS acts as a front-end processor. Messages created by a terminal operator or an application in a front end IMS are

queued for transmission to the receiving subsystem and sent asynchronously with respect to the terminal or application that sent the message. A terminal attached to an IMS front-end is not held in response mode during the receiving (back-end) subsystem's processing, unless it uses the Front-End Switch exit routine (DFSFEJ0).

Distributed services

IMS supports distributed services by providing remote access to an IMS message queue. Although IMS does not support distributed DL/I calls, a DL/I database in one subsystem can be updated by another subsystem; this action requires the sending subsystem to invoke an updating application in the receiving subsystem.

IMS transaction types

IMS supports the following transaction types:

- Recoverable-update (includes Fast Path)
- Recoverable-inquiry
- Irrecoverable-inquiry

IMS execution modes

IMS supports the following execution modes:

- Response mode (includes Fast Path)
- Conversational mode
- Exclusive mode
- Transaction preset mode
- Test mode
- Non-response or non-conversational mode

IMS editing facilities

Messages transmitted on the ISC session and processed within an IMS subsystem are edited by the following editing facilities:

- ISC edit (the default editor)
- Message Format Service (MFS)
- Basic edit

These editors can be selected on a message-by-message basis.

ISC message integrity

Message integrity is provided to prevent loss or duplication of input or output messages between IMS and another half session during session restart.

Message integrity and recovery are increased by log write-ahead (LWA). This facility is invoked during system definition by an option on the TRANSACT macro. LWA ensures that sync point information is written to the log (and thus available to IMS restart procedures) before IMS acknowledges the message. More information on message integrity can be found in the Resynchronizing sessions and Sync point and response requirements topics.

ISC security

IMS security facilities control access to IMS resources by another subsystem. Before implementing ISC security in IMS, examine the data protection facilities of the subsystems and associated operating system components. Some important ISC-environment security factors are described there.

Terminal operator verification and authorization are provided by the subsystem controlling that terminal connection.

Both IMS terminal and password security can be defined for static ISC terminals by using resource access control facility (RACF). If password security is defined, a password must be provided with the input message.

ISC in an XRF complex

ISC VTAM sessions are handled as class-2 or class-3 terminals. If the active IMS system fails, all LU 6.1 sessions are terminated. The alternate (now active) IMS system reestablishes the sessions after it takes control of the complex. Users might be without service for a short time.

Related concepts:

“Establishing a connection with an XRF complex” on page 505

“Planning for security” on page 410

“ETO and LU 6.1 (ISC) terminals” on page 129

“Resynchronizing sessions” on page 507

“Relationship of ISC and IMS execution modes” on page 475

 [Extended Recovery Facility Overview \(System Administration\)](#)

 [IMS security \(System Administration\)](#)

Related tasks:

Chapter 26, “Editing and formatting IMS messages,” on page 437

“IMS transaction types and transaction states” on page 432

Related reference:

“Sync point and response requirements” on page 527

 [Front-End Switch exit routine \(DFSFEJ0\) \(Exit Routines\)](#)

Sample system configurations

These figures illustrate sample IMS Intersystem Communication (ISC) configurations. The first two figures use IBM CICS Transaction Server for z/OS as an example of an ISC node.

Existing or new IMS transactions invoked from CICS

The following figure illustrates ISC's distributed transaction processing capability—the ability to invoke existing or new IMS transactions from a CICS application or from a CICS application on behalf of a terminal attached to CICS (using transaction routing).




Figure 55. Existing or new IMS transactions invoked from CICS

In figure above, ISC provides transaction-to-transaction capability, because an application program can be written as two complementary transactions: one executing in an IMS system, the other in a CICS system. User functions can be distributed between systems as required.

IMS MFS DPM mapping function distributed to CICS's BMS

The following figure illustrates ISC VTAM's ability to distribute device mapping function from an IMS Message Format Service (MFS) system to a CICS Basic Mapping Support (BMS) system. MFS partially maps the data stream that is sent to a CICS application. The application is responsible for processing the input data stream to a form that is acceptable to BMS. BMS can then be used to complete the device mapping.




Figure 56. IMS MFS DPM mapping function distributed to CICS's BMS

IMS-to-IMS with ISC VTAM

The following figure illustrates how two IMS systems can be connected using ISC VTAM links. Interconnection of two IMS systems allows a new or modified transaction (TRAN1) in one IMS system to invoke an existing transaction (TRAN2) in another IMS system. Using MFS, replies can be routed to a new transaction (TRAN3) in the initiating IMS system, or to a new instance of the originating transaction (TRAN1). Without MFS, replies are treated as message switches and routed to the source terminal.




Figure 57. IMS-to-IMS with ISC VTAM

CICS-to-IMS MSC using ISC

The following figure illustrates how a CICS subsystem can communicate with an IMS subsystem that has MSC links to other IMS subsystems. In this case, MSC is used to couple the IMS subsystems, while the CICS-to-IMS session is supported by ISC. The CICS subsystem has a single-system view of the multiple IMS subsystems with which it can communicate. MSC distributes the load between the multiple IMS subsystems.




Figure 58. CICS-to-IMS MSC using ISC

Related reference:

ISC between IMS and CICS

The following table identifies the ISC functions that are provided by CICS and IMS, and the IMS facilities that are available to users of CICS START or RETRIEVE, and SEND or RECEIVE.

Table 68. Facilities available to CICS START RETRIEVE and SEND RECEIVE USER.

Function	Front-end system	Back-end system
Distributed Transaction Processing	CICS ¹	CICS ¹
Initiate remote transaction	CICS ¹	IMS non-response type
	IMS all transaction types ²	CICS ¹
	IMS all transaction types except conversational ²	IMS all transaction types except conversational ²
Distributed Transaction Processing	CICS ³	CICS ³
Application-to-Application Conversation	CICS ³	IMS Response Conversational Fast Path
Distributed Presentation Management	IMS MFS map	IMS MFS map
Using distributed transaction processing	CICS user application ⁴ IMS MFS map	IMS MFS map CICS user application ⁴

Note:

1. Using CICS START or RETRIEVE commands.
2. If the transaction in the processing subsystem is a response mode or Fast Path transaction, a nonresponse transaction type is required in the originating system to handle the reply from the processing subsystem.
3. Using CICS SEND or RECEIVE commands.
4. The user application can optionally invoke BMS to aid in mapping functions between CICS and the terminal.

Related concepts:

Chapter 35, “How IMS and CICS use the ISC interface,” on page 607

Terminal device-dependent data

ISC, when used with IMS, is a technique that allows programs in various subsystems (IMS, CICS, and user-written) to exchange messages with application programs in other subsystems. To minimize or eliminate the need to modify these programs, the messages that are exchanged should not contain device control characters that are unique to the originating or destination terminal.

The device control editing facilities of IMS (MFS) or CICS (BMS) should be used to remove these device control codes prior to the initiating program creating its ISC message. Likewise, the program that sends the response to a terminal should use the formatting service of its own subsystem to properly format the message before sending it. This way, all inter-program message exchanges can be independent of the originating and destination terminals. This means that each terminal's owning system should understand all formats for transactions that can be entered or received by terminals on that system.

Passing CICS data to IMS

Users want their front-end application programs to receive 3270 data streams from an attached terminal and to simply pass through this data stream (device control characters) to the back-end IMS.

The resulting output from IMS should then be a 3270 data stream that could be passed through the front-end application and out to the 3270.

Pass-through input to IMS

It is possible to send a 3270 data stream from a front-end application to IMS. However, IMS does not recognize the ISC session as a 3270 device and does not edit the device control characters from the data stream. Your IMS application program must provide this editing function.

Output from IMS

Similarly, IMS does not place 3270 device control characters in an output data stream that is destined for an ISC LTERM, because the ISC LTERM is not a 3270 device. Your application program must then construct a 3270 data stream to be sent to a front-end application that can, with minimal editing, send this data stream to an attached terminal.

This effort needs to be repeated for each device and each application program, and again after any change to a device format.

Chapter 28. VTAM facilities used for ISC connections

VTAM controls the physical transmission of data between IMS and a logical unit.

Both VTAM applications (such as IMS) and other subsystems can be viewed as VTAM logical units. A *logical unit* is an addressable resource, such as an application program or a subsystem. A logical unit can also be a component of a general-purpose terminal system.

The VTAM concepts and facilities used by IMS that are particularly relevant to ISC include:

- Connection, disconnection, and establishing logon mode.
- Messages and responses.
- Request definite response 1 (RQD1) or request definite response 2 (RQD2) and the associated definite responses (DR1 or DR2). Note that definite response 1 and definite response 2 have been separated and redefined for LU 6.1 protocols. RQD2 requests and DR2 responses are now known as sync-point requests and responses and are functionally independent of those responses associated with DR1.
- Sequencing and chaining.
- Orderly session disconnection from stop bracket initiation (SBI) and bracket initiation stopped (BIS) commands.
- Facilities for ensuring orderly communication, including the use of brackets and change-direction indicators.
- Sequence number recovery.
- Receiving input and sending messages.
- Conditional bracket termination.
- Extended Error Recovery Procedure (EERP).
- Use of parallel sessions between ISC nodes.
- Support for both negotiable and nonnegotiable session bind parameters.

IMS also supports the class-of-service (COS) and session-outage-notification (SON) facilities.

Related reading: For more information about the communication concepts and facilities that govern data transmission between a VTAM application program (such as IMS) and another subsystem, see *z/OS Communications Server: SNA Programming*.

The ISC system programmer and system analyst must be familiar with these concepts and facilities in order to design and implement a communications interface between IMS and a remote subsystem using SNA LU 6.1 protocols.

Related concepts:

“Resynchronizing sessions” on page 507

“Using SON/COS support in IMS” on page 430

Related reference:

“Symmetrical session shutdown for LU 6.1 (SBI and BIS)” on page 558

VTAM commands and indicators

VTAM commands and indicators (communication control information) are necessary for data transmission between an IMS application program and another VTAM logical unit.

The following table shows which VTAM commands and indicators that IMS sends and receives during an IMS session (X=supported). Results are unpredictable if any unsupported commands or indicators are used.

Table 69. VTAM commands and indicators sent and received by IMS.

VTAM command or indicator	IMS sends as Primary half session, Receives as secondary half session	IMS sends as Secondary half session, Receives as primary half session
Independent of session		
Initiate session		Note 4 on page 471
Procedure error	Note 1	Note 1
Terminate session		Note 5 on page 471
Session control commands		
Bind	X	
Bind response		X
STSN	X	
STSN response		X
Start data traffic (SDT)	X	
SDT response		X
Unbind	X	X
Normal (synchronous) flow indicators		
Begin bracket (BB)	X	X
End bracket (EB)	X	X
Change direction (CD)	X	X
Commands		
BID	Note 3 on page 471	X
CANCEL	X	X
CHASE	X	X
Logical unit status (LUSTATUS)	X	Note 2
Ready to receive (RTR)		X
Bracket initiation stopped (BIS)	X	
Expedited (asynchronous) flow commands		
Request shutdown (RSHUT)	X	Note 2
Signal	X	X
Stop bracket initiation (SBI)		X

Notes:

1. Sent by VTAM.
2. Not sent by IMS, but optionally sent by the other subsystem and received by IMS when IMS is a primary half session.

3. Not sent by IMS, but optionally sent by the other subsystem and received by IMS when IMS is a secondary half session.
4. Supported by internal VTAM forms of CINIT (LOGON exit).
5. Supported by internal VTAM forms of CTERM (NS and LOSTERM exits).

Using the VTAM application programming interface

The VTAM application program interface (API) used by IMS for ISC consists of macro instructions and control blocks.

The macro instructions and control blocks allow IMS to:

- Establish connection or disconnection.
- Request and control the transfer of data between IMS and another logical unit.

To provide application-to-application communication, the VTAM API provides both a primary and a secondary session application interface. The primary interface is essentially the same as that used by IMS to support other VTAM nodes. However, the interface introduces some new macros and some additional parameters on existing macros.

The VTAM API supports single sessions as well as parallel (multiple, simultaneously active) sessions between two logical units. The IMS interface for this support is primarily through the session qualifier fields of CINIT and BIND, which are available through the LOGON and SCIP exit routines. Single-session support requires static definition of message queues during IMS system definition; parallel-session support permits message queues to be dynamically allocated at session initiation.

For more information about VTAM, see *z/OS Communications Server: SNA Resource Definition Reference*.

Related concepts:

 [VTAM interface considerations \(System Administration\)](#)

Specifying logon modes when establishing a connection

When establishing a connection, IMS requires that session parameters define the rules a logical unit must follow when communicating with IMS. These session parameters are contained in the VTAM mode table.

The use of the default logon mode table entry can be overridden in one of two ways:

- At system definition, the logon mode entry can be specified on the `TERMINAL` macro using the `MODETBL=keyword` or an ETO logon descriptor.
- The `MODETBL` keyword on the `/OPNDST` command can specify the logon mode table entry to replace the table entry defined at system definition.

Related reading:

- For information on specifying the `MODETBL` keyword, see *IMS Version 11 System Definition*.
- For information on specifying the `/OPNDST` command, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

If neither method is used to specify a logon mode table entry, VTAM uses the first entry in the default logon mode table, unless it is overridden by VTAM node mode table entry definition. If a logon mode table entry is specified, VTAM searches the default or user-specified logon mode table for the specified entry.

IMS examines the set of session parameters within the indicated VTAM mode table entry and overlays only those parameters on which IMS has dependencies. The remaining bytes are not changed. IMS ignores any unformatted user data transmitted with the session parameters.

Related reading: For more information on establishing logon mode tables and defining logon mode table entries, see *z/OS Communications Server: SNA Resource Definition Reference*.

Related reference:

-  Bind parameters for SLU P and LU 6.1 (System Programming APIs)
-  Macros used in IMS environments (System Definition)
-  /OPNDST command (Commands)

Design considerations for secondary logical units

When IMS and another subsystem establish an ISC session, one session partner is the bidder and the other the first speaker (the contention winner).

In an IMS ISC session, the bidder is always the primary logical unit (the bind sender), and the first speaker is always the secondary logical unit (the bind receiver). When providing a system design for a secondary logical unit that is to communicate with IMS within an ISC network, consider the following points:

- The first speaker (secondary logical unit) must resolve bracket contention and be prepared to receive interrupts (for example an VTAM command or bracket indicator) at all times. If the first speaker is not prepared to do so, a deadlock can result (for example, if both the primary logical unit and the secondary logical unit use the VTAM SEND macro, and issue POST=RESPONSE rather than POST=SCHED).
- When terminating a session using the VTAM TERMSESS macro, the system designer should be aware of the differences involved in issuing CONDITIONAL versus UNCONDITIONAL. Sending CONDITIONAL permits the primary logical unit to clean up and control session termination. Sending UNCONDITIONAL causes immediate termination of the session by VTAM prior to notification of the primary logical unit. If the primary logical unit is in a suspended state as, for example, when waiting for a response, the system designer should be aware that UNCONDITIONAL session termination can have unanticipated results.

The term “secondary logical unit” is synonymous with “first speaker” and the term “primary logical unit” with “bidder”.

Chapter 29. IMS facilities affected by ISC

The following topics explain how IMS facilities are affected by ISC processing.

Editing messages

Within IMS, the communication interface is transparent to application programs.

In support of this, IMS provides some common editing that:

- Appends protocols and function management headers to outbound messages
- Strips function management headers from incoming messages and moves the messages to appropriate points within the IMS control blocks

Subsequent to this common editing, messages received by IMS are edited by ISC edit, Message Format Service (MFS), or basic edit in accordance with the input ATTACH or SCHEDULER DPN parameter. Messages sent by IMS are edited by Message Format Service before they are sent. The edit functions of ISC edit and basic edit are unique. However, the associated protocols are the same.

- ISC edit is the default editor. When ISC edit is requested and that request does not contain the SNA-defined PRN parameter following the input data, the message is edited with partial input data transparency. That is, only the IMS-required destination code and optional password, which must occur at the beginning of the message, are subjected to editing by ISC edit. The balance of the input message is not edited prior to processing unless the transaction is defined to translate to uppercase. Full input data transparency is provided through ISC edit using the SNA-defined PRN parameter on the ATTACH or SCHEDULER FM header. This enables the IMS transaction code to be external to the message and enables IMS to receive input and route it to its proper destination (terminal or transaction) without examining the input message itself. Because the transaction code is external to the data stream, the optional password is not available to the session. Therefore, password security should not be defined. If password security is defined, an error results.
- Message Format Service's Distributed Presentation Management (DPM) option divides responsibility for message formatting between MFS and a program residing in another subsystem. When using DPM-Bn, physical terminal characteristics are not defined to MFS. Instead, IMS sends MFS-formatted data to a program in the other subsystem. That program must complete the formatting, if necessary, and present the data to the physical device. DPM also allows user involvement in specifying input and output ISC message routing parameters. You can also use the PRN parameter on the ATTACH or SCHEDULER FM header to specify the IMS destination externally, relative to the resulting MFS-formatted input data stream.
- Basic edit provides standard editing for terminal input. It is used to edit device and operator control characters and can be used in an ISC session when it is important to maintain input compatibility for existing applications. Protocols for basic edit are the same as those for ISC edit. Use of the ATTACH or SCHEDULER FM header PRN parameter has no effect when basic edit is used. The message destination is determined from the first field of the input data stream.

IMS permits basic edit, ISC edit, or MFS to be invoked in order to handle communication of transactions, commands, and message switches. MFS DPM facilities are defined as optional on a component-by-component basis. Any of these available processes are selectable on a message-by-message basis by using the ATTACH or SCHEDULER FM header.

ISC input and output can be edited by most IMS data communication exit routines. The exceptions are the MSC-related exit routines, and the hardware-required routines.

Related tasks:

Chapter 26, “Editing and formatting IMS messages,” on page 437

Related reference:

“FM headers for message routing” on page 484

“Using FM headers to invoke ISC edit” on page 561

 How MFS formats input messages (Application Programming APIs)

Issuing IMS commands from an ISC session

Although available to the ISC session, IMS operator commands are primarily intended for use by appropriately authorized operators of IMS master terminals and remote terminals that are directly attached to an IMS system.

Terminal operators of other (non-IMS) subsystems do not communicate directly with IMS, but rather with the control program or user-provided applications within the other subsystem. That subsystem must determine and provide procedures for its terminal operators and application programmers to follow when communicating with IMS. If an application in a non-IMS subsystem is permitted to issue IMS commands through the ISC session, IMS dependencies are introduced into that application. An ISC half session cannot be defined as the primary or secondary master terminal.

Effects on parallel sessions

Include the USER keyword when using authorized IMS commands with the NODE keyword.

Equivalent action is taken against all active, or potentially active, parallel sessions of the indicated node when a specific parallel session instance is not identified to IMS by the USER keyword.

For example, issuing the /CLSDST NODE x command without specifying USER schedules termination of all active parallel sessions of NODE x. Issuing the /STOP NODE x command without specifying USER schedules termination of all active parallel sessions of NODE x and also prevents any new parallel sessions from being initiated on NODE x.

Using IMS test mode for ISC VTAM sessions

You can test ISC VTAM communication protocols and editing facilities by putting a back-end IMS system into test mode.

While in test mode, IMS receives an input message, checks input protocols and FM header parameters, performs input and output editing, inserts appropriate output protocols and FM header parameters, and returns the message to the front-end

subsystem. Noncommunication error analysis, transaction code verification, and transaction scheduling are bypassed. Also, IMS does not send asynchronous output while in test mode.

You can place one or all of the ISC sessions within a specific subsystem in test mode when the IMS /TEST command is received on the ISC sessions. The /TEST command can only be entered by the terminal that is being put into test mode. You can terminate test mode by sending an end-bracket (EB) indicator on an SNA LUSTATUS or CHASE command. Test mode can also be terminated when an /END command is received on the ISC session or is entered locally in the back-end IMS processor by an authorized terminal operator.

In an IMS-to-IMS environment, use the ISC message switch facility to place the back-end IMS system into test mode, to send the message and receive the echo reply, and to terminate test mode. This provides an efficient means of testing without requiring a user application in either subsystem.

IMS control block storage on ISC parallel sessions

The storage required for the IMS control block structure representing potential sessions is greater than the storage required for control blocks representing input and output message queues.

Installations needing only a few logical units type 6.1, each having a relatively small number of associated parallel sessions, but each requiring that a large number of ISC users (subpools or LTERM sets) be dynamically allocated to them, can have lower storage requirements. For static ISC terminals, these users (subpools) are still statically defined. The IMS control block structure for ISC single session (statically defined and allocated LTERMs) is the same as for other static VTAM terminal types within IMS.

You do not need to define a maximum number of ETO-ISC sessions. You can continue to add sessions until you run out of storage and processor capacity. Therefore, the SESSION= keyword is not one of the supported keywords on the ETO descriptor. To define the maximum number of sessions for static ISC sessions, use the SESSION keyword on the TERMINAL macro.

Related reference:

Chapter 36, "ISC data flow control examples," on page 653

Relationship of ISC and IMS execution modes

Because the terms "synchronous" and "asynchronous" have slightly different connotations within IMS and ISC, the following topics explain the relationship of these execution modes.

External specification of execution modes

Messages in ISC VTAM use the SNA-defined function management headers. ISC messages that are routed by VTAM can be processed either synchronously or asynchronously, as determined primarily by the type of FM header and secondarily by the VTAM bracket protocols that are sent with the message.

The following table provides a summary of supported VTAM protocols. Using FM headers and VTAM bracket protocols to establish the mode of the message (synchronous or asynchronous) can be considered to be external to the message.

Table 70. Processing mode requested by FM headers

FM header with VTAM bracket protocol	Synchronous	Asynchronous
ATTACH with CD	X	
ATTACH with EB		X
ATTACH without either EB or CD	X	
SCHEDULER with CD	X	X
SCHEDULER with EB		X
SCHEDULER without either EB or CD	X	X

Messages that are received by IMS with the ATTACH FM header are processed:

- Synchronously, if the session state is left “in-brackets” after the message is received.
- Asynchronously, if either of the following occurs:
 - The session state is left “between-brackets” after the message is received.
 - The message is sent with EB on the first- or only-in-chain.

During synchronous processing, unsolicited or asynchronous output is not sent.

All messages that are received by IMS with the SCHEDULER FM header are processed asynchronously by the receiver with respect to the session.

Related concepts:

“Handling IMS response mode or conversational output errors” on page 523

Related reference:

“Function management headers” on page 559

Internal definition of execution mode

The IMS internal execution mode determines how IMS processes a transaction.

Response mode and conversational mode, for example, are used to ensure synchronism between a given input message and its associated reply. Although a transaction uses an IMS input and output message queue between the session and the application, these modes ensure that associated messages are always processed synchronously with relation to the source session. While these messages are being processed, IMS does not send asynchronous or unsolicited output on that session.

Internal IMS definitions are either synchronous or asynchronous.

The following modes are synchronous:

- Response mode
- Conversational mode
- Fast Path
- Commands
- Test mode

The following modes are asynchronous:

- Nonresponse mode
- Nonconversational mode
- Message switch

Resultant processing mode during ISC VTAM communications

During ISC VTAM communications, when the relationship between the externally requested execution mode and the internally understood processing mode are consistent, the message is processed exactly as requested. In the case in which the two specifications (internal versus external) are not consistent, the message's execution with respect to the session is synchronous.

In either case, if the generated reply is returned on the same session as the request, it is sent using the same type of FM headers that were received with the input message. If the reply is returned on a session other than the one on which the input message was received, it is considered unsolicited asynchronous output and is sent with the SCHEDULER FM header. If the subsystem that is to receive the reply does not support receipt of the SCHEDULER header, the reply is sent with an ATTACH header that terminates the bracket at the end of the IMS message. Therefore, except for a nonlast conversation, nonlast MFS demand-paged messages, or test mode output, ATTACH is used instead of ATTACH SCHEDULER and EB is indicated on the first or only chain of the message. A Nonlast conversation, nonlast MFS demand-paged messages, and test mode output are sent carrying ATTACH with CD indicated on the first or only chain of the message.

However, when a message is sent from the front-end subsystem to the back-end subsystem that is externally requesting synchronous processing, no assumptions should be made by the front-end subsystem as to the timing of the output reply or the availability of any other output for that session. As a result, requests and replies cannot be correlated within the front-end subsystem, even though execution in the back-end subsystem might be synchronous. This is summarized in the following table.

The special support for front-end and back-end system utilization Front-End Switch exit routine provides.

Traffic between two IMS systems is always sent in asynchronous format. If the internal definition of the transaction within the back-end system is synchronous, it is processed synchronously and the ISC session is held for synchronous output. However, this same relationship does not apply for the source terminal. The front-end IMS does not hold that terminal in response mode until the reply is received from the back-end subsystem, unless you are using the Front-End Switch exit routine.

Table 71. Internal versus external execution mode specification.

Internal specification (system definition)	External specification (FMH)	
	Synchronous ¹	Asynchronous ²
Synchronous	Synchronous	Asynchronous ³
Asynchronous	Synchronous ^{4, 5}	Asynchronous

Notes:

1. Synchronous specification does not include ATTACH with EB indicated on the first or only chain of the message.
2. Asynchronous specification includes ATTACH with EB indicated on the first or only chain of the message.
3. Synchronous operation in IMS is allowed in this case, because the SNA-defined external asynchronous format allows IMS as the message receiver to execute an inbound message according to its own interpretation. Thus, IMS's internally

synchronous definition overrides the external asynchronous format. If, however, IMS's internal (system definition) mode of execution is synchronous and the message is received with ATTACH EB, IMS generates an error message and terminates the session.

4. This support allows response mode to be established dynamically (on a message-by-message basis), even though the message was internally defined to be asynchronous (nonresponse mode).
5. Messages cannot be externally defined as synchronous when `OPTIONS=NORESP` is specified internally on the IMS `TERMINAL` macro or ETO user descriptors. This conflict between system definition and FM header intent causes session termination.

Traffic between two IMS systems is always sent in asynchronous format. If the internal definition of the transaction within the back-end system is synchronous, it is processed synchronously and the ISC session is held for synchronous output. However, this same relationship does not apply for the source terminal. The front-end IMS does not hold that terminal in response mode until the reply is received from the back-end subsystem, unless you are using the Front-End Switch exit routine.

Related reference:

 Front-End Switch exit routine (DFSFEBJ0) (Exit Routines)

LTERM users (subpools) and components

IMS user blocks are sets of IMS logical terminals (LTERMs) defined by the `SUBPOOL` macro during IMS system definition or dynamically created from ETO user descriptors.

Subpools defined at system definition cannot be used with ETO LU 6.1 terminals. Users (subpools) defined for ISC are separate from the users defined for dial-type terminals and are permitted only in conjunction with ISC parallel session support.

Definition: The collection of all static ISC users is known as the *VTAMPOOL*. LTERMs defined within the *VTAMPOOL* are reassignable only between users within the *VTAMPOOL*. LTERMs not defined within the *VTAMPOOL* cannot be assigned into the *VTAMPOOL*.

ISC users are dynamically assigned to an ISC session instance as a result of session initiation. That is, these parameters define the user that is available to the given session instance. This user remains allocated to the named parallel session instance even across session and subsystem failures until released through normal termination by mutual agreement of IMS and the other subsystem. For a single nonparallel session, the allocation of a set of LTERMs is fixed during system definition.

Each IMS LTERM is associated with one input and one output IMS component. The input and output components can be the same component, or different components can be specified. Conversely, IMS does not prevent multiple input or output LTERMs from being associated with a single component. However, doing so can cause problems with input component determination or output presentation.

IMS uses the input component ID to identify the LTERM that is to be associated with the input message. For other terminal support, IMS assumes that all input is from the first LTERM in the list that passes the necessary operational and security

checks. However, for input from an ISC node, the input component is determined based on the component that is indicated in the ATTACH FM header. If no FM header is received, IMS assumes the input is to be associated with the LTERM for input component (ICOMPT) one. After the component value is determined, if the associated LTERM cannot be found, is stopped, or is not ready, or if that LTERM cannot pass security checks, the message is rejected.

When output is sent, it is sent to the component (COMPT) identified with the output LTERM. Message switches, broadcast messages for specific LTERMs, and data replies from transactions are directed to the component that is associated with the specified output LTERM. The user-written MPP can insert to the I/O PCB and default to the output component that is associated with the selected input LTERM. It can also address specific components through the appropriate LTERM name by an insert to an alternate PCB.

You can establish proper relationships between input and output components through the NAME macro during IMS system definition or LTERM keyword on the ETO user descriptor. This enables a logical unit to indicate its input component and causes output to be returned to the associated output component that was indicated during IMS system definition or on an ETO user descriptor. Proper definition and use of components can reduce or eliminate the need for LTERM naming conventions, DL/I CHNG calls, and inserts to alternate PCBs.

Recommendation: Incorrectly specifying the message delete system definition parameter, MSGDEL, or the ETO user descriptor while defining ISC users (subpools) can prevent a session from being initiated. An ISC session can be initiated only if the MSGDEL specifications on the TERMINAL or ETO user descriptor and SUBPOOL macros match.

Also, when using the ISC message switch support or an alternate PCB insert from a local transaction in order to route messages to another subsystem and return replies to a source terminal operator, MSGDEL=SYSINFO must be specified on the IMS system definition TERMINAL macro or ETO logon descriptor associated with both the ISC session and the source terminal. Specifying MSGDEL=NONIOPCB for the ISC session prevents ISC-message routing to other subsystems. Specifying MSGDEL=NONIOPCB for the source terminal prevents message replies from being routed to the source terminal operator.

The output bracket and send and receive protocols used by IMS depend on a combination of the system definition output component specifications, the protocol used to input the message to IMS, and the type of IMS message received. For example, IMS synchronous message types have a predefined protocol for output replies. Associated with the output component is a definition of send and receive protocols for asynchronous output.

Related concepts:

“IMS use of routing parameters” on page 485

➞ Defining VTAM terminals (System Definition)

➞ ETO descriptors (System Definition)

Related reference:

“Bracket protocol for IMS output” on page 536

➞ VTAMPOOL macro (System Definition)

➞ SUBPOOL macro (System Definition)

 **TERMINAL** macro (System Definition)

Chapter 30. Designing communications using the ISC protocol

The following topics provide an overview of subsystem-to-subsystem communications that use the ISC protocol.

Determining output protocols

Traffic on the ISC session is controlled by the VTAM protocols associated with the message.

The primary bracket protocols used are:

- **Begin bracket (BB)**
Signals the beginning of a bracket. Begin-bracket is an unconditional request when issued by the first speaker (secondary half session) and a conditional request when issued by the bidder (primary half session).
- **End bracket (EB)**
Signals the end of a bracket, places the session in contention state, and allows either session partner to request a new bracket.
- **Change-direction (CD)**
Turns control of the session over to the session partner and allows the session partner to send traffic on the session.

Traffic across an ISC session can be sent and received synchronously to the session. The processing of messages is performed either synchronously or asynchronously to the session. That is, the sending subsystem initiates a message and waits for a reply (synchronous) or does not wait (asynchronous).

Within ISC, the external synchronous mode of operation (as specified by FM headers) requires transmitted messages and associated replies, if applicable, to occur within a single bracket. The end of synchronous mode is signaled by EB. When in synchronous mode, IMS must be able to reply, if required, within the same bracket (before EB is received) to any outstanding message traffic with the exception of a message switch. If the input message protocol prohibits IMS from sending any required replies within the same bracket, the input message is rejected.

Within IMS, the protocols used for output that must be sent with ATTACH (because the other subsystem lacks SCHEDULER support) are predefined, regardless of whether the originating input transaction and resulting output replies are synchronous (ATTACH) or asynchronous (SCHEDULER). This output includes:

- Nonlast chains (pages) of MFS paged output
- The last (MFS paged) or only chain of response mode, conversation mode, test mode, and IMS command replies
- Asynchronous output

You must define to IMS the protocols to be used with the last (MFS paged) or only chain of other asynchronous output. These are defined on the COMPTn keyword of the TERMINAL macro or on an ETO logon descriptor. Four parameters are supplied:

SINGLE1

Asynchronous output for this component is sent one message per bracket.

SINGLE2

Asynchronous output for this component is sent one message at a time with the VTAM begin-bracket (if necessary) and change-direction indicators to allow the receiving subsystem to optionally send its communication traffic.

MULT1

All asynchronous messages for a given LTERM are sent before the bracket is ended.

MULT2

All asynchronous messages for a given LTERM are sent before change-direction is sent.

For IMS-to-IMS sessions, the selection of SINGLE1, SINGLE2, MULT1, or MULT2 can be related to the transaction types and characteristics of the receiving IMS subsystem. The following considerations apply:

SINGLE1

This protocol is appropriate for sending message switches, nonresponse transactions, or nonconversational transactions to the receiving IMS.

SINGLE2

This protocol is appropriate for sending response mode (including Fast Path) transactions, commands, and test mode, because synchronous communication assumes that input generates corresponding output. Conversational mode is not supported on an IMS-to-IMS session.

MULT1 or MULT2

These protocols treat asynchronous traffic the same way as SINGLE1 or SINGLE2 treat synchronous traffic. They are useful in suppressing contention when large amounts of asynchronous traffic are queued for transmission on one (MULT1) or both (MULT2) subsystems.

Related concepts:

“Relationship of ISC and IMS execution modes” on page 475

Related reference:

“Bracket and half-duplex protocol” on page 533

“Bracket protocol for IMS output” on page 536

Accessing existing application programs with ISC

During an ISC session, you can access application programs written prior to your installation's support of ISC.

Consider the information in the following topics when planning to access these programs.

Accessing programs that use MFS

If a program that uses MFS program is sensitive to the MFS MODname, MFS must be used when input is received from an ISC session.

Your program is sensitive, for example, if the program tests the MODname in the IOPCB as part of its logic, or uses ISRT calls to the IOPCB that include a MODname as part of the call.

If your program is sensitive to cursor position, it cannot be used to receive ISC session input. An ISC LTERM is not a hardware device and, as a result, cursor position has no meaning.

For output, MFS provides support for attribute and extended attribute bytes.

ISC supports the system control area (SCA) and default system control area (DSCA) in the following manner. All functions allowed for the 3270 display can be specified by the application program in a message field (MFLD) defined as an SCA. A device field (DFLD statement) can be defined as an SCA in the device output format (DOF). For ISC subsystems, MFS does not interpret the SCA specifications; it merely relays them in the user-defined device field SCA that it sends to the remote program or ISC subsystem. The MFS-supported SCA, or DSCA, for ISC session output can be used by the receiving application program to intelligently handle a terminal device attached to its subsystem.

If your application program is not sensitive to MFS formatting, it is possible for an ISC-connected subsystem to send an input data stream to your program without invoking MFS. To accomplish this, the sending subsystem must construct an input data stream that is exactly the same as your program's I/O area. This "bypassing" of MFS would result in a reduction of the processor overhead associated with MFS. Bear in mind that bypassing MFS in this manner can result in the transmission of inefficient data streams from the sending subsystem and also makes the sending subsystem sensitive to changes in your program's I/O area.

It is probably easiest, from an implementation point of view, to use MFS with ISC when MFS is already used for input to your program. In this case, you simply create ISC (DPM-Bn) device input formats (DIF) and DOFs.

Related reference:

 DFLD statement (System Utilities)

Accessing programs that do not use MFS

The two input edit choices for existing application programs that do not use MFS are ISC edit and basic edit.

ISC Edit is the default edit process for ISC input to IMS. An existing application program expects the transaction code to be at the start of the input message and to be followed by a blank. When this is true, the sending subsystem should ensure that the transaction code is placed at the start of the data stream, followed by a blank. Password security is for statically defined terminals only and can only be invoked by placing the password after the transaction code in the data stream and ensuring that the SNA-defined PRN in the FM header is not present. When ISC Edit is used, the input data stream sent to IMS should match the application program's I/O area exactly.

Basic edit for ISC input provides standard editing as if the input comes from a terminal. It is used to edit device and operator control characters. The sending subsystem could simply pass through to IMS a data stream from a device that would normally use basic edit if it were attached directly to IMS. In this situation, IMS's ISC basic edit support eliminates application-provided editing in the sending subsystem. Basic edit is selected by the sending subsystem when it places the characters 'BASICEDT' in the DPN field of the FM header.

Routing messages

The following topics introduce the routing fields in the SNA-defined FM headers and describes how IMS uses these fields to route messages (and thus support distributed processing).

FM headers for message routing

ISC message routing information is provided within SNA-defined function management (FM) headers.

On input, IMS automatically processes the information in these FM headers and then removes them from the input message. IMS then passes the message to the appropriate input edit process. On output, IMS automatically builds the necessary routing FM headers based upon information provided by IMS system definition, session bind parameters, the input FM header, or MFS.

For this explanation of distributed transaction processing, the portion of the ATTACH and SCHEDULER FM headers of interest are the routing parameters that they carry. Four fields are used to route messages to appropriate control blocks. These routing fields, present in both headers, are:

DPN (Destination Process Name)

The DPN parameter names an input process to be invoked synchronously to the session. Within IMS, the processes that can be invoked are basic edit, ISC edit, or MFS. For MFS, the input DPN value is the MFS message input descriptor (MID) name. For other subsystems, the named process can be the message destination (such as a transaction code).

PRN (Primary Resource Name)

The PRN parameter names a resource to be associated with the attached process. This parameter occurs on a reply returned as the result of processing a message in a remote subsystem.

Within IMS, the input PRN can name either an LTERM message queue for terminal output (IMS message switch) or an input transaction message queue. When the PRN is not available on input, IMS defaults to the normal method of using the message destination contained in the first data field (or optionally inserted by MFS formatting) or a preset transaction code established by an IMS operator command.

RDPN (Return Destination Process Name)

The RDPN parameter defines a suggested return destination process name to be associated with an output message. It should be returned as the DPN on resulting message replies to facilitate the processing of replies returned to the source subsystem. The input RDPN is discarded when replies or message switch output is sent on another session. The RDPN is sent by IMS only for MFS-formatted output, wherein the RDPN is the optional chained MID name specified on the MFS message output descriptor (MOD).

RPRN (Return Primary Resource Name)

The RPRN parameter defines a suggested return primary resource name to be associated with an output message. It should be returned on resulting message replies to the subsystem that was the source of the message to facilitate return reply processing within that subsystem. The input RPRN is discarded when replies or message switch output is sent on another session. The RPRN can optionally be set by use of MFS formatted output if specified on the MFS MOD.

Related reference:

“Function management headers” on page 559

“ATTDPN” on page 581

“ATTRPN” on page 583

“ATTRDPN and ATTRPRN” on page 583

“ATTDQN and ATTDQ” on page 585

IMS use of routing parameters

IMS ISC without the Front-End Switch exit routine does not provide for automatic transaction routing when IMS is the front-end subsystem.

These transactions can be sent as alternate PCB output from a user-written application program within IMS or by using an ISC message switch. In the latter case, the source terminal operator provides the output routing information through an ISC destination LTERM name which is the required first data field on IMS message switches.

Except when the output is formatted by MFS (by using a MOD chained from the input MID associated with the message switch), IMS message routing support for ISC removes the first data field before sending the balance of the message text to the back-end subsystem. When MFS is used to format message switch output, you should use MFS editing to remove the output LTERM name.

Further, ISC message routing support within IMS defaults the RPRN within the output FM header sent with the message to the LTERM name associated with the terminal that was the source of the message switch. For alternate PCB output, the RPRN is set to the I/O PCB name. The RPRN allows IMS to automatically route message replies from the remote subsystem back to the source terminal in the form of a message switch, provided the other subsystem has returned the RPRN parameter as the PRN parameter on the reply. MFS DPM can be used to alter the FM header destination and return routing parameters. When using IMS as a front-end subsystem to route messages to another subsystem or route replies back to the source terminal, MSGDEL=SYSINFO must be specified on the IMS system definition TERMINAL macro statement or an ETO user descriptor associated with both the ISC session and the source terminal.

Routing examples

The following routing examples show typical functions that can be accomplished between two subsystems using ISC and how routing parameters are used to accomplish these functions.

In the examples:

- The ISC edit alias is ISCE.
- The LTERM name for input and output associated with the terminal on the left in the diagrams is T.
- The LTERM name for input and output associated with the terminal on the right in the diagrams is T2.
- The LTERM name for the ISC session between the two subsystems is LTISC1.


Related reading: For an example of the Front-End Switch exit routine, see *IMS Version 11 Exit Routines*.

Related reference:

 Front-End Switch exit routine (DFSFEBJ0) (Exit Routines)

Example 1. IMS-to-IMS message switch routing:

The following figure shows the terminal attached to IMSA uses the ISC message switch to input TRANX to be executed in IMSB.



Note to figure: The name ISCE was defined for ISC edit during system definition by specification of EDTNAME=ISCE on the COMM macro in both subsystems.

Figure 59. ISC example for IMS-to-IMS message switch routing

An MPP in IMSB processes TRANX, and the reply is routed back to the terminal in IMSA, which was the source of the message switch. The output terminal reply appears as an output message switch.

1. Assume the terminal is a 3270 display device. Based on this assumption, basic edit can only be invoked by entering input from a cleared screen. Therefore, in this example, the terminal operator enters:

LTISC1 | TRANX | Data...

LTISC1 is the IMSA logical terminal name associated with an ISC session between IMSA and IMSB.

2. Basic edit edits the input data stream from the terminal, and the message is placed on the IMSA message queue with a destination of LTISC1.
3. On output to LTISC1, ISC support in IMSA:

- a. Strips the destination LTERM name (LTISC1) from the data stream
- b. Builds the FMH required to send the transaction to IMSB

4. The data stream that is sent to IMSB looks like:

FMH: DPN=SCHEDULER

FMH: DPN=ISCE,PRN=,RDPN=,RPRN=T | TRANX | Data...

- a. DPN=ISCE because IMS ISC support supplies this value as a default if no DPN is supplied when output is to be sent to another subsystem.
- b. PRN= is not supplied and also not required for this example. TRANX is part of the user data. Only through the use of MFS can PRN= be supplied.
- c. RDPN= is not supplied and also not required for this example. Only through the use of MFS can RDPN= be supplied.

- d. RPRN=T is automatically inserted by IMSA as a default function of the message switching logic incorporated in IMS ISC.
5. The TRANX data stream is edited by ISC edit (ISCE) on input to IMSB, because the FMH specified DPN=ISCE. After editing, the data stream is placed on the message queues, and looks like:
TRANX | Data...
6. When scheduled, an MPP retrieves TRANX from the message queues and processes the transaction. Output inserted (ISRT) by the MPP to the originating input terminal looks like:
Data...
This output can be inserted (ISRT) by the MPP to its I/O PCB (on the same parallel session) or to an alternate PCB. The I/O PCB or an alternate response PCB must be used if TRANX is a response-mode transaction. An alternate PCB could be used if TRANX were a nonresponse transaction. Alternate PCB output is sent on the same or a different parallel session, depending on the session to which the output LTERM is assigned.
7. On output from IMSB to IMSA, the FMH is built and sent with the data:
FMH: DPN=SCHEDULER
FMH: DPN=ISCE,PRN=T,RDPN=,RPRN | Data...
The output is sent on the same session as that on which the input was received.
 - a. DPN=ISCE is specified, because ISC support supplies this value as a default if no DPN is available from input, or through MFS when asynchronous output is to be sent to another system.
 - b. PRN=T is supplied, because the input FMH to IMSB specified RPRN=T. IMS ISC support automatically wraps an input FMH RPRN value to the output FMH PRN field.
 - c. RDPN= is not required but can be added by MFS DPM. The RPRN is not supplied, because the reply is returned on the same session as the input transaction. However, MFS can also be used to set the RPRN.
8. On input to IMSA, the reply from IMSB is edited by ISC edit (ISCE), because the FMH specified DPN=ISCE.
Because a PRN is supplied in an input FMH, ISC edit uses the PRN as the IMS destination and appends the PRN value to the input data. After ISC edit, the message is placed on the input message queues and looks like:
T | Data...
If MFS had been used to process the reply within IMSA, the insertion of the LTERM name can be suppressed. See Example 6.
9. An IMS default MOD is used to format the data for output to the terminal. The data displayed to terminal T is:
T | Data...

Example 2. IMS-to-IMS application routing:

The following figure illustrates IMS-to-IMS routing using an application program in both IMS subsystems.




Figure 60. ISC example for IMS-to-IMS application routing

The input format from the terminal operator differs from that in Example 1 in that the routing can be supplied by TRANX in IMSA rather than by the terminal operator. Processing in IMSB and on the session is the same as that described for Example 1. Also see Examples 6 and 7 later in this topic to understand the interaction of MFS in the routing scenario.

Example 3. IMS-to-other-subsystem message switch routing:

The function to be accomplished with the example shown in the following figure is the same as that in Example 1. However, the backend subsystem is either CICS or a user-written subsystem.




Figure 61. ISC example for IMS-to-other-subsystem message switch routing

The description of the activities taking place in IMS for this example parallels that in Example 1. The “other” subsystem is of interest and is described below.

1. Assume the other subsystem is CICS. Remember, from Example 1, the data stream sent from IMS looks like this:
`FMH: DPN=ISCE,PRN=,RDPN=,RPRN=T | Data...`
 The data is formatted in a way that ISCE expects to receive it. The IMS terminal operator who is entering the transaction should understand what that format is.
2. CICS must have a user-written transaction defined with transaction code ISCE, because the DPN represents a transaction code to CICS. CICS, after receiving the data stream described in item 1, invokes a transaction named ISCE, which must be specifically written for the ISC environment. This transaction examines the FMH values supplied and must use the RETRIEVE interface to obtain the user data and process TRANX.
 Because CICS supports transaction codes having a maximum of 4 characters, ISC edit has been given the alias ISCE.
3. The output from the transaction in CICS must look like the output that IMSB sent back to IMSA in Example 1. That is:
`FMH: DPN=ISCE,PRN=T,RDPN=,RPRN | Data...`
 The CICS transaction ISCE actually gives CICS the format to use to build the output FMH (in contrast to Example 1, wherein IMS ISC support builds the FMH). In this example, the CICS transaction must supply DPN=ISCE and PRN=T to CICS to enable the proper FMH to be built.
4. When IMS receives this output data stream from CICS, all activities parallel those in Example 1.

A user can implement the functions of this example in a user-written subsystem by understanding the previous explanation of CICS functions and relating it to the functions implemented in the system.

Example 4. IMS-terminal-to-IMS-terminal message switch routing:

The following figure shows a terminal T connected to IMSA effecting a message switch to terminal T2, which is connected to IMSB.




Figure 62. ISC example for IMS-terminal-to-IMS-terminal message switch routing

1. The terminal T enters the following data stream:
`LTISC1 | T2 | Data...`

LTISC1 is an LTERM name that IMSA has associated with an ISC session to IMSB. The terminal operator (T) must know this ISC LTERM name as well as the LTERM name of the destination terminal (T2) attached to IMSB.

2. ISC message switch support removes the LTERM name (LTISC1) before sending the balance of the message on the ISC session to IMSB. Therefore, the data stream that is sent on the ISC session looks like:

FMH: DPN=SCHEDULER

FMH: DPN=ISCE,PRN=,RDPN=,RPRN=T | T2 | Data...

- a. DPN=ISCE is specified, because IMS ISC support supplies this value as a default if IMS does not supply a DPN when output is sent to another subsystem.
 - b. PRN= is not required, because the destination, terminal T2, is part of the data stream, but could be supplied or modified by MFS.
 - c. RDPN= is not supplied and is also not required for this example. This value could have been supplied by MFS.
 - d. RPRN=T, as in Example 1, is automatically inserted by IMSA as a default function of the message switching capability of IMS ISC. Because no reply is returned, this parameter is discarded by IMSB. MFS can be used in IMSB to make this parameter available to the operator of terminal T2.
3. Before the data stream is placed on IMSB's message queues, it is edited by ISC edit. In IMSB, ISC edit examines the data (because no PRN is available in the input FMH) to determine the destination. T2 is found to be the destination and the input message is placed on the message queues with a destination T2.
 4. On output from IMSB to terminal T2, the data stream now looks like this:
T2 | Data...
 5. A default system MOD would be used for output if T2 were a required MFS device.

Example 5. IMS-terminal-to-other-terminal message switch routing:

The following figure is similar to Example 4; however, the back-end subsystem is CICS or a user-written subsystem.




Figure 63. ISC example for IMS-terminal-to-other-terminal message switch routing

The description of the activities taking place in IMS in this example parallels the description in Example 1 exactly. The other subsystem of interest is described.

1. Assume the other subsystem is CICS. Remember from Example 1 that the data stream sent from IMS looks like this:

FMH: DPN=ISCE,PRN=,RDPN=,RPRN=T | Data...

2. As in Example 3, CICS must have a transaction code defined as ISCE.

The data is formatted in the way in which ISCE expects to receive it. The IMS terminal operator who is entering the transaction should understand what that format is to be. CICS, upon receiving the data stream, must invoke a transaction named ISCE, specifically written for the ISC environment. This transaction uses the data in the input FMH and obtains the input data through the RETRIEVE interface. Because the PRN field in the input FMH is not initialized by IMSA to be equal to T2, the CICS application must initiate a message switch to terminal T2 from the transient data queue or start a new transaction naming T2 as the primary resource.

In this example, the output terminal T2 must be identified in the input data stream. If IMSA had initialized the PRN field in the input FMH to T2, CICS would have attached the transaction ISCE with T2 as the primary resource, and a SEND to terminal T2 could have been made directly.

3. In order for IMSA to set the PRN field, MFS would be required.

Example 6. IMS-to-IMS message switch routing with MFS:

In the following figure, terminal T enters a transaction to be processed in the backend IMS subsystem.




Figure 64. ISC example for IMS-to-IMS message switch routing with MFS

The reply from the back-end subsystem is to be routed back to the original terminal that entered the transaction (T).

1. Terminal T enters its input from a formatted screen. The terminal only needs to enter data.
2. After editing by MID1, the message is placed on the message queues and looks like this:

LTISC1 | TRANX | Data...

The values LTISC1 and TRANX were appended to the data by MFS.

3. MOD2, to be used for editing the data stream that is sent to IMSB, is chained from MID1, and is chained to another MID (MID4) within IMSA. The data stream to be sent to IMSB looks like this:

FMH: DPN=SCHEDULER

FMH: DPN=MID3, PRN=, RDPN=MID4, RPRN=T

TRANX | Data...

- a. DPN=MID3 can be supplied by MOD2 or from a literal in the DOF associated with MOD2.
 - b. PRN= is not needed, because TRANX is already in the data stream. It could have been supplied from MOD2 or its associated DOF.
 - c. RDPN=MID4 is specified, because MID4 is chained to MOD2.
 - d. RPRN=T is again automatically inserted by IMSA as a default function of message switching and not overridden by MFS.
 - e. LTISC1 again has been stripped from the output message by the MFS MOD.
4. Upon receipt of the data stream from IMSA, IMSB edits the data using MID3, because the DPN in the FMH specified MID3.

5. The application receives and processes the following:

TRANX | Data...

6. Output from the application program is sent by IMSB to IMSA in the following format:

FMH: DPN=SCHEDULER

FMH: DPN=MID4, PRN=T, RDPN=, RPRN= | Data...

7. DPN=MID4 and PRN=T are automatically wrapped by IMS from the RDPN and RPRN values in the original FMH.
8. On receipt of this data stream, MID4 in IMSA edits the input and places the message on the queues for final output to terminal T in the following format:
Data...

Because PRN=T is supplied, IMS uses it as the destination. Because MFS is also used to format the input reply, the PRN is not appended to the data.

9. Because MOD5 is chained to MID4, MOD5 is used to format the output to terminal T.

Several observations can be made about this example:

- MID3 in IMSB is not chained to a MOD. Therefore, the I/O PCB that the application program in IMSB sees does not contain a MODname to be used for output. If MID3 should have a chained MOD, the application ISRT call should have a blank MOD name to negate MFS editing on output.
- The application in IMSB has not changed the MOD name for output. If it does, additional MFS format design is required.
- The MPP in IMSB can be an already-existing program that also handles transactions from terminals connected to IMSB. The use of MFS can make the use of an ISC session transparent to the application.
- The original input terminal is not held in response mode.

Example 7. IMS-to-IMS application routing with MFS:

The following figure can be understood in two contexts based upon an understanding of previous examples—the case shown in which ISC edit is used to edit the input in IMSB or the case in which MFS DPM-Bn is used to edit the input in IMSB.




Figure 65. ISC example for IMS-to-IMS application routing with MFS

Example 7 emphasizes several points:

- Terminal T is completely handled by the front-end subsystem. Therefore, in this example, Terminal T can be held in response mode with IMSA.
- The data sent from IMSA is inserted (ISRT) by an alternate PCB and is asynchronous to the response mode transaction.
- IMSA could be connected to many IMS subsystems and, based on the input transaction, route the ISC traffic to the appropriate back-end IMSB.
- The MPP in IMSA can provide any required message routing information, such as selecting the appropriate back-end subsystem through an ISRT to an alternate PCB.
- MFS can be used in either or both subsystems to provide data stream formatting and routing.

Routing messages through MSC to an ISC LTERM

You can route messages across an MSC link to an ISC half session.

Suppose you have a setup similar that shown in the following figure. IMS A and IMS B connected by an MSC link. SYS C is connected to IMS B with an ISC link. If you want to route input messages and output replies between IMS A and SYS C, place the ISC routing parameters in the IMS prefix portion of the message. The MSC link-processing preserves these routing parameters.




Figure 66. Routing of MSC to an ISC LTERM

Considerations for IMS-to-IMS ISC sessions

Consider the following information before configuring an ISC session between two IMS subsystems using static terminals.

Coexistence of ISC and MSC VTAM within one IMS subsystem

During IMS system definition, the `TERMINAL`, the `MSPLINK` macro statement, or an ETO logon descriptor of one IMS subsystem must be defined with the same name as that defined on the `COMM` (ACB name) macro of the other IMS system.

IMS system definition allows both the `MSPLINK` and `TERMINAL` macros or an ETO logon descriptor to indicate the same remote subsystem name. However, the session qualifier information (the name used on the `SUBPOOL` macro and the partner-id parameter on the `MSPLINK` macro) must be unique. This allows MSC and ISC sessions to be defined and simultaneously active between two IMS subsystems. The number of ISC parallel sessions that can be active simultaneously between two IMS subsystems is the smaller of the values defined for the `SESSION=` parameter of the `TERMINAL` macro statements of the two IMS subsystems. (The `SESSION=` parameter is for statically defined terminals only.) The number of MSC parallel sessions that can be active simultaneously between two IMS subsystems is the smaller of the values defined for the `SESSION=` parameter of the `MSPLINK` macros of the two IMS subsystems. The total number of MSC and ISC parallel sessions that can be simultaneously active is the sum of these two smaller numbers.

Defining single and parallel sessions

During IMS system definition, the `TERMINAL` macro statements or the ETO logon descriptors defining the ISC session(s) between the IMS subsystems must be defined identically as either single session or parallel sessions. However, when both subsystems are defined for parallel sessions, the number of concurrently active sessions or number of `SUBPOOL` macro statements need not be identical.

Without ETO, you can define up to 4,095 parallel sessions. With ETO, the number of ISC parallel sessions is limited primarily by the processor capacity and the amount of virtual storage above and below the line available to the control region.

Ensuring compatible buffer sizes

The VTAM output buffer size defined on the `TERMINAL` macro statement or an ETO logon descriptor of one IMS subsystem must be compatible with the “receive-any” buffer size defined on the `COMM` macro statement of the other subsystem.

Remote control of IMS through ISC

The IMS ISC message switch support or user-written applications allow a terminal operator within one IMS subsystem to effectively control the operation of another IMS subsystem. The IMS automated operator interface (AOI) facility can also be used to aid in the control of either IMS subsystem.

For statically defined terminals, you can authorize commands or transactions to ISC sessions within the remote IMS subsystem; all terminal operators and application programs within the local IMS subsystem that are authorized to send data on the ISC session can then access the remote subsystem. Therefore, use password security to restrict remote subsystem access to only specific authorized individuals and applications as a master terminal within the IMS subsystem.

You can provide additional security by specifying that all commands on the ISC session be issued between automated operator interfaces in both subsystems or issued to a single automated operator interface in the back-end subsystem.

A logical unit type 6.1 cannot be defined as the IMS master terminal during IMS system definition, nor can it be assigned as a master terminal using the IMS /ASSIGN command.

Restriction on IMS-to-IMS conversation mode

Conversation mode between two IMS subsystems is not supported. This is because conversations between IMS subsystems that are connected using ISC produce unpredictable session protocols and conversational transaction input when only one half session is in conversation mode. Further, the conversation terminates if both half sessions are in conversation mode.

Routing transactions to the back-end IMS

IMS ISC support allows the following to route transactions to a back-end IMS subsystem:

- Terminal operators using ISC message switches.
- Application programs using alternate PCB inserts within the front-end IMS subsystem.

In either case, the default action by the ISC support is to route a resulting transaction reply, in the form of a message switch, back to the originating terminal operator or transaction. Under these conditions, the transactions accessed within the back-end IMS subsystem should generally be defined as recoverable, because the response required for the message switch reply is also recoverable.

Irrecoverable transactions can be accessed, but require the use of MFS DPM-Bn to change the default destination set in the ISC message routing parameters from the LTERM associated with the originating terminal operator to an irrecoverable transaction within the front-end IMS. This avoids the protocol errors that occur when sending irrecoverable reply messages. The destination can be changed by MFS DPM either within the front-end IMS to modify or delete the default RPRN parameter sent on the transaction to the back-end IMS, or within the back-end IMS to modify or delete the wrapped PRN parameter sent on the reply to the front-end IMS. This irrecoverable transaction can then route the reply to the appropriate terminal operator by inserting to a modifiable alternate PCB.

Special support exists for front-end/back-end system utilization provided by the Front-End Switch exit routine.

Sending messages between the subsystems

Communication from an IMS front-end is always asynchronous. Therefore, messages are required to be sent and received with both the ATTACH and SCHEDULER FM headers. The exception to this is that system messages are sent with the ATTACH FM header only.

Protocol restrictions on IMS-to-IMS sessions

When an ISC session is between two IMS subsystems, certain types of protocols are not sent between the subsystems. This subtopic provides these protocols. If you

are designing IMS-to-IMS sessions, you can skip the information on these protocols and headers provided in Chapter 31, “ISC protocols for VTAM connections,” on page 503.

- The data flow control commands BID, RTR, and RSHUT are not sent between the two IMS subsystems.
- The reset-attached process (RAP) FM header is not sent between IMS subsystems.
- The queue model (QMODEL) headers are not sent between IMS subsystems. This precludes IMS output demand-paging and all associated input QMODEL paging requests. Also, the ATTDQN parameter on the ATTACH FM header and the SCDDQN parameter on the SCHEDULER FM header are not supported in IMS-to-IMS sessions.
- The following error codes are not sent between two IMS subsystems:
 - LUSTATUS commit and function abort X'0864' and X'0866'
 - All sender ERP sense codes except X'0813' and X'0846'
 - Selective receiver ERP sense codes, except as listed under “Selective receiver ERP” on page 549.

Related reference:

“System message process (SYSMSG) and related FM headers” on page 564

 Front-End Switch exit routine (DFSFEBJ0) (Exit Routines)

Statically defining an ISC node to IMS

A subsystem that is statically defined at system definition to IMS as an ISC node appears to IMS to be a terminal and to an IMS application program to be one or more logical terminals (LTERMs). Thus, an ISC node is defined using those system definition macros applicable to defining other VTAM terminals.

The following table shows the macro statements that define an ISC node.

Table 72. Macro statements for defining an ISC node

Macro statement	Use
COMM macro	Names the IMS subsystem to VTAM; names the ISC edit process and defines the receive-any (RECANY) buffer size.
TYPE macro	Identifies the unit type as an ISC node (UNITYPE=LUTYPE6).
TERMINAL macro	Identifies the other ISC node to IMS. Also provides sizes for the output and Fast Path buffers, limits for the input segment size, and processing options. Defines the number of sessions possible between two ISC nodes. Describes component characteristics.
VTAMPOOL macro	Begins the definition of the ISC LTERM users (subpools) used with parallel sessions.
SUBPOOL macro	Defines a set of LTERMs within the VTAMPOOL.
NAME macro	Names the LTERMs associated with a given terminal for a single session or a given user (subpool) for parallel sessions.

This topic more fully describes the macros that are key to defining an ISC session between two systems and some of the implications of selecting certain system definition options instead of others.

The following figure shows two IMS systems, NYIMS in a New York data center and SFIMS in a San Francisco data center. NYIMS and SFIMS are each defined to

the other as an ISC node. The following paragraphs details the way in which these IMS systems would be defined during IMS system definition. For ease of understanding, this description focuses on defining the IMS SFIMS to the IMS NYIMS. A similar system definition would be required to define IMS NYIMS to IMS SFIMS. Only those parameters of interest on the COMM, TERMINAL, and NAME macro statements are described.




Figure 67. Two IMS systems defined to each other as ISC nodes

The parameter APPLID=NYIMS on the NYIMS system definition COMM macro makes that subsystem known by the node name NYIMS to VTAM, to itself, and to any other subsystems that need to access it. A similar system definition must occur for the San Francisco IMS SFIMS.

NAME=nodename on the TERMINAL macro statement of the NYIMS system definition is coded to show the node name of the SFIMS system (NAME=SFIMS). If SFIMS is an XRF complex, the node name should be the USERVAR or MNPS ACB associated with the SFIMS system.

The SESSION= keyword is related to parallel sessions rather than to a single session. (The SESSION= keyword applies to statically defined terminals only. It does not apply to the ETO environment.) SESSION= specifies the maximum number of parallel sessions that are allowed from NYIMS to SFIMS. It is possible to code one "parallel" session. A single session and one parallel session are not the same.

By coding SESSION= as greater than one, the NYIMS might maintain multiple concurrent sessions with SFIMS, up to the number specified by the SESSION= parameter. In the figure above, four paths are drawn between the two IMS systems. Each path represents a parallel session. To allow four parallel sessions in the NYIMS, code SESSION=4.

A major advantage of parallel session support is the ability to dynamically allocate LTERMs to that session based on the IMS /OPNDST command parameters.

Each parallel session is required to be identified to NYIMS using the SUBPOOL macro statement. In the example, to allow four concurrently active parallel sessions, four users (subpools) must be defined as follows in the NYIMS system definition:

```

VTAMPOOL
  SUBPOOL  NAME=SF1
  SUBPOOL  NAME=SF2
  SUBPOOL  NAME=SF3
  SUBPOOL  NAME=SF4

```

The VTAMPOOL macro statement begins the definition of the ISC subpools and must be coded if parallel sessions are to be used. The VTAMPOOL macro statement has no operands.

In Figure 67 on page 497 SFLT1, SFLT2, SFLT3,...SFLTn represent LTERM names that have been defined to the first SUBPOOL, SF1. The definition of LTERMs within a given subpool is accomplished by coding the NAME macro. Defining multiple LTERMs to a subpool provides several advantages. The first is that an LTERM or LTERMs defined to a given subpool can be reassigned (using the /ASSIGN command) to any other subpool. This capability gives flexibility in the operation of an ISC network.

The second advantage is with the COMPTn= keyword on the TERMINAL macro statement. This keyword can be used to define or assign certain characteristics to a given LTERM. COMPT= and ICOMPT= are the keywords on the NAME macro statement and are used to associate a component (COMPTn) with an LTERM.

The definition of SFIMS to NYIMS now looks like the following code:

```

TYPE      UNITYPE=LUTYPE6
TERMINAL  NAME=SFIMS
          SESSION=4

```

```

VTAMPOOL
  SUBPOOL NAME=SF1
    NAME SF1LT1
    NAME SF1LT2
  :
    NAME SF1LTN
  SUBPOOL NAME=SF2
  :
  SUBPOOL NAME=SF3
  :
  SUBPOOL NAME=SF4
  :
  :

```

Related concepts:

“LTERM users (subpools) and components” on page 478

Related reference:

 Macros used in IMS environments (System Definition)

Choosing parameters: system design considerations

The following topics further describe the parameters of the COMM, TERMINAL, NAME, and SUBPOOL macro statements.

COMM macro statement

The key parameters on the COMM macro statement are APPLID, RECANY, and EDTNAME.

The APPLID parameter defines the VTAM ACB name for IMS. The name specified for APPLID= must be the same as the ACBNAME= parameter of the VTAM APPL definition statement. If the ACBNAME= parameter is not specified, the APPLID= name must be the same as the name of the APPL definition statement that is used when defining IMS to VTAM. If the other subsystem is also an IMS, the name specified for APPLID= must be the same as that supplied on the NAME= keyword parameter of that IMS TERMINAL macro statement.

If you choose not to define the APPLID parameter on the COMM macro statement and instead use the job step name of the EXEC statement when initiating the IMS control region, VTAM attempts to match that name with either the ACBNAME= parameter or the name of the VTAM APPL definition statement, as described in the preceding paragraph.

When specifying the receive-any (RECANY) buffer size on the COMM macro, remember that a 28-byte overhead applies to the size of the receive buffers for ISC sessions. Therefore, the RECANY buffer size specified on the NYIMS system definition must be at least 28 bytes larger than the size specified for OUTBUF= on the TERMINAL macro of the SFIMS system definition for NYIMS. When defining SFIMS, the same considerations apply.

Other factors affect the size of the RECANY buffers. Because the specified RECANY buffer size applies to all VTAM devices, the size specified for RECANY might be larger than the size required for ISC due to the needs of other VTAM devices. Also, the FM header size must be added to the maximum data record size for both systems to determine the maximum ISC RECANY size.

The EDTNAME parameter specifies the alias that can be used for ISCEDT in your IMS system. In this system definition, the default name, ISCEDT, is used, making the specification of the EDTNAME= parameter unnecessary.

Related reference:

 COMM macro (System Definition)

NAME macro

The key system definition options for defining an ISC session on the NAME macro are: COMPT, OUTPUT, EDIT, and ICOMPT.

COMPT= and ICOMPT= specify the output and input components respectively to be associated with the LTERM being defined. The system definition example uses components defined as COMPT=1 and COMPT=2.

OUTPUT= should not be used on an ISC session.

Recommendation: If you do not want translation of ISC output to the session, specify ULC. If session output can be uppercase or lowercase data, or binary data, use ULC.

SUBPOOL macro

The SUBPOOL macro identifies ISC subpools to a system.

The key system definition options on the SUBPOOL macro are as follows:

- NAME
- MSGDEL

The NAME parameter identifies the ISC subpool.

In this example, the MSGDEL specification is the default, SYSINF0. When MSGDEL is specified on the SUBPOOL macro, it must match the MSGDEL specification on the TERMINAL macro.

Related concepts:

“LTERM users (subpools) and components” on page 478

TERMINAL macro

Several system definition keyword parameters on the TERMINAL macro are principal for defining an ISC session.

The principal system definition keyword parameters for defining an ISC session on the TERMINAL macro are as follows:

- OUTBUF
- MSGDEL
- COMPT1
- COMPT2
- COMPT3
- COMPT4

Additional parameters you must specify include:

- SINGLE or MULT
- VLVB or DPM-B1...DPM-B15

Specifying the OUTBUF= keyword parameter

IMS does not negotiate buffer sizes when initiating an ISC session, nor does IMS support RU truncation. When sending a negotiated bind, IMS sets its send (OUTBUF=) and receive (RECANY=) sizes in the bind.

When receiving the bind reply, IMS verifies that:

- Its send (OUTBUF=) size has not been reduced by the other subsystem.
- Its receive (RECANY=) size is not exceeded by the other subsystem's send (OUTBUF=) size.

When receiving a negotiated bind, IMS:

- Verifies that the send (OUTBUF=) size of the other subsystem does not exceed IMS's receive (RECANY=) size.
- Inserts the send (OUTBUF=) size into the bind reply.

The specified OUTBUF size must be large enough to include the largest output segment size plus the overhead for message headers. When defining the SFIMS system, the same considerations apply.

Specifying the MSGDEL= keyword parameter

A description of the ramifications of choosing SYSINFO or NONIOPCB is found in “LTERM users (subpools) and components” on page 478. In the example in that topic, the default for SYSINFO is used for both the TERMINAL and the SUBPOOL macros.

Specifying the COMPTn= keyword parameter

Up to four components can be defined for a given ISC node on the `TERMINAL` macro by using the keyword parameters `COMPT1=` through `COMPT4=`. `COMPTn=` allows flexibility in assigning characteristics to LTERMs. In the system definition example, component 1 is specified as having the characteristics `SINGLE1`, `DPM-B1`, and `IGNORE`. Component 2 is specified as having the characteristics `SINGLE2`, `DPM-B2`, and `IGNORE`.

Specifying the SINGLE or MULT parameter

The characteristics of `SINGLE1`, `SINGLE2`, `MULT1`, and `MULT2` are described in “Determining output protocols” on page 481.

Specifying the VLVB or DPM-B1...DPM-B15 parameter

Selecting `VLVB` or `DPM-Bn` determines whether the component can use the Distributed Presentation Management feature of MFS. Although MFS is specified for a component by specifying `DPM-Bn`, its use is optional, on a message-by-message basis. `DPM-Bn` is used for this system definition example.

The selection of `VLVB` precludes the use of MFS-DPM for a component and indicates that variable-length, variable-blocked format is to be used instead of MFS for both input and output.

Specifying the IGNORE or 1, 2,...10 parameter

This parameter is used to specify a user-defined feature code for MFS `DPM-Bn`. The designated feature is used to select an MFS format (`DPM-B1...DPM-B15`) that contains the matching feature specification. `IGNORE` is used to specify that an MFS format (`DPM-Bn`) with the `FEAT=IGNORE` of the `DEV` statement is to be selected. `IGNORE` has been used for this system definition example.

Related reference:

“COMM macro statement” on page 498

System definition summary

SFIMS from the previous example can be defined to NYIMS as shown in the following example.

```
TYPE          UNITYPE=LUTYPE6
TERMINAL NAME=SFIMS
  COMPT1=(SINGLE1, DPM-B1, IGNORE)
  COMPT2=(SINGLE2, DPM-B2, IGNORE)
  SESSION=4
VTAMPOOL
SUBPOOL NAME=SF1
  NAME SF1LT1, COMPT=1
SUBPOOL NAME=SF2
  NAME SF2LT1, COMPT=2
SUBPOOL NAME=SF3
  NAME SF3LT1, COMPT=2
SUBPOOL NAME=SF4
  NAME SF4LT1, COMPT=2
```

This definition does not correspond exactly to the figure in “Statically defining an ISC node to IMS” on page 496. In that figure, the SUBPOOL `SF1` contains multiple LTERMs (to indicate that multiple LTERMs can be assigned to a SUBPOOL). Based

on SINGLE1, SINGLE2, MULT1, and MULT2, the definition can now be simplified to show just one LTERM being defined per SUBPOOL.

The first SUBPOOL, SF1, has LTERM SF1LT1 defined with COMPT=SINGLE1. For asynchronous output from NYIMS to SFIMS, LTERM SF1LT1 is used as the destination LTERM for all:

- Alternate PCB output from application programs in NYIMS that generates asynchronous transactions in SFIMS
- Message switches originated by a terminal connected to NYIMS
- Replies to nonresponse transactions received by NYIMS from SFIMS

The characteristics of SINGLE1 (BB/message/EB) should allow one parallel session to handle all asynchronous input and output between NYIMS and SFIMS.

Three SUBPOOLS (SF2, SF3, SF4) have been defined with single LTERMs, each defined with a COMPT defined as SINGLE2. These three SUBPOOLS (sessions) can be used to send response mode transactions to SFIMS. Because of the nature of response mode transactions (session tied up from transaction origination until receipt of reply), three parallel sessions are defined to allow the user to minimize response mode bottlenecks between NYIMS and SFIMS.

In most IMS ISC definitions, the assignment of one LTERM to each SUBPOOL is sufficient to handle the communications traffic between the two nodes.

Not shown, but certainly to be considered, is the possibility of defining a fifth SUBPOOL within NYIMS to handle incoming SINGLE1 traffic from SFIMS. By design then, NYIMS could receive all SINGLE1 input from SFIMS on the session defined by the fifth SUBPOOL and send all SINGLE1 output on the session defined by the first SUBPOOL (SF1). This plan would result in minimal contention for asynchronous output between the two systems.

Related concepts:

 Defining VTAM terminals (System Definition)

Chapter 31. ISC protocols for VTAM connections

IMS uses ISC protocols to control sessions, data flow, and message routing over ISC VTAM connections. The following topics include the specific protocol information that you need to send and receive data with an ISC link.

Operating the network

Because ISC permits multiple sessions between logical units, the balance of this explanation of ISC differentiates between the terms logical unit, session, and half session.

Definitions:

- The term *logical unit* is used when referring to characteristics common to all sessions between two session partners or when it is not necessary to differentiate between individual sessions.
- The term *session* is used when describing a characteristic unique to a given connection (session instance) between logical units.
- A *half session* describes characteristics unique to one of the session partners.

Making IMS ready

Use the IMS /START command with the DC keyword to make IMS ready to receive VTAM logon or BIND SCIP (session initialization) requests.

The DC keyword initiates IMS data communications processing, opens the VTAM access method control block (ACB) if it is not already open, and enables the IMS VTAM logon exit. Any logon requests received by VTAM before the IMS /START DC command is issued but after the ACB has been opened are queued in VTAM until the /START DC command is completed. If VTAM is active when IMS is initialized, and the DFSDCxxx PROCLIB member keyword VACBOPN=INIT, then the IMS VTAM ACB is opened. If DFSDCxxx PROCLIB member keyword VACBOPN=DELAY, then the IMS VTAM ACB open is delayed until the /START DC command is processed.

The /START DC command also tells VTAM to pass any queued VTAM logon requests to IMS.

Starting an IMS network for ISC

Before any sessions can be established, VTAM and NCP must be active. In addition, all logical units must be online (activated by the VARY command).

You do not need to perform message resynchronization after a cold start.

Related concepts:

 Starting Transaction Manager (Operations and Automation)

Related tasks:

“Controlling the session (session control protocols)” on page 504

Shutting down an IMS network for ISC

Use the IMS /CHECKPOINT command to terminate the network and shut down IMS.

The format of the /CHECKPOINT command determines whether the network termination occurs immediately or waits for processing to complete:

- /CHECKPOINT FREEZE|DUMPQ|PURGE immediately terminates the session for all logical units, as follows:

FREEZE Immediately after current input/output message

DUMPQ After control blocks are past checkpoints

PURGE After all queues are empty

- /CHECKPOINT FREEZE|DUMPQ|PURGE QUIESCE allows all network nodes to complete normal processing before shutting down.

The IMS command /STOP DC also shuts down an IMS network at completion of processing.

Also, certain VTAM commands, such as VTAM HALT NET, shut down the network.

Related tasks:

 Shutting down an IMS network (Operations and Automation)

Controlling the session (session control protocols)

Session initiation includes initiating a session instance, binding the session, and, if necessary, ensuring that the half sessions are in sync.

When both half sessions have agreed to the bind and are in sync, normal traffic flow is initiated using the VTAM start data traffic (SDT) command.

Initiating an ISC VTAM session

A session must be established before data can be transmitted between another logical unit type 6.1 and IMS. If message resynchronization is required, it is performed after the bind.

IMS can be requested to initiate a session in one of the following ways:

- An ISC logical unit requests session initiation by sending the “initiate-self” command. VTAM verifies the command and passes the request to IMS.
- An ISC logical unit sends BIND to initiate a session with IMS.
- The z/OS VTAM network operator requests session initiation on behalf of the logical unit by using the z/OS VARY command with the LOGON option. VTAM processes the request and passes it to IMS. The VARY command cannot be used to initiate parallel sessions.
- VTAM passes to IMS a logon or BIND SCIP request for each logical unit defined to VTAM as belonging to IMS. (This method cannot be used to initiate parallel sessions.)
- An authorized IMS terminal operator requests session initiation for an ISC logical unit by entering the IMS /OPNDST command.

IMS can initiate (send BIND) or accept (receive BIND) a session if all the following conditions are met:

- The IMS master terminal operator has issued a /START DC command.
- The logical unit name is known by IMS.
- The logical unit is not stopped (/STOP command) within IMS.

- The logical unit has not reached the maximum allowable sessions defined to IMS. (This applies to statically defined terminals only.)
- If the session is to be parallel, the CINIT or BIND session qualifier field must contain a valid LTERM subpool name that defines the message queue set to be used. A valid subpool name is one that is not stopped or allocated (except during session restart). The session qualifier fields must not be supplied when requesting session initiation with a logical unit defined to IMS as single session.
- For statically defined terminals, the MSGDEL option for the LTERM subpool as specified on the SUBPOOL macro does not conflict with the MSGDEL option specified for the session on the TERMINAL macro.
- No invalid or conflicting parameters are indicated on either a CINIT, BIND, or negotiable BIND reply.

Related concepts:

“Statically defining an ISC node to IMS” on page 496

Related reference:

 Bind parameters for SLU P and LU 6.1 (System Programming APIs)

Establishing a connection with an XRF complex

An XRF complex can be a session partner.

There are two types of XRF complexes: those that use a USERVAR to manage sessions during an XRF takeover and those that use an MNPS ACB to manage sessions during an XRF takeover.

In either case, if the active IMS subsystem in the XRF complex fails, the alternate IMS subsystem re-establishes the session during takeover. Until the ISC session between the old active IMS and the remote system terminates completely, session initiation fails. Upon takeover, the alternate IMS subsystem becomes the active subsystem and retries session initiation every 30 seconds (for a maximum of 20 attempts) until the session initiates successfully.

Connecting to an XRF complex that uses an MNPS ACB

To establish a connection with a session partner that is an XRF complex that uses an MNPS ACB, simply initiate the session with the XRF complex by specifying the MNPS ACB name as you would to connect to any other non-XRF IMS system.

Connecting to an XRF complex that uses a USERVAR

To establish a connection with a session partner that is an XRF complex that uses a USERVAR, the steps you take differ depending on whether the session is initiated from the XRF complex or the IMS system outside of the XRF complex.

- If the non-XRF half session initiates the session with IMS, the non-XRF half session should:
 1. Identify the APPLID of the currently active IMS system by issuing the following macro:
`INQUIRE OPTCD=USERVAR,AREA=area,AREALN=8`
 If the VTAM network is Release 3.2 or later and USERVAR Management Enhancement is installed, you do not need to issue the INQUIRE macro. With this request, the AREA parameter specifies the address of the area with the USERVAR name for the active XRF system. The USERVAR is then replaced with the APPLID of the currently active IMS system.
 2. Issue a session-initiation request by specifying the APPLID of the active IMS system.

If the VTAM network is Release 3.2 or later and USERVAR Management Enhancement is installed, issue a session-initiation request by specifying the USERVAR of the XRF system.

- If the IMS system in the XRF complex initiates the session, the IMS system:
 1. Issues the session-initiation request, adding the USERVAR assigned to the BIND in the user data field.
 2. Validates the session-initiation request by referencing the USERVAR in the user data field located at the end of the request. Initiation is not apparent to the user.

Binding the session

In an ISC session, IMS can assume either the primary half-session role (send the BIND) or the secondary half-session role (receive the BIND). However, when session restart and recovery are required, session polarity must be maintained.

Definition: *Session polarity* means that the same session role (primary versus secondary) is in effect at the point of failure and is reestablished by the session initiation request. Otherwise, the request is rejected. When the session is initiated using an IMS /OPNDST command and session restart and recovery are not required, IMS requests to be the primary half session.

Negotiable versus nonnegotiable BIND

As the primary half session, IMS sends either a negotiable or nonnegotiable BIND, depending on parameters in the VTAM mode table.

The mode table entry is indicated on the VTAM CINIT or IMS /OPNDST commands, or defined on the TERMINAL macro during IMS system definition.

Definitions:

- When sending a *negotiable BIND*, IMS sets the bind parameters. Because the secondary half session can change some parameters, all the parameters are checked for validity by IMS when a negotiable BIND response is received.
- Prior to sending a *nonnegotiable BIND*, IMS checks all parameters from the mode table entry for validity. The session is terminated if IMS finds any incompatible parameters.

As the secondary half session, IMS receives both the negotiable and the nonnegotiable form of BIND. When receiving a negotiable BIND, IMS checks the bind parameters, except the secondary network addressable unit (NAU) protocol field, for validity. IMS then sets the secondary NAU protocol field for the negotiable BIND response.

When receiving a nonnegotiable BIND, IMS checks the bind parameters, except for “STSN required” and “BIS sent,” for validity prior to accepting the BIND. IMS must then operate under the secondary NAU protocol definition provided in the bind.

Related reference:

➡ IMS as primary half session (System Programming APIs)

➡ Bind parameters for SLU P and LU 6.1 (System Programming APIs)

Binding single or parallel sessions

Requirements for binding a session differ for single and parallel sessions.

To bind a single session, either negotiable or Nonnegotiable BIND is sent. The ISC logical unit with which IMS is communicating must have been defined with a static set of LTERMs during IMS system definition. IMS ignores the session qualifier pair (SQP) field on both CINIT and BIND.

To bind a parallel session, the CINIT and BIND parameters must include a SQP field to identify the specific parallel half-session instance between IMS and a logical unit. (All sessions created dynamically using the ETO feature are parallel sessions.) This field in the bind parameters contains both the primary and the secondary session qualifiers. The half-session name of the ISC node communicating with IMS consists of the logical unit name concatenated with its associated session qualifier. The IMS half-session name is the IMS ACB name concatenated with the session qualifier associated with IMS. The IMS session qualifier is the subpool name. Both half-session names are saved across session and IMS subsystem failures and are used to allocate, or validity-check when warm starting, the LTERM subpool to be used for the session.

Resolving a bind race

A race occurs when IMS and another logical unit simultaneously send BIND requests to each other and the two half-session names are mirror images.

That is, the primary logical unit name and session qualifier concatenated with the secondary logical unit name and session qualifier of one side is equal to the primary logical unit name and session qualifier concatenated with the secondary logical unit name and session qualifier of the other side. The logical unit that wins and becomes primary is the one whose name (the primary logical unit name taken from each half-session name) is at the top of the standard collating sequence. This test prevents both sides from rejecting each other, resulting in no session.

After a successful session bind

Following a successful session bind, both half sessions must perform a resynchronization process. Message resynchronization is not performed when IMS is cold started or when a previous session instance between IMS and another logical unit has been normally shut down using the Stop Bracket Initiation/Bracket Initiation Stopped (SBI/BIS) procedure.

Related concepts:

“Resynchronizing sessions”

Related tasks:

“Terminating an ISC VTAM session” on page 516

Related reference:

“Symmetrical session shutdown for LU 6.1 (SBI and BIS)” on page 558

Resynchronizing sessions

To maintain the integrity of recoverable resources, messages, and queues in IMS across both subsystem and session failures, both half sessions must maintain the session information required for the resynchronization process.

Message integrity cannot be maintained without user intervention when either or both subsystems incur an error or when a user restart procedure destroys this information.

Resynchronization is required when it is possible for a recoverable work unit to be indoubt on a flow (primary-to-secondary or secondary-to-primary).

Definition: A *work unit* includes all transmissions between sync points within a bracket, as illustrated in the following figure. A sync point might have been requested by one or both half sessions without having been acknowledged. These conditions can be caused by a subsystem or session failure or an abnormal completion of a shutdown sequence.

Performing message resynchronization is unnecessary following a normal shutdown sequence (unless nonnegotiated BIND was sent), because both half sessions can come to a controlled, mutual understanding that no additional normal message traffic or sync-point requests are to occur prior to session termination. Message resynchronization is always required following nonnegotiable BIND to allow error conditions detected by the secondary half session to be communicated to the primary half session.

The half-session pairs resynchronize with the VTAM BIND, set-and-test-sequence-numbers (STSN), and start data traffic (SDT) commands. The STSN command allows both half sessions to reestablish sync-point information (session sequence numbers), which is being maintained by both half sessions.

When message resynchronization is necessary, it must be completed successfully before either half session can resume normal data transmission.

The following figure shows two work unit examples.




Figure 68. Work unit examples

In work unit 1 of the preceding figure, a reply to the exception response request and CD creates an implied sync point. In work unit 2, sending the DR2 response to the RQD2 creates a sync point.

Related concepts:

“Keeping half sessions synchronized” on page 525

Related reference:

“Sync point and response requirements” on page 527

Designing restart resynchronization procedures

Resynchronization might or might not be required when IMS is restarted and when a session is established or reestablished. This topic addresses considerations for resynchronization at IMS restart and at session restart.

IMS restart provides for message recovery during:

- A normal IMS subsystem restart procedure, by restoring the IMS message queues and session restart information to the last or specified checkpoint
- An emergency IMS subsystem restart procedure by restoring the IMS message queues and session restart information to their states just prior to failure

If the failure occurred prior to a normal shutdown sequence between IMS and another logical unit, ensure that the session is restarted using the same half-session pairs (half-session names) that were in session at the time of failure. Further, when the session is resumed, the half-session pairs must maintain their previous relationship—that is, the former primary half session must again be primary, and the former secondary half session must be the current secondary.

Maintaining sequence numbers

To allow session recovery and message resynchronization across session and subsystem failures, both half sessions must maintain a checkpoint of sequence numbers and other indicators.

Specifically, a checkpoint must be maintained of the following:

- Three sequence numbers that are produced when the session was last active
The three sequence numbers are the potential (pending) and committed sequence numbers for the flow sent by the half session and the last committed sequence number for the flow received by the half session.
- An indicator produced when a unilateral decision was made by a half session to back out or commit a work unit that was left pending during the session outage
- An indication of the direction of the decision

Sequence number mismatches and incorrect decisions to commit or back out a recoverable work unit are detected by comparing the sequence numbers sent or received on the VTAM set-and-test-sequence-number (STSN) command with the checkpointed resynchronization information. Each half session can detect invalid sequence number mismatches on their outbound flow and reject the resynchronization request. Each half session can also agree or disagree to unilateral decisions made by the other half session to commit or back out a work unit. This agreement or disagreement is based on the detection of a sequence number mismatch on the inbound flow and involves a second STSN when an incorrect decision is made on the STSN receiver's inbound flow. The STSN command receiver does this before responding to the second STSN command (which is sent in response to the receipt of TEST NEGATIVE on the first STSN). This second STSN informs the STSN receiver that a wrong decision had been made on its inbound flow.

Recovering sessions with cold start

Invalid sequence number mismatches and mismatches where no unilateral decisions have been made often indicate a subsystem restart from an incorrect log or incorrect checkpoint on the log. This situation can require forcing the recovery of the ISC session by cold starting one or both half sessions.

IMS allows authorized terminal operators to change IMS ISC session state from recoverable to cold start when necessary by using the IMS /ASSIGN (subpool to VTAMPOOL) command before attempting to initiate a session. When a session is changed to a cold-start state, the subpool associated with that session is made available for allocation to any ISC cold-startable parallel session with any node.

Session recovery where one or both half sessions are cold started is not considered a major error because, by definition, the cold starting half session has no information from the previous active session with which to negotiate (agree or disagree) recovery or resynchronization.

Controlling unilateral decisions about pending work units

To prevent resources from accidentally getting out of synchronization between subsystems, you can specify whether to continue session resynchronization after a session outage, where the other half session has made a unilateral decision to commit or back out a pending unit of work.

IMS allows system definition to allow resynchronization without regard to inbound sequence number mismatches (OPTIONS=FORCESS on the TERMINAL macro statement or an ETO user descriptor), or to only allow resynchronization when the inbound sequence numbers agree (OPTIONS=SYNCESS). A keyword on the IMS /CHANGE command allows an authorized terminal operator to override the system definition specification for a single attempt to initiate a session. The effects of the CHANGE command are reset to the original system definition specification after one session initiation attempt. IMS does allow unilateral decisions to back out during session outages. However, using the IMS /DEQUEUE command for a session (terminal, and optionally, the subpool parameter) or LTERM during a session outage where an output-message sync point is pending is considered to be a unilateral decision by an authorized terminal operator to commit the pending output.

A pending output sync-point response can be determined by using the IMS /DISPLAY command. When two IMS subsystems are connected by an ISC session, the FORCESS option must be in effect on the opposite IMS subsystem from one where a pending output message was dequeued (committed) using a /DEQUEUE command, or session resynchronization fails with message DFS2065. Using FORCESS rather than SYNCESS has no other effect between two IMS subsystems.

Recovering from in-bracket failures

When a session has failed while in-brackets, it might be necessary to restart the failed bracket when the session is restarted.

The bracket state manager and the half-duplex reset flags in the bind can be set in either the BIND request or the negotiable bind response by either half session. These flags are set to in-brackets/SEND or RECEIVE to indicate the possibility that the process previously attached must be restarted.

Response or conversational output available at restart

If, during session initiation, IMS has response mode or conversational output available to be sent or is in a sync-point response pending state, the bracket state manager reset and half-duplex flags within the IMS portion of the bind or bind response are set to in-brackets/SEND by IMS.

If IMS is in conversation mode but has no conversational output available or pending, these same bind or bind reply flags are set to in-brackets/RECEIVE, because input is required to continue the conversation. IMS does not accept a

nonnegotiated BIND or negotiated BIND response indicating IMS to be in-brackets/RECEIVE when conversation or response mode output is available. Unless a sync point is pending from a previous session, the session is resumed and the STSN that is sent must receive a TEST POSITIVE reply to the SET AND TEST option in order for processing to continue.

When IMS is in conversation or response mode and receives a nonnegotiated BIND or a response to a negotiated BIND indicating between-brackets, IMS attempts to terminate the conversation or response mode. This termination is just like that which occurs if end-bracket is received on an FMH7 or LUSTATUS while IMS is in conversation or response mode during normal data flow active state (after SDT).

Conversation and response mode termination is only assured when the output reply message is available to be sent or the half session is waiting for the next conversational input. Therefore, IMS restricts attempts to terminate conversation or response mode by binding between-brackets when the output reply is available to be sent or no conversational input is required. If no output reply is available to be sent or no conversational input is required, a warning message is sent to the IMS master terminal operator with instructions to retry session initiation later.

If IMS receives a negotiated BIND indicating in-brackets/SEND and is not in conversation or response mode, a BIND response is sent indicating between-brackets. IMS sends LUSTATUS - NO-OP indicating end-bracket immediately after SDT when IMS is not in conversation or response mode and when receiving a non-negotiated BIND or negotiated BIND response indicating in-brackets/SEND. LUSTATUS - NO-OP is sent because transaction restart is not possible under these circumstances.

It is possible that IMS is left in-brackets/SEND after a BIND or BIND response, because a conversation or response mode output sync-point response is pending from the previous session. The output message might be dequeued if STSN processing indicates that the response was sent. In this case, an LUSTATUS - NO-OP indicating end-bracket is to be sent after SDT if the output message is a response mode reply. An LUSTATUS - NO-OP indicating change-direction is sent if the output message is a conversational reply, because input is required to continue the conversation.

If STSN indicates that the conversation or response mode message was not acknowledged, IMS again places the message in a sync-point response-pending state as if it had been sent with change-direction and as if it had requested an exception sync-point response. Any normal flow reply is then considered implicit acknowledgment of the pending output message and causes it to be dequeued by IMS even if the reply was an LUSTATUS - Abort or an ATTACH ATTDPN=SYSMSG. These replies indicate additional error conditions associated with the message's scheduling or execution. Had the session not failed, these error conditions would have been reflected by an exception response to the sync-point request rather than the LUSTATUS - Abort or ATTACH ATTDPN=SYSMSG. Depending on the sense code used, the exception response might have caused the message to be returned to the message queue (backed out) for retransmission rather than dequeued (committed) as occurs for LUSTATUS - Abort or ATTACH ATTDPN=SYSMSG resulting when both a session and an application failure occur together.

Session failures without IMS failure

A session failure not also involving an IMS subsystem failure might occur before IMS returns a requested sync-point response for a response mode or conversational input message.

When the session is reestablished, the DFC state set by IMS using the bind is IMS in-brackets/SEND. The subsequent STSN sequence number recovery might reflect either that the input message has been committed (the transaction has reached a sync point after inserting a reply message) or that the input message has not yet been committed. In either case, IMS being bound as in-brackets/SEND indicates that the input message has been received and the other half session should wait for the reply message. Because session restart resets the original sync-point request, the reply message now becomes an implicit sync-point response to the original input message.

IMS recovers the fact that the session was in conversational mode (and also recovers associated input messages and output replies) across IMS outages. The fact that the session was in response mode (and any associated output replies) is only recovered across an IMS outage if the failure occurs after the transaction sync point that made the response mode reply available.

Session failures because of IMS failure

A session failure that also involves an IMS subsystem failure might occur before IMS responds to a synchronization request on response mode (if the message was recoverable) or conversational mode input.

When the session is reestablished, the DFC state set by IMS by the bind indicates IMS as bound in-brackets/SEND. In either case, because of the IMS restart process from the subsystem log, the subsequent STSN sequence number recovery reflects that the input message has been committed even if that commit has not yet actually occurred. However, this condition produces inconsistent results only if the transaction subsequently abnormally terminates after being restarted. In this case, the response mode output reply message is made available for asynchronous (ATTACH EB or ATTACH ATTDPN=SCHEDULER) delivery on the recovered ISC session.

Recoverability of commands and execution modes

IMS commands (except /DIS, /RDIS, and /FOR) and test-mode input are executed immediately without being placed on an input message queue. These commands complete all processing prior to causing output to be made available using a transaction sync point. Any failure causes the command or test mode transaction input and associated output message to be backed out and discarded.

IMS commands /DIS, /RDIS, and /FOR produce asynchronous queued output. Output that is enqueued prior to the failure is recovered and made available for asynchronous transmission when the session is reestablished.

Fast Path, recoverable response mode, and recoverable conversational mode transactions are backed out and discarded if the failure occurs before a transaction sync point. The fact that the session was in response mode is not recovered.

Coordinating the restart process

Half sessions use rules to coordinate restart after a session failure.

The rules used by the half sessions include:

- The restart always occurs from the most currently completed sync point.
- The session bind establishes the half-duplex state at the current sync point.
- When the primary half session wants to restart:
 - It sends a BIND request to select the proper half-duplex state.

- The secondary half session cannot change this state in the BIND response, unless the secondary half session does not want to restart.
- When the secondary half session does not want to restart, it sends a BIND response to preclude restart by setting the bracket state to between-brackets.
- When the primary half session does not want to restart:
 - The primary half session sends a BIND request setting the bracket state to between-brackets.
 - If the secondary half session wants to restart, it sends a BIND response that establishes the proper half-duplex state and sets the bracket state to in-brackets.
 - If the secondary half session does not want to restart, the secondary half session's response to the BIND does not change (it matches the state set by the primary half session).
- STSN is used to resynchronize any pending requests for sync-point responses.
 - If no mismatch occurred, agreement exists on the current sync point. The restart is attempted after sending start data traffic (SDT).
 - It might be necessary for the half session in send state to send LUSTATUS - NO-OP, CD to adjust the half-duplex state to the current sync point based on the STSN.
 - If a mismatch occurred and the session continues, the half session in send state sends LUSTATUS - NO-OP, EB to place the session into contention.
- The session in send state at the final restart point reestablishes the session by sending an explicit ATTACH.

Half sessions assuming a secondary role must reject a session bind in these situations:

- Session bind parameters indicate in-brackets, "other half session speaks first", and response mode or conversational output (or equivalent) are immediately available at data traffic active state. This is a session restart logic error.
- On restart following an abnormal failure or shutdown sequence, session bind parameters indicate half-session names different from those in effect during the previous session between IMS and another logical unit.

In an ISC session, either subsystem can assume a primary or a secondary half-session role, and either subsystem can request message resynchronization; therefore, a set of rules must be established by which these half-session pairs can maintain message or sync point integrity.

These topics describe the format of the STSN and how the STSN is used to complete resynchronization and recovery.

Related reference:

"LUSTATUS protocol" on page 546

Determining session synchronism using STSN

The need to resynchronize can be communicated during bind negotiation; two flags in the BIND request are used to determine the requirement to resynchronize and to return the half sessions to the state that existed at the time of session termination.

If both half sessions have terminated normally (that is, neither had any outstanding traffic to handle), the session is restarted as though it were a new session—bind parameters are sent, and no requirement exists that the relative

positions of the session partners be maintained, nor that the same half-session names be used. The two flags upon which this determination is made are:

- Sequence number indicator
 - 1 = Sequence numbers available
 - 0 = Sequence numbers not available
- Bracket initiation stopped (BIS) indicator
 - 1 = BIS sent
 - 0 = BIS not sent

The following table is a matrix describing the states set by these two flags in relation both to the bind sender (primary half session, or PHS) and the bind receiver (secondary half session, or SHS).

When both half sessions are in a “COLD START” state, no sequence numbers are available or required to be sent. Session shutdown is such that resynchronization is not required. The bind is negotiated and the session started in accordance with the bind parameters. For a nonnegotiable BIND, some information required by both half sessions is not available until STSN flows exist. Therefore, for a nonnegotiable BIND, STSN is always sent.

Table 73. BIND action/response matrix

Bind sender (PHS):	Bind receiver (SHS): Numbers not Available	Bind receiver (SHS): Numbers available, BIS sent	Bind receiver (SHS): Numbers available, BIS not sent (See Table 74 on page 518)
Numbers Not available	COLD STSN not required Send RESET to NOBB	COLD STSN not required Send RESET to NOBB	COLD/WARM mismatch STSN sent
Numbers available BIS sent	COLD STSN not required Send RESET to NOBB	COLD STSN not required Send RESET to NOBB	WARM STSN sent
Numbers available BIS not sent (Table 75 on page 519)	WARM/COLD mismatch STSN sent	WARM STSN sent	WARM STSN sent

IMS sets the “BIS sent” and “sequence numbers available” flags in the negotiable BIND and BIND response (or PHS and SHS, respectively) as follows:

- Session cold start following IMS cold start or normal session termination:
 - BIS not sent
 - Sequence numbers not available
- Session restart
 - Previous session terminated normally following attempted normal termination by IMS:
 - BIS sent
 - Sequence numbers available
 - Previous session terminated abnormally (normal termination by IMS not attempted):
 - BIS not sent

Sequence numbers available

Related reference:

“Performing the resynchronization”

“Symmetrical session shutdown for LU 6.1 (SBI and BIS)” on page 558

“STSN command format” on page 520

Performing the resynchronization

When a session is reestablished (or established for the first time), the state of each flow is represented by the sequence numbers that are available for synchronization in each half session.

The sequence numbers available for synchronization in each half session can represent the following states for each flow:

- **COLD**

This session has no sequence numbers to send or to compare. The primary half session indicates this by resetting the sequence number indicator bit in the BIND request. If the BIND is negotiable, the secondary half session indicates this by the same bit setting in the BIND response. If STSN is required (as for nonnegotiable BIND), the primary half session sends a SET AND TEST action code with zeros as the sequence number values.

- **NOT PENDING**

This session has no work units that are indoubt (that is, waiting to be committed)

- **PENDING**

One work unit is in doubt (waiting to be committed). The sequence numbers available to the session are the committed number (the sequence number of the last message to be committed) and the potential number (the sequence number of the last message to be issued).

- **DECISION TO COMMIT**

An indoubt work unit is committed during the session outage. Potential and committed numbers are still available. (A decision to commit occurs only as the result of a /DEQ command being issued during a session outage.)

- **DECISION TO BACK OUT**

An indoubt work unit is backed out during the session outage. Potential and committed numbers are still available. (IMS does not back out work units, but other session partners might).

- **INVALID**

Invalid sequence numbers are detected. (An invalid sequence number is one that should not occur, as, for example, a sequence number sent on the secondary-to-primary flow that does not match the committed or potential numbers saved by the STSN receiver.) This type of condition usually indicates an incorrect log data set; the master terminal operator must intervene to recover the session.

- **RECEIVED PENDING (Sync-Point Response Lost)**

The receiver of this flow receives the entire work unit that the sender sends and generates a DR2 that might be lost in the network.

- **NOT RECEIVED PENDING (Sync-Point Response Not Lost)**

The receiver of this flow never receives the RQ*2 that the sender sends and therefore backs out the work unit that the sender is pending on.

The half sessions exchange information on their respective states by using bind negotiation and the action codes in the STSN request and response.

Related reference:

“Using STSN to resynchronize sessions” on page 517

Completing session initiation

After the half sessions are in sync, either half session can decide to accept or reject the session.

- If the primary half session does not want to continue the session, it sends UNBIND.
- If the primary half session wants to continue the session, it sends start data traffic (SDT).
 - If the secondary half session responds positively, the session is formally bound and traffic on the session can begin.
 - If the secondary half session wants to discontinue the session, it rejects the primary half session's SDT, whereupon the primary must respond with an UNBIND and the session terminates.

Session initiation, including allocation of the LTERM subpool, is completed when a start data traffic (SDT) response has been received if IMS is the primary half session or when an SDT response is sent if IMS is the secondary half session. The LTERM subpool remains allocated to the other half-session name, even across session or subsystem failures, until the session is terminated by mutual consent of both half sessions (symmetrical shutdown). This requires that, after they have been allocated, all subsequent session binds for the same half session must specify the same bind session qualifiers that were active at abnormal session termination until the subpool is released through normal session termination.

If the session is terminated prior to the completion of SDT, any newly attempted subpool allocation caused by a session cold start is backed out, and the subpool is returned to an available-for-allocation status.

In IMS, the master terminal operator must always be notified of any session that is rejected prior to a start data traffic (SDT) completion. In an ISC session, notification to the IMS master terminal operator is optional for normal initiation and termination sequences. Operator notification is specified by the OPTIONS keyword parameters MTOMSG and NOMTOMSG on the TYPE and TERMINAL macros or by an ETO logon descriptor.

Running the session

Use the data flow control protocols to control the flow of data in an ISC session.

Related concepts:

“Handling IMS response mode or conversational output errors” on page 523

Related reference:

“Data flow control protocol reference” on page 533

Terminating an ISC VTAM session

Terminating an ISC VTAM session releases a logical unit from its current connection to the VTAM application program, making the LU available for sessions with other VTAM applications, or terminating communications altogether.

Definitions: The two types of session termination are *normal* and *abnormal*. Normal termination allows both half sessions to complete normal processing before the session is terminated. Abnormal termination forces the session to terminate unconditionally.

Session termination can be invoked by the IMS master terminal operator, the VTAM network operator, or either half session.

Because of this variety of session termination methods, each IMS network installation must determine specific procedures for session termination. When developing these procedures, consider the requirements for session-termination processing.

Normal termination

Normal termination of an ISC session occurs with the flow of the data flow control indicators stop bracket initiation (SBI) and bracket initiation stopped (BIS). Normal termination can be initiated by either half session.

Normal session termination also occurs when the MTO invokes an orderly termination of the IMS network by the IMS /CHECKPOINT command with the FREEZE, PURGE, or DUMPQ parameter and the QUIESCE parameter. The QUIESCE parameter ensures that message queues are emptied before the session is terminated.

When all terminals have indicated that shutdown is complete, IMS:

- Performs checkpoint processing and then issues the VTAM CLSDST macro instruction, when acting as primary half session.
- Awaits UNBIND, if acting as secondary half session.

CLSDST causes VTAM to send the UNBIND command. CLSDST or UNBIND releases the logical units from session with IMS. Any further data transmission to the logical units is prohibited.

During the processing of an orderly session termination, the IMS master terminal operator can terminate the network unconditionally by using an IMS /CLSDST, /STOP, or /CHECKPOINT command.

Abnormal termination

Abnormal session termination can occur as a result of transmission or protocol errors, or errors in data that make that data unacceptable to the receiving message processing program. Because an ISC session involves two peer-level systems, error recovery processing and abnormal session termination processes can differ. IMS-detected error conditions requiring abnormal session termination result in IMS issuing a VTAM CLSDST macro when IMS is the primary half session, or TERMSESS with OPTCD=UNCOND when IMS is the secondary half session.

Related reference:

“Symmetrical session shutdown for LU 6.1 (SBI and BIS)” on page 558

“Data flow control protocol reference” on page 533

Using STSN to resynchronize sessions

The tables in the following topics show the results and actions related to using STSN for primary-to-secondary (P-S) and secondary-to-primary (S-P) flows.

Primary-to-secondary flow matrix

The following table illustrates the actions taken when resynchronizing a session and the STSN command flows from the primary half session to the secondary half session.

The first result in each table cell represents the STSN command action; the second result represents the response returned to the STSN command. These are separated by a slash (/). When two STSNs are sent, they are separated by a comma.

Table 74. STSN primary-to-secondary flow

Primary half session	Secondary half session, session cold	Secondary half session, not received: pending	Secondary half session, received: pending
Pending	1 Set and Test / Reset	2 Set and Test / Negative	3 Set and Test / Positive
Decision to commit	4 Set and Test / Reset	5 Set and Test / Negative, Set / Positive	6 Set and Test / Positive
Decision to back out	7 Set and Test / Reset	8 Set and Test / Positive	9 Set and Test / Negative, Set/Positive
Not pending	10 Set and Test / Reset	11 Set and Test / Positive	Not Applicable

Notes to the preceding table

1 The STSN sender is PENDING. The STSN receiver has no sequence numbers. A COLD/WARM mismatch has occurred. The session can continue. The STSN sender should continue the session.

2 The STSN sender is PENDING. The STSN receiver never received the pending RQ*2 and returns a TEST NEGATIVE in the STSN response. The STSN sender backs out the pending work unit to the last commit point.

3 The STSN sender is PENDING. The STSN receiver received the pending RQ*2 and sent a + DR2. The STSN receiver responds TEST POSITIVE, and the STSN sender commits the pending work unit.

4 The STSN sender wants to release the locks and commit resources; however, an RQ*2 is outstanding. The STSN receiver has no sequence numbers. A COLD/WARM mismatch has occurred. The session can continue. The STSN sender should continue the session.

5 The STSN sender wants to release the locks and commit resources; however, an RQ*2 is outstanding. The STSN receiver never received the pending chain, as indicated by TEST NEGATIVE in the STSN response. The STSN sender sends a second STSN with the SET action code to inform the STSN receiver that a wrong decision was made. The STSN receiver can continue the session by sending TEST POSITIVE or decline by sending INVALID.

6 The STSN sender wants to release locks and commit resources; however, an RQ*2 is outstanding. The STSN receiver received the pending chain and committed it, as indicated by the TEST POSITIVE in the STSN response.

7 The STSN sender wants to back out the pending work unit, but a pending RQ*2 is outstanding. The STSN receiver has no sequence numbers. A COLD/WARM mismatch has occurred. The session can continue. The STSN sender should continue the session.

8 The STSN sender wants to back out the pending work unit, but an RQ*2 is outstanding. The STSN receiver never received the pending chain, so the sequence numbers match. A TEST POSITIVE is returned on the STSN response.

9 The STSN sender wants to back out the pending work unit, but an RQ*2 is outstanding. The STSN receiver received the pending chain and committed resources. Because the STSN sender sent out the old committed sequence numbers and the STSN receiver received the chain, a TEST NEGATIVE is returned. The STSN sender informs the STSN receiver that a wrong decision has been made by sending a second STSN with a SET action code. The STSN receiver can continue the session by sending TEST POSITIVE or decline by sending INVALID.

10 The STSN sender is not pending on the P-S flow. The STSN receiver has no sequence numbers. The session continues.

11 The STSN sender is not pending on the P-S flow. The STSN receiver agrees with the sequence numbers on the STSN and returns a TEST POSITIVE.

Secondary-to-primary flow matrix

The following table illustrates the actions taken when resynchronizing a session and the STSN command flows from the secondary half session to the primary half session.

The first result in each table cell represents the STSN command action; the second result represents the response returned to the STSN command. These are separated by a slash (/). When two STSNs are sent, they are separated by a comma.

Table 75. STSN secondary-to-primary flow

Primary	Secondary, session cold	Secondary, pending	Secondary, decision to commit	Secondary, decision to back out	Secondary, invalid	Secondary, not pending
Not received: pending	1 Set and Test / Reset	2 Set and Test/ Positive	3 Set and Test/ Negative	4 Set and Test/ Positive	5 Set and Test / Invalid	6 Set and Test / Positive
Received: pending	7 Set and Test / Reset	8 Set and Test/ Positive	9 Set and Test/ Positive	10 Set and Test/ Negative	11 Set and Test / Invalid	Not Applicable

Notes to the preceding table

- 1** The STSN sender is not COLD. The STSN receiver has no sequence numbers. The session continues.
- 2** The STSN receiver has a pending work unit and has held onto locks and not committed resources. The sequence number sent on the SET AND TEST action code indicates that the STSN sender did not receive the

pending RQ*2. Therefore, the STSN receiver backs out the pending work unit to the last commit point and returns a TEST POSITIVE on the STSN response.

- 3** The STSN receiver had a work unit pending and decides to commit resources and release locks. The STSN sender did not receive the RQ*2 sent by the STSN receiver. The STSN receiver returns a TEST NEGATIVE on the STSN response to inform the STSN sender that a wrong decision was made. The STSN sender can continue the session by sending SDT or decline by sending UNBIND.
- 4** The STSN receiver had a work unit pending and decides to back out the pending work unit. The STSN sender never received the pending work unit, and the STSN receiver returns a TEST POSITIVE on the STSN response.
- 5** The STSN receiver detects a severe loss of synchronization (possible log data set mismatch) and returns an INVALID on the STSN response.
- 6** The STSN receiver has no pending work unit and agrees with the sequence numbers sent on the STSN; therefore, it replies with a TEST POSITIVE.
- 7** The STSN sender is not COLD. The STSN receiver has no sequence numbers. The session continues.
- 8** The STSN receiver has a pending work unit and still holds resource locks. The number sent by the STSN sender on the S-P SET AND TEST indicate to the STSN receiver that the STSN sender received the pending work unit. The STSN receiver commits the pending work unit and responds TEST POSITIVE.
- 9** The STSN receiver had a pending work unit and decides to commit the pending work unit. The number sent by the STSN sender on the S-P flow indicates that the STSN sender received the RQ*2 and a TEST POSITIVE is sent on the STSN response.
- 10** The STSN receiver had a pending work unit. The receiver decides to back out that work unit and release the locks. The number sent by the STSN sender on the S-P flow indicates to the STSN receiver that the RQ*2 was received and processed by the STSN sender. The STSN receiver indicates the wrong decision by a TEST NEGATIVE on the STSN response. The STSN sender can continue the session by sending SDT or decline by sending UNBIND.
- 11** The STSN receiver detects a severe loss of synchronization (possible log data set mismatch) and indicates this to the STSN sender by an INVALID in the STSN response.

STSN command format

Both the STSN command and STSN response contain a 5-byte data field.

The format of the STSN command is:

Byte 0 Action code

Bytes 1, 2

Sequence number of the last inbound sync point message sent to the PHS

Bytes 3, 4

Sequence number of the last outbound sync point message sent from the PHS

The primary half session uses the action code to ask the secondary half session to verify the VTAM sequence numbers. The bits of the action code byte are as follows:

Bits 0 and 1

Refer to the inbound sequence-number field

Bits 2 and 3

Refer to the outbound sequence-number field

Bits 4 through 7

Reserved

These values are acceptable for bits 0, 1, 2, and 3 of the STSN command action code:

00 IGNORE

Do not alter value. IMS does not send this action code value as PHS. IMS as SHS returns an INVALID response code for this action code when resynchronization is required and TEST POSITIVE when resynchronization is not required.

01 SET

Set the appropriate sequence number to the value indicated in the sequence number field. For ISC, this code only occurs when the PHS must send a second STSN to indicate a unilateral decision was made to commit or back out a pending work unit.

10 SENSE

Do not alter value. The SHS should return its version of the sequence number in the command response. IMS as SHS returns an INVALID response to this action code.

11 SET AND TEST

Set the appropriate sequence number to the value indicated in the sequence number field. For ISC, this code always occurs for both flows on the first STSN sent by the PHS. The SHS must indicate in the command response whether the sequence number values are acceptable.

When the secondary half session receives the STSN command, it must be able to:

- Verify the SHS-outbound (PHS-inbound) sequence number and arrange to retransmit a message to the PHS if required.
- Verify the SHS-inbound (PHS-outbound) sequence number and inform the PHS whether the number matches the SHS-saved number. Further, the SHS might or might not agree with unilateral decisions by the PHS to commit or back out a pending work unit from the previous session.
- Return to the PHS a DR1 and a 5-byte data response to the STSN command.

The format of the STSN response has the same format as the STSN command:

Byte 0 Action code

Bytes 1, 2

Sequence number of the last recoverable message the SHS sent to the PHS

Bytes 3, 4

Sequence number of the last recoverable message the SHS received from the PHS

The secondary half session (SHS) uses the action code to indicate the test results. The action code must be returned to the primary half session (PHS). Returning the sequence numbers (bytes 1 through 4) is optional. However, for debugging recovery and restart problems, always return these to the PHS whenever the STSN numbers do not match those maintained by the SHS. The bits of the action code byte are as follows:

Bits 0 and 1

Refer to the SHS inbound sequence number

Bits 2 and 3

Refer to the SHS outbound sequence number

Bits 4 - 7

Reserved

These values are acceptable for bits 0, 1, 2, and 3 of the STSN response action code:

00 RESET

Returned on both sequence number flows in response to the SET AND TEST option to indicate that the SHS must cold start and has no sequence number information. The PHS should continue session initiation and should not treat this response as an error. IMS sends this code when cold-starting and replying to STSN, and continues the session when receiving it.

01 TEST POSITIVE

Returned in response to the SET AND TEST option to indicate that the sequence number agrees with the sequence number that the SHS checkpointed. This code should be returned in response to the SET option when the SHS agrees to continue session initiation following a unilateral decision by the PHS to commit or back out a pending work unit on its outbound flow during the session outage.

10 INVALID

Returned in response to the SET AND TEST option to indicate that the STSN SHS outbound sequence number does not agree with the sequence number that the SHS checkpointed (major session restart mismatch). INVALID is also returned in response to the SET option to indicate disagreement with a unilateral action by the PHS to commit or back out a pending work unit on its outbound flow. The PHS should not continue the session when receiving an INVALID response to either flow.

11 TEST NEGATIVE

Returned in response to the SET AND TEST option on the SHS inbound flow to indicate that the sequence number does not agree with the sequence number checkpointed by the SHS. The PHS responds with SDT if the reason for the mismatch is a pending work unit sent by the PHS, but is not received by the SHS; with a second STSN indicating the SET option for both flows if the PHS had made a unilateral decision to commit or back out a pending work unit; or UNBIND if no work unit was pending and no unilateral decision had been made (major session restart mismatch). TEST NEGATIVE can also be returned in response to the SET AND TEST option to indicate that the SHS had made a unilateral decision to commit or back out a pending work unit on its outbound flow. In this case, the PHS can optionally agree to continue session initiation by sending SET or not to continue session initiation by sending UNBIND.

Related tasks:

“Controlling unilateral decisions about pending work units” on page 510

Handling IMS response mode or conversational output errors

This topic describes how IMS handles response and conversational mode errors during an ISC session and how to keep the half sessions in sync.

IMS messages cannot be canceled by session termination or protocols after they have been received and enqueued (made available for scheduling). However, errors can be detected by either half session after the input message has been enqueued. If the error is detected by IMS while processing the transaction, but before the reply message is committed (for example, a DFS555 transaction abend occurs), both half sessions are able to back out, because IMS holds the input sync-point response until the reply message is available for output. An exception response and appropriate error recovery process (ERP) message are returned. Errors occurring after the reply message is made available for output result in backout to the last application sync point (the one that made the reply message available). These errors are not communicated to the other half session.

After an exception response to a DFC command, you might need to purge before further sends or receives.

Related reading: For information on the DFC protocols, see *z/OS Communications Server: SNA Programming*.

Related reference:

“Data flow control protocol reference” on page 533

“ERP purging” on page 542

Response mode errors

Errors detected by the other half session cannot be communicated to IMS until IMS attempts to send the reply.

At this point, errors must be communicated to IMS before the requested sync-point response is returned, because all response mode output is sent indicating RQD and EB. Errors are reflected by the return of an exception response with an optional ERP message, which might contain appropriate sense data and protocols.

The exception response or ERP FM header sense codes can cause the output message to be backed out (dequeued) or retransmitted, or they can cause the session to terminate with the message still on the queue. ERP message protocols can cause the message to be backed out (dequeued) regardless of the sense code through use of EB or can leave IMS in send state through use of CD. This latter case is only allowed if the sense code used results in the output message being dequeued or retransmitted.

An LUSTATUS - abort with EB is sent as the next output on the session if CD is used on the ERP message and the sense code results in dequeuing the response mode output message. However, in all of the cases where the error is detected by the other half session, the reply message is the only resource that IMS can back out. The database updates and other messages initiated by the response mode transaction cannot be backed out by IMS.

Conversational mode errors

Errors detected by the other half session cannot be communicated to IMS until IMS attempts to send the output reply.

Errors must be communicated to IMS before the requested sync-point response is completed. For nonlast conversational output, the sync-point requested using RQE2 with CD and any returned normal flow data implicitly completes the sync point. The last conversational output message is sent requesting RQD2 with EB. Errors are reflected either by returning an exception response with an optional ERP message (which might contain appropriate sense data and protocols) or by returning an LUSTATUS with appropriate sense codes and protocols.

The sense codes can cause the message to be backed out (dequeued) or retransmitted, or they can cause the session to terminate with the message still on the queue. ERP and LUSTATUS message protocols can cause the message to be backed out (dequeued) using EB, or they can leave IMS in send state using CD. This latter case is only allowed if the sense code retransmits the message or if the last conversational output for the sense code dequeues the output message. An LUSTATUS - abort with EB is sent as the next output on the session if CD is used on the LUSTATUS or ERP message, and the sense code results in dequeuing the last conversational mode output message.

When the error is detected by the other half session, the output reply message is the only resource that IMS can automatically back out. The database updates, conversational SPA, and other messages initiated by the conversational transaction cannot be automatically backed out by IMS. However, for conversations, an internal IMS /EXIT command is scheduled to invoke the user's Conversational Abnormal Termination exit routine, which can optionally schedule a transaction to reverse necessary database changes based on the conversational SPA content. The interface to the exit routine is the same as if an /EXIT command had been received on the session.

Related reference:

 Conversational Abnormal Termination exit routine (DFSCONE0) (Exit Routines)

Normal conversation termination extension with ISC

In a non-ISC environment, normal termination of IMS conversation mode occurs when the IMS transaction creates a blank in the transaction-code field of the conversational SPA prior to committing the reply message. The conversation ends when the reply message has been successfully dequeued.

An extension for normal termination has been made for ISC, because the peer-level application in the other half session can also now end the conversation. However, IMS supports this request for normal termination by receiving only an SNA-defined commit request - LUSTATUS X'0006' with EB. This occurs because of a stand-alone commit request (or normal termination) by the remote application to its subsystem. This form of conversational termination requires notification of the user using the Conversational Abnormal Termination exit routine (new input vector of X'28' in byte 3 of register 1), because committing changes might have been deferred by the completed conversational steps. This is possible by recording the deferred information within the conversational SPA. The exit routine can schedule an appropriate transaction to commit the deferred changes upon being invoked with the new vector.

Restriction: IMS does not support an input message with EB to terminate the conversation.

Keeping half sessions synchronized

Sync-point responses (DR2) are used between ISC session partners to ensure that both partners' sync-point managers can commit or back out recoverable resources synchronously.

All messages sent or received on an IMS ISC session are defined as either recoverable or irrecoverable, depending on the message type. The session-response protocols are used to ensure that both ends of an ISC session mutually understand the recoverability attributes associated with each message. The response protocol used must be consistent with the IMS message type.

The ISC sync point can be explicit or implicit:

Explicit

The input requests an RQD2 response.

Implicit

The input requests the sending of an exception response (RQE2) and change-direction on output made available by the transaction, or the sending of an LUSTATUS in lieu of output data from the transaction).

To further increase integrity and recoverability, IMS can supplement the sync point facility by logging the information. Use of log write-ahead ensures that the sync point indication is recorded on the log. This makes sync point information available to IMS restart procedures before the sync-point response is actually sent or change-direction is replied to (implied sync point).

Before reading these topics, be sure to understand the definitions and relationships of the ISC sync point and associated commit and backout processes to the IMS application program, described in these paragraphs.

Related concepts:

 Keeping half sessions synchronized (Communications and Connections)

Sync points requested on input to IMS

For any type of input, IMS does not schedule the intended transaction until the complete input message is successfully received.

Sender-detected errors, errors resulting from processing of IMS input, and session failure prior to the receipt of the complete input message cause the entire message to be discarded or backed out. However, the input message cannot be canceled by ISC session failures or protocols after the complete message is received and made available for scheduling.

When an input message is backed out because of errors detected during IMS input processing or during synchronous transaction execution, the session partner is notified, either by means of session termination or by an exception response to the ISC input. These events occur for this backout:

- Backout results in resetting the associated DFC and ATTACH states to those of the last sync point.
- Backout during transaction execution results in backing out application updates and messages, except express messages made after the last application sync point.
- The application sync point and ISC sync point are not necessarily the same.

- Backout during input to IMS has no effect on other recoverable IMS resources, such as databases, because input messages are not available for scheduling or execution until they are received completely and without error.
- The result of the backout is only the current input message, even if several consecutive input (irrecoverable) messages were received and executed or enqueued for scheduling after the last input sync point was requested from IMS.

The definition and relationship of successful ISC input sync points to IMS application sync points depend upon whether the ISC input was synchronous or asynchronous. For asynchronous input, the sync point reflects only that IMS is now responsible for message recovery. No implications exist relative to the scheduling or execution of transactions or to the availability of transaction output. After the message is successfully enqueued, the DFC and ATTACH sync point information is updated and the requested sync-point response is returned to the session partner.

Although some IMS exceptions for synchronous input exist, the sync point is intended to reflect the success of the IMS transaction execution and sync point. IMS updates the DFC and ATTACH sync point information as appropriate and commits all associated transaction resources (for example, DL/I databases) and the output transaction message reply when responding to the ISC attached input sync-point request.

These exceptions apply to sync points for synchronous input to IMS:

- The ISC sync-point response is returned by IMS when a transaction-inserted response mode or conversational reply message (first message inserted to the I/O PCB or alternate response PCB) is made available for output using a transaction sync point. Additional transaction processing and sync points are not reflected in the ISC session. No ISC sync point-response occurs for transaction sync points prior to the transaction's inserting the reply message, as in a program-to-program switch. Nonfirst messages inserted to the I/O PCB or alternate response PCB, messages inserted to nonresponse PCBs, and IMS express messages intended for the ISC session are queued for future asynchronous delivery after successful delivery of the reply message.
- Except during program-to-program switches to another conversational transaction, IMS generates a subsystem error message as the result of abnormally terminating a conversational transaction that attempts to cause a transaction sync point without first having inserted an output conversational reply message. The subsystem error message causes an exception response to be sent to the input sync-point request and the input message to be backed out.

Sync points requested on output by IMS

IMS commits the output message when the requested sync-point response is returned by the other session partner. The message might also be committed as the result of some sense codes that are returned on an exception response to the requested sync-point response.

Definition: The term *commit* means that the message has been successfully sent and dequeued, and sync point information has been updated as appropriate.

Use of the IMS /DEQUEUE command during a session outage while an output sync-point response is pending is considered a unilateral decision by an authorized terminal operator to commit the sending output message.

Depending upon the sense code used, IMS backs out the output message when an exception response is returned to the sync-point request or when the IMS /DEQUEUE command is issued by an authorized terminal operator before IMS requests a sync-point response.

Definition: The term *backout* means that a recoverable message is returned to the message queue (unless dequeued by a /DEQUEUE command) for subsequent retransmission. An irrecoverable message is either dequeued or returned to the message queue for subsequent retransmission, depending upon the type of error. Normally, an exception response or IMS failure causes an irrecoverable message to be dequeued. Some session failures that do not result from a subsystem failure cause an irrecoverable message to be retransmitted at the first opportunity after resynchronization. Backout results in resetting the associated DFC and ATTACH states to those of the last sync point.

Sync point and response requirements

The IMS input/output message flow can be represented as an input/output flow from a sequential queue data set.

In order to maintain integrity and recoverability of this queue data set, IMS requires that each recoverable input and output message establish a sync point between both half sessions before continuing the flow. This allows both half session sync-point managers to commit or back out resources in synchronism.

The sync point facility for ISC redefines and separates the DR1 and DR2 requests and responses. The DR2 requests and responses are known as sync-point requests and responses and are functionally independent from those associated with DR1.

Recoverable messages

To ensure that a recoverable transaction can be recovered, IMS requires the following response protocols for each recoverable message sent or received.

- Messages that are not MFS-paged messages:

An exception DR2 (RQE2) must be requested on each nonlast (or when it is not the only) RU of an SNA chain.

Except when change-direction is sent, each last or only RU of an SNA chain must request a DR2 (RQD2). Either exception DR2 (RQE2) or RQD2 can be requested when change-direction is sent.

A sync-point response on the last or only RU of an SNA chain always indicates end-of-message. End-of-message always occurs at end-of-chain for single chain (non-MFS-paged) messages.

- First chain of asynchronous (ATTACH SCHEDULER) demand-paged message:

The first chain (OIC) of asynchronous demand-paged messages is an ATTACH for the SCHEDULER model. This OIC requests a DR2 (RQD2) with end-bracket. This allows the application to be asynchronously scheduled to receive the message and leaves the session in a state to be allocated to the scheduled application.

- Nonlast pages of MFS demand-paged (output using ATTACH or ATTACH SCHEDULER) messages, and last pages of MFS-operator logical-paged (OLP) output:

An exception DR1 (RQE1) is requested on each nonlast (or when it is not the only) RU of an SNA chain (MFS demand-page).

Each last or only RU of an SNA chain (MFS demand-page) requests exception DR1 (RQE1) with change-direction.

- Last pages of MFS demand-paged (output) messages:
An exception DR2 (RQE2) is requested on each nonlast (or when it is not the only) RU of an SNA chain (MFS demand-page).
Except when change-direction is sent on the last page, each last or only RU of an SNA chain (MFS demand-page) requests a DR2 (RQD2). Exception DR2 (RQE2) with change-direction is requested.
A sync-point response on the last or only RU of an SNA chain always indicates end-of-message. A sync point must always be requested on the last page of MFS demand-paged messages to ensure end-of-message.
- First, nonlast pages of MFS-autopaged messages:
An exception DR1 (RQE1) is requested on each nonlast (or when it is not the only) RU of an SNA chain (MFS autopage) sent to IMS and must be requested on each nonlast (or when it is not the only) RU of an SNA chain (MFS autopage) received by IMS.
RQD1 is requested on the last or only RU of the first SNA chain of an MFS-autopaged output message sent by IMS.
RQD1 must be requested on the last or only RU of the first SNA chain of an MFS-autopaged input message received by IMS when the other half session is initialized as the primary half session (bidder). Either exception DR1 (RQE1) or DR1 (RQD1) can be requested on the last or only RU of this SNA chain when the other half session is initialized as the secondary half session (first speaker). Change-direction is not sent and must not be requested on nonlast pages of MFS-autopaged input and output messages.
- Nonfirst, nonlast pages of MFS-autopaged messages:
An exception DR1 (RQE1) is requested on each RU (including last or only) of an SNA chain (MFS autopage) sent by IMS. An exception DR1 (RQE1) must be requested on each nonlast (or when it is not the only) RU of an SNA chain received by IMS. Either exception DR1 (RQE1) or RQD1 can be requested on each last or only RU of an SNA chain received by IMS. Change-direction is not sent and must not be requested on nonlast pages of MFS-autopaged input and output.
- Last pages of MFS autopaged messages:
An exception DR2 (RQE2) is requested on each nonlast RU of an SNA chain (MFS autopage) sent by IMS and must be requested on each nonlast RU of an SNA chain (MFS autopage) received by IMS.
Except when change-direction is sent on the last page, each last or only RU of an SNA chain (MFS autopage) requests a DR2 (RQD2). Exception DR2 (RQE2) is requested when change-direction is sent.
Except when change-direction is received on the last page, each last or only RU of an SNA chain of an MFS-autopaged input must request a DR2 (RQD2). Either exception DR2 (RQE2) or definite response 2 (RQD2) can be requested when change-direction is indicated.
A sync-point response on the last or only RU of an SNA chain or on an LUSTATUS - commit always indicates end-of-message. A sync point must always be requested on the last page or on an LUSTATUS - commit following the last page of autopaged messages to ensure end-of-message.

For the case when allowing RQE1 on the last or only RU of an SNA chain that does not indicate a change-direction for MFS-autopaged input or output, an exception to the preceding protocols occurs when the session bind indicates DEFINITE RESPONSE CHAINS. If the definite response chains parameter is set for the IMS half session, each last or only RU defined that does not also indicate

change-direction is sent requesting RQD1 or RQD2. RQD1 or RQD2 must be requested by the other half session under these same conditions if its session bind indicates DEFINITE RESPONSE CHAINS. Both exception and definite response (DR1 and DR2) are valid if change-direction is indicated under definite response chain rules.

Related reference:

“LUSTATUS protocol” on page 546

Irrecoverable messages

IMS treats an irrecoverable message in the same manner as a recoverable one, except that all processing required to achieve recoverability is eliminated. As a result, irrecoverable messages require less processing time, but can be lost in the event of a failure.

Irrecoverable, non-MFS input and output messages basically have the same requirements as those for recoverable messages, except that DR1 and exception DR1 can optionally be requested instead of DR2 and exception DR2, respectively. Irrecoverable MFS-paged input and output messages have the same requirements as recoverable ones.

When messages are being sent, only one message can be outstanding at a time. This means that the sender can send one message and must wait for the response or reply before sending another.

The required response and sync-point protocol allows message integrity to be maintained by allowing the half sessions' sync-point managers to mutually understand when messages are accepted (committed) or rejected (backed out). This also allows change-direction to be solicited using SIGNAL RCD as required. RQE1 or RQE2 is recommended where change-direction is indicated, because these capabilities are automatic. An RQD1 or RQD2 are valid, but reduce performance because of the unnecessary response. A response or sync point is implied when a reply is received to a sent message indicating change-direction and RQE1 or RQE2. That is, the reply is an implied DR1 or DR2 response. Requesting either a definite response (RQD1 or RQD2) or an exception response (RQE1 or RQE2) with change-direction is valid for the session bind option DEFINITE RESPONSE CHAINS.

A failure can occur between sending a recoverable message and receiving the sync-point response (or reply) and receiving the sync-point response (or reply message). During the session restart procedure, the STSN command is used to inform both half sessions of the sequence number of the last sent or received sync-point message. If either half session has not completely received a given message, that message can then be retransmitted.

Sync-point indicators on messages

The following topics describe the sync-point indicators sent on IMS input and output messages.

Requests on IMS input messages

The response and sync-point requests for input messages to IMS are summarized in the following table.

In the table, an “X” in the table indicates the entry is supported by IMS. An “S” in the table indicates the corresponding entry is suggested.

Table 76. Response and sync-point requests for IMS input messages.

Input message type	VTAM indicators with message							
	RQE1 ¹		RQD1		RQE2 ¹		RQD2	
	CD	-CD	CD	-CD	CD	-CD	CD	-CD
Nonlast MFS page (autopage)		S		X				
Last MFS page (autopage)								
Recoverable and nonrecoverable					S		X	S
Fast Path conversational recoverable and nonrecoverable transaction								
Response mode transaction	S		X	X	S		X	X
Nonrecoverable, nonresponse, nonconversational mode transaction	S	S ²	X	X	S	Note 2 on page 531	X	X
Recoverable, nonresponse, nonconversational mode transaction						S	X	S
MFS paging control request: SNA-formatted QMODEL FMHs	S							
MFS paging control request: SNA RAP FMH					S		X	
IMS message switch								
IMS message switch using ATTACH SCHEDULER					S		X	S
IMS message switch using ATTACH (no SCHEDULER)								Note 3 on page 531
ATTACH SYSMMSG ⁴								Note 3 on page 531
IMS command	S		X	X	S		X	X
VTAM command ⁵								
LUSTATUS - commit					S		X	S
Other	X		X	X				
Test mode ⁶	S		X	X	S		S	X
FMH7 (ERP) messages	X	Note 3 on page 531	X	Note 3 on page 531				

Notes:

1. A response is implied by a reply to the change-direction indicator; therefore, RQE1 or RQE2 is recommended when change-direction is sent. Either method of requesting a response is supported for the session BIND option of DEFINITE RESPONSE CHAINS. When IMS is running as a secondary system, the other half session (as the primary system), if sending a BB-CD chain, must indicate RQD1 to allow proper DFC bracket and send/receive synchronization.
2. Supported only for irrecoverable-inquiry input requesting EB or BB/EB.
3. Sent with EB. BB is also sent as needed.
4. Because IMS SYSMMSG is handled in the same manner as is a message switch, the same protocols apply.
5. LUSTATUS might indicate RQE1 with EB or BB/EB. CHASE, LUSTATUS, and CANCEL commands can optionally request RQE1 with CD. LUSTATUS - commit can also request RQE2 or RQD2 for specific conditions. All other normal flow VTAM commands must request RQD1.
6. Applies to test "echo" mode only. Not applicable to /TEST MFS.

Related reference:

"Bracket and half-duplex protocol" on page 533

Requests on IMS output messages

The response and sync-point requests for IMS output messages are summarized in the following table.

An "X" in the table indicates that the corresponding entry is supported by IMS.

Table 77. Response and sync-point requests for IMS output messages.

Output message types	VTAM indicators with message							
	RQE1 ¹		RQD1		RQE2 ¹		RQD2	
	CD	¬CD	CD	¬CD	CD	¬CD	CD	¬CD
Update, recoverable					X			X
Inquiry, recoverable					X			X
Inquiry, nonrecoverable	X			X				
Fast Path (recoverable)								X
Nonlast MFS page (autopaged)		See notes ²		See notes ²				
First chain demand-paged output								
ATTACH (no SCHEDULER)	X							
ATTACH SCHEDULER								X
Nonlast MFS page (demand-page)	X							

Table 77. Response and sync-point requests for IMS output messages (continued).

Output message types	VTAM indicators with message							
	RQE1 ¹		RQD1		RQE2 ¹		RQD2	
	CD	-CD	CD	-CD	CD	-CD	CD	-CD
Last MFS page (autopaged or without OLP) recoverable, nonrecoverable					X			X
Last MFS page (with OLP) ³	X							
IMS command reply: /DISPLAY, /RDISPLAY, /FORMAT					X			X
IMS command reply: /TEST	X							
Other IMS command replies				X				
Test mode output ⁴	X			X				
Broadcast output (ATTACH SYMSG)					X			X
System error messages (ATTACH SYMSG)	X			X	X			X
Message switch output					X			X
VTAM commands ⁵	X			X				
FMH7 (ERP) messages	X			See notes ⁶				
QSTATUS (QMODEL) ⁷	X				X			X

Notes:

1. A response is implied by a reply to the change-direction indicator; therefore, IMS indicates RQE1 or RQE2 when sending the change-direction indicator. Either method of requesting a response is supported for the session BIND option DEFINITE RESPONSE CHAINS.
2. RQD1 occurs on each chain rather than RQE1 when IMS is bound as "RQD only," and on the first page of MFS autopaged output to prevent unnecessary ERP overhead if errors or contention is detected by the receiver.
3. MFS operator logical paging is in effect if the MOD specifies PAGE=YES and autopaged output is not indicated in the MFS system control area.
4. Applies to test "echo" mode only. Not applicable to /TEST MFS.
5. CHASE is always sent RQE1/CD.
6. Sent with EB.
7. CD/RQE1 is sent on QSTATUS, which results from an invalid cursor in an MFS DPM demand-page request for output sent as ATTACH SCHEDULER. CD/RQE2 or EB/RQD2 is sent on QSTATUS, which results from the receipt of QPURGE during demand-paged output. CD/RQE2 is sent:

- When operating in IMS conversation or in test mode
- When the input is sent with ATTACH SCHEDULER and the associated output component is SINGLE2 or MULT2.

All other cases result in EB/RQD2.

Related reference:

“LUSTATUS protocol” on page 546

“CANCEL protocol” on page 541

“Bracket and half-duplex protocol”

Data flow control protocol reference

The following topics describe the byte-level protocols for data flow control (DFC). The protocols are presented here.

BID protocol

IMS does not send the SNA BID command, but receives the command when acting as a secondary half session.

When receiving the BID command in a between-bracket state, IMS responds DR1 and enters a receive state for the pending input. If the BID command is received while IMS is in-brackets, IMS rejects the BID with an exception DR1 response (X'08130000').

Bracket and half-duplex protocol

IMS uses SNA bracket protocol to resolve contention, and uses the change-direction indicator of the half-duplex protocol to control normal flow send/receive mode while within bracket state.

IMS uses these options when in session with another ISC logical unit:

- Bracket reset state can be selected at data traffic active state. This state can be set to between-brackets (BETB) or in-brackets (INB).
- Half-duplex send or receive state can be selected at data traffic active state.
- Either half session can send end-bracket (EB). IMS must always be bound to allow EB.
- Half-duplex flip-flop is used for normal flow send/receive states.
- Bracket termination rule 1 (conditional rule) is used.

Related reading:

- For information on bracket termination rules, see *z/OS Communications Server: SNA Programming*.

Related reference:

 Bind parameters for SLU P and LU 6.1 (System Programming APIs)

Bracket protocol for IMS input

One or more input messages can be received within a bracket.

For example:

- Any number of nonresponse, nonconversation, noncommand, or nontest mode input messages can be sent within the bracket. The input, the subsequent

asynchronous output, or an LUSTATUS can end the bracket. Further, a single nonresponse, nonconversation, noncommand, or nontest mode input message with ATTACH EB can be sent.

- One response mode or command message, except /TEST, can optionally occur within the same bracket after any of the input messages listed in the previous list bullet. The response mode or command message output ends the bracket.
- Any number of conversation or test mode inputs followed by conversation or test mode outputs can optionally begin within the same bracket after the input messages listed in the first list bullet. The bracket ends on the last conversation output message (or command complete message from /END for test mode). The bracket can also be ended by the other half session's sending an input LUSTATUS or CHASE with EB, but not by input data sent with EB.
- The bracket can be ended by an EB occurring on an FMH7 or LUSTATUS resulting from an error.

The following series of tables summarizes the bracket and send/receive indicators acceptable for the various message types input to IMS.

An “X” in the table indicates that the corresponding entry is supported by IMS.

In the tables the VTAM indicators are represented by the following abbreviations:

- BB: Begin-Bracket
- EB: End-Bracket
- CD: Change-Direction. Allowed on Last- or Only-in-Chain.
- BB/EB: Begin- and End-Bracket. Allowed on First- or Only-in-Chain.

Table 78. VTAM bracket and send/receive indicators that are sent with messages to IMS: input message type using ATTACH SCHEDULER

Input message type using ATTACH SCHEDULER	BB	EB	BB/EB	BB/CD	CD	¬BB, ¬EB, ¬CD
MFS-autopaged input: First page (of multiple page input)	X					X
MFS-autopaged input: Nonfirst, nonlast page						X
MFS-autopaged input: Last page: see transaction types						
Nonresponse transaction						
Nonconversational mode transaction	Note 1 on page 535	Note 2 on page 535	Notes1 on page 535, 2 on page 535	Note 1 on page 535	X	X
Response mode (including Fast Path) transaction	Note 1 on page 535			Notes 1 on page 535, 3 on page 535	Note 3 on page 535	X
Conversational transaction	Note 1 on page 535			Notes 1 on page 535, 3 on page 535	Note 3 on page 535	X
IMS message switch	X	X	X	X	X	X
IMS command	X	Note 4 on page 535	Note 4 on page 535	Note 3 on page 535	Note 3 on page 535	X
Input while in test mode ⁴					Note 3 on page 535	X

Notes:

1. Not valid for the last page of MFS-autopaged input.
2. Not valid for the first page of MFS-autopaged input.
3. Denotes the optimal end-of-message indicators to prevent IMS from soliciting change-direction using SNA SIGNAL.
4. Valid only for /DIS, /RDIS, and /FOR commands. Also optimal for these commands to keep IMS from immediately forcing a between-brackets state by sending LUSTATUS - NO-OP (X'0006') with EB. The output from these commands is always queued and sent asynchronously.
5. FMH7 messages and LUSTATUS NO-OP indicating EB force end of conversation, response mode, and test mode and cause the associated IMS output message to be dequeued. Further, the Conversational Abnormal Termination exit routine is invoked.

Table 79. VTAM bracket and send/receive indicators that are sent with input messages to IMS using ATTACH (no SCHEDULER)

Input message type using ATTACH (no SCHEDULER)	BB	EB	BB/EB	BB/CD	CD	BB, EB, CD
MFS-autopaged input: First page (of multiple page input)	X					X
MFS-autopaged input: Nonfirst, nonlast page						X
MFS-autopaged input: Last page: see transaction types						
Nonresponse, nonconversational mode transaction		X	X			
Response mode (including Fast Path) transaction ⁶	Note 1			Notes 1, 3	Note 3	X
Conversational transaction	Note 1			Notes1, 3	Note 3	X
IMS message switch		X	X			
ATTACH SYMSMSG		X	X			
IMS command	X	Note 4	Note 4	Note 3	Note 3	X
Input while in test mode ⁵					Note 3	X

Notes:

1. Not valid for the last page of MFS-autopaged input.
2. Not valid for the first page of MFS-autopaged input.
3. Denotes the optimal end-of-message indicators to prevent IMS from soliciting change-direction using SNA SIGNAL.
4. Valid only for /DIS, /RDIS, and /FOR commands. Also optimal for these commands to keep IMS from immediately forcing a between-brackets state by sending LUSTATUS - NO-OP (X'0006') with EB. The output from these commands is always queued and sent asynchronously.
5. Applies to test "echo" mode only. Not applicable to /TEST MFS.
6. For ISC, response mode operation is forced for attached transactions not indicating EB (regardless of their definition on the TRANSACT macro statement) when the TERMINAL macro statement or ETO user descriptor

indicates either TRANRESP or FORCRESP. The input is rejected if the TERMINAL macro statement or ETO user descriptor indicates NORESP.

Table 80. Vtam bracket and send/receive indicators sent with other input messages sent to IMS

Other input message types	BB	EB	BB/EB	BB/CD	CD	BB, EB, CD
VTAM normal flow command/ indicator ^{1, 2}	X				X	X
FMH7 (ERP) messages		Note 5 on page 535	X	X	X	
IMS MFS paging control request (QMODEL FMHs)				X	X	
RAP FMH					X	

Notes:

1. LUSTATUS, CHASE, and CANCEL can indicate EB or CD. Other VTAM normal flow commands or indicators must not indicate either EB or CD.
When IMS is waiting for conversational input, LUSTATUS and CHASE cannot be sent to IMS unless they also indicate end-bracket. EB causes the conversational output message to be dequeued, the conversation mode to be terminated, and the Conversational Abnormal Termination exit routine to be invoked. The session is terminated if an LUSTATUS or CHASE is received by IMS with any other protocol.
2. FMH7 messages and LUSTATUS NO-OP indicating EB force end of conversation, response mode, and test mode and cause the associated IMS output message to be dequeued. Further, the Conversational Abnormal Termination exit routine is invoked.

Related concepts:

“Handling IMS response mode or conversational output errors” on page 523

Related reference:

“Data flow control protocol reference” on page 533

Bracket protocol for IMS output

The output bracket and send/receive protocol used by IMS and the number of output messages sent per bracket are dependent on a variety of factors.

The output bracket and send/receive protocol used by IMS and the number of output messages sent per bracket depend upon one or more of the following:

- The IMS message type
- Whether the message is to be MFS paged
- Whether the other half session supports the SCHEDULER model
- The output component specified during IMS system definition
- The bracket and ATTACH protocol associated with the originating input transaction

Within IMS, the protocols used for nonlast chains (pages) of MFS-paged output and the last (MFS-paged) or only chain of response mode, conversation mode, test mode, or command replies are predefined regardless of whether the originating input transaction and resulting output replies are synchronous (ATTACH) or asynchronous (SCHEDULER). This is also true for output resulting from input sent with ATTACH EB and for asynchronous output that must be sent ATTACH because the other half session lacks SCHEDULER support.

CD is used for all nonlast chains of MFS demand-paged messages, the last (MFS-paged) or only chain of test mode, and nonlast conversational mode messages. CD allows the paging operation or synchronous event between IMS and the other half session to continue.

EB ensures that the end of the synchronous event occurs by placing the session in reset state for the last (MFS-paged) or only chain of response mode replies, last conversational mode replies, and command replies. This is also true for replies resulting from input sent with ATTACH EB and for asynchronous output that must be sent ATTACH because the other half session lacks SCHEDULER support. All nonlast chains (pages) of MFS-autopaged messages are sent without either CD or EB.

Using ATTACH SCHEDULER, your installation must define to IMS the protocols to be used for the last (MFS-paged) or only chain of other asynchronous output. IMS allows you to specify single or multiple messages for each component defined for a session and whether IMS should send more than one asynchronous message before indicating change-direction or end-bracket. The possible component definitions are:

SINGLE1

Asynchronous output for this component is sent one message per bracket. Each message begins a bracket (if necessary) and always ends a bracket.

SINGLE2

Asynchronous output for this component is sent one message at a time with the VTAM begin-bracket (if necessary) and change-direction indicators to allow the receiving subsystem to optionally send its communication traffic.

MULT1

All asynchronous messages for a given LTERM are sent before the bracket is ended. Traffic is sent BB (if necessary), message1, message2,...messageN, EB. EB occurs after the last message for the LTERM is acknowledged and dequeued.

MULT2

All asynchronous messages for a given LTERM are sent before change-direction is sent. Traffic is sent BB (if necessary), message1, message2...messageN, CD. CD occurs after the last message for the LTERM is acknowledged and dequeued.

Consider the definition (on the IMS system definition TERMINAL macro statement or an ETO logon descriptor), and use of protocols for ISC components to be used for transmission of asynchronous messages. This is particularly true in an IMS-to-IMS environment, because most messages between the two subsystems are asynchronous. Incorrectly defining or using these ISC protocols can:

- Require extra transmissions to occur in order to acquire the flow or to end the bracket
- Cause unnecessary bracket contention error recovery operations
- Produce output protocols unacceptable to the receiving subsystem.

A further explanation follows:

SINGLE1

Use of SINGLE1 results in EB on each message and might cause bracket contention if the other subsystem has data to send.

EB on each message can cause process errors within the other subsystem if the transaction must reply synchronously within the initiating bracket. An error of

this type can occur in the receiving subsystem when a SINGLE1 component is used to send response mode, conversation, or test mode messages, or IMS commands from one IMS subsystem to another.

SINGLE2 or MULT2

Change-direction sent on messages from a component defined as SINGLE2 might result in unnecessary transmissions to end the bracket when either no output is available or no reply is available within the other subsystem.

MULT1 or MULT2

Messages sent on components defined as MULT1 OR MULT2 indicate change-direction or end-bracket only after the last message on a queue has been dequeued. Therefore, extra transmissions might result if the other subsystem must signal for the flow to return synchronous replies.

For components defined as MULT1 and MULT2, IMS suppresses the queue rotation that normally occurs between output messages. This suppression allows all messages from a queue selected for output to be sent before IMS initiates output on other queues. Selection of the next queue with available output occurs when the previous queue has emptied or when input changes the active output queue. The queues are not rotated if an LUSTATUS - queue empty or another input that indicates change-direction occurs for the same queue.

The following tables summarize the bracket and send/receive indicators that are acceptable for IMS output messages.

In the tables the VTAM indicators are represented by the following abbreviations:

- INB** Indicates that synchronous or asynchronous output occurs within the same bracket as the previous input
- BETB** Indicates that IMS initiates output while in a between-brackets state
- BB** Begin-bracket
- EB** End-bracket
- CD** Change-direction (allowed on last- or only-in-chain)
- N** No bracket or send/receive indicators

Table 81. VTAM bracket and send/receive indicators sent with output message using ATTACH SCHEDULER

Message type sent using attach scheduler	SINGLE1 msg.init INB BETB	SINGLE2 msg.init INB BETB	MULT1/2 ¹ msg.init INB BETB
MFS demand-paged output:First SNA chain (ATTACH SCHEDULER) ²	EBBB/EB	EBBB/EB	EBBB/EB
MFS demand-paged output:Nonlast page			
MFS demand-paged output:Last page if OLP ³	CD	CD	CD
MFS demand-paged output:Last page without OLP (See appropriate message type)			
MFS autopaged output: First page	NBB	NBB	NBB
MFS autopaged output: Nonfirst, nonlast page	N	N	N
MFS autopaged output: Last page (See appropriate message type)			
Response mode output (Including Fast Path)	EB	EB	EB
Nonlast conversational output message	CD	CD	CD
Last conversational output message	EB	EB	EB

Table 81. VTAM bracket and send/receive indicators sent with output message using ATTACH SCHEDULER (continued)

Message type sent using attach scheduler	SINGLE1 msg.init INB BETB	SINGLE2 msg.init INB BETB	MULT1/2 ¹ msg.init INB BETB
Message switch output	EBBB/EB	CDBB/CD	NBB
Command output: /FOR, /DIS, /RDIS	EBBB/EB	CDBB/CD	NBB
Command output: /TEST ⁴	CD	CD	CD
Other Command output	EB	EB	EB
Test mode output ⁴	CD	CD	CD
None of the above message types - asynchronous	EBBB/EB	CDBB/CD	NBB

Notes:

1. LUSTATUS - queue empty is used at the end of a queue to send EB for MULT1 and CD for MULT2. EB is sent when an ERP backout has reset the DFC and ATTACH states to between-brackets and no component 1 (COMPT1) is defined during IMS system definition. Also, the BB indicator in the BETB column occurs only for the first message or MFS page that must initiate an output bracket asynchronously.
2. Although IMS is between-brackets after having sent stand-alone ATTACH SCHEDULER for response mode, Fast Path, or conversational demand-paged output, IMS does not accept input until the preceding output is successfully transmitted and dequeued using the appropriate paging requests.
3. MFS operator logical paging (OLP) is in effect if the MOD specifies PAGE=YES and autopaged output is not indicated in the MFS system control area (SCA).
4. Applies to test "echo" mode only. Not applicable to /TEST MFS.

Table 82. VTAM bracket and send/receive indicators sent with output message using ATTACH (no SCHEDULER)

Message type sent using ATTACH (no SCHEDULER)	SINGLE1 msg.init INB BETB	SINGLE2 msg.init INB BETB	MULT1/2 ¹ msg.init INB BETB
MFS demand-paged output: First SNA chain (ATTACH)	CD	CD	CD
MFS demand-paged output: Nonlast page, last page if OLP ³	CD	CD	CD
MFS demand-paged output: Last page without OLP (See appropriate message type)			
MFS autopaged output: First page	N	N	N
MFS autopaged output: Nonfirst, nonlast page	N	N	N
MFS autopaged output: Last page (See appropriate message type)			
Response mode output (including Fast Path) ⁵	EB	EB	EB
Nonlast conversational message	CD	CD	CD
Last conversational message	EB	EB	EB
Nonresponse, nonconversational output (No SCHEDULER model defined for other half session)	EB BB/EB	EB BB/EB	EB BB/EB
ATTACH SYMSG	BB/EB	BB/CD	BB
Command output			
Command output: /TEST ⁴	CD	CD	CD

Table 82. VTAM bracket and send/receive indicators sent with output message using ATTACH (no SCHEDULER) (continued)

Message type sent using ATTACH (no SCHEDULER)	SINGLE1 msg.init INB BETB	SINGLE2 msg.init INB BETB	MULT1/2 ¹ msg.init INB BETB
Other command output (except /FOR, /DIS, and /RDIS) ⁶	EB	EB	EB
Test mode output ⁴	CD	CD	CD

Notes:

1. LUSTATUS - queue empty is used at the end of a queue to send EB for MULT1 and CD for MULT2. EB is sent when an ERP backout has reset the DFC and ATTACH states to between-brackets and no component 1 (COMPT1) is defined during IMS system definition. Also, the BB indicator in the BETB column occurs only for the first message or MFS page that must initiate an output bracket asynchronously.
2. Although IMS is between-brackets after having sent stand-alone ATTACH SCHEDULER for response mode, Fast Path, or conversational demand-paged output, IMS does not accept input until the preceding output is successfully transmitted and dequeued using the appropriate paging requests.
3. MFS operator logical paging (OLP) is in effect if the MOD specifies PAGE=YES and autopaged output is not indicated in the MFS system control area (SCA).
4. Applies to test "echo" mode only. Not applicable to /TEST MFS.
5. Response mode operation always occurs for input transactions that are directly attached.
6. See "LUSTATUS protocol" on page 546 for more information on /DISPLAY, /RDISPLAY, and /FORMAT command output.

Table 83. VTAM bracket and send/receive indicators sent with other output message types

Other output message types	SINGLE1 msg.init INB BETB	SINGLE2 msg.init INB BETB	MULT1/2 ¹ msg.init INB BETB
VTAM command/indicator ²			
FMH7 (ERP) messages: During conversation	CD	CD	CD
FMH7 (ERP) messages: During test mode	CD	CD	CD
FMH7 (ERP) messages: Other ³	EB BB/EB	CD BB/CD	Note 4
QSTATUS (QMODEL)	Note 5 on page 541	Note 5 on page 541	Note 5 on page 541

Notes:

1. LUSTATUS - queue empty is used at the end of a queue to send EB for MULT1 and CD for MULT2. EB is sent when an ERP backout has reset the DFC and ATTACH states to between-brackets and no component 1 (COMPT1) is defined during IMS system definition. Also, the BB indicator in the BETB column occurs only for the first message or MFS page that must initiate an output bracket asynchronously.
2. BIS is sent without CD or EB. CHASE is always sent CD.
3. An FMH7 requests the protocol associated with the resulting between-brackets reset state of the last committed input component after backout of the DFC and ATTACH states due to the ERP operation.
4. FMH7 is sent with EB for MULT1, CD for MULT2.

5. CD/RQE1 is sent on a QSTATUS that results from an invalid cursor in an MFS DPM demand-page request for output. CD/RQE2 or EB/RQD2 is sent on QSTATUS, which results from the receipt of QPURGE during demand-paged output. CD/RQE2 is sent: 1) when in IMS conversation or test mode or 2) when the input is sent with ATTACH SCHEDULER and the associated output component was SINGLE2 or MULT2. All other cases result in EB/RQD2.

Related reference:

Chapter 36, "ISC data flow control examples," on page 653

"LUSTATUS protocol" on page 546

"CANCEL protocol"

CANCEL protocol

IMS sends and receives the SNA CANCEL command, which allows the sender to terminate an output SNA chain without having to send all of the chain following a sender- or receiver-detected error.

The CANCEL command that is sent because of either a sender- or receiver-detected error is sent by IMS with the VTAM indicators shown in the following table.

The CANCEL command that is received by IMS because of either a sender- or receiver-detected error can indicate whatever protocol (consistent with current DFC states) is appropriate for the message sender.

In the table the VTAM indicators that are sent with the CANCEL command are represented by the following abbreviations:

- EB: End-Bracket
- CD: Change-Direction

Table 84. VTAM indicators sent with the CANCEL command

Output (determined in following order)	¬CD or ¬EB	CD	EB
After /DEQ (LUSTATUS ABORT is next)	X		
When primary half session and FIC/OIC was between bracket (conditional BB)	X		
After receiving selective receiver ERP sense code		X	
After receiving SIGNAL		X	
	X		
When conversational or response mode output is still available			
FIC/OIC was EB RQD* (conditional EB)			X
While in test (echo) mode		X	
Non-MFS demand-paged output for SINGLE1 component	See indicators for FIC/OIC was between bracket or FIC/OIC was EB		
MFS demand-paged output for SINGLE1 component			X
Output for SINGLE2 component		X	
Output for MULT1 or MULT2 component	X		

Chaining protocol

Whether operating as the primary or secondary half session, IMS sends both single and multiple RU chains. Half sessions can operate with either single or multiple RU chains as specified in the bind parameters.

IMS messages are usually sent and received as single SNA chains where only-in-chain (OIC) or first-in-chain (FIC) indicates beginning-of-message, and OIC or last-in-chain (LIC) indicates end-of-message. However, each page of an MFS-paged input or output message is sent or received as a single SNA chain. The beginning-of-message occurs when the first page is sent or received, and the end-of-message occurs at sync point (RQD2 or RQE2, CD) requested at the end of the last input or output page (except for MFS operator logical paging).

Optionally, end-of-message can be requested using LUSTATUS - commit immediately following the last page of MFS DPM-autopaged input.

Use the CANCEL command when errors are detected on multiple RU input chains.

Related reference:

“LUSTATUS protocol” on page 546

“ERP purging”

CHASE protocol

IMS sends and receives SNA CHASE.

IMS sends CHASE requesting RQE1/CD to cause a synchronizing event after the receipt of an exception response when it is between pages (chains) of MFS-autopaged output and is not bound with Definite Response Chains Only.

IMS responds with DR1, as necessary, when receiving CHASE. IMS can receive CHASE between any two messages within a bracket. IMS can also receive CHASE as a synchronizing event after IMS has sent an exception response to MFS-autopaged input. (The sender of the MFS-autopaged input receives the exception response while between output pages or after a chain indicating RQE1.)

Only CHASE carrying EB is accepted while IMS is waiting for conversational input. EB causes the conversational output message to be dequeued, conversation mode to be terminated, and the Conversational Abnormal Termination exit routine to be invoked, just as if an /EXIT command had been received. The session is terminated if CHASE is received with any protocol other than EB.

ERP purging

After sending an exception response and before continuing to send or receive, the exception response sender might need to enter error recovery process (ERP) PURGE mode until the DFC state managers of both half sessions are synchronized.

Additionally, ERP PURGE must be complete prior to sending the FMH7 ERP message. ERP purging can occur for either single or multiple SNA chains. Single chain purge occurs when DFC states are synchronized within the same chain receiving the exception response. Multichain purge occurs when more than a single SNA chain must be purged before the DFC state managers are synchronized. (The CANCEL and CHASE commands are logically considered part of the same chain when they occur immediately after the SNA chain that resulted in the exception response.)

Single- and multiple-chain purging occur within the data flow control support layer and are independent of SNA presentation layer functions such as MFS.

ERP PURGE must be entered when a half session sends an exception response to an RU (OIC, MIC, or LIC) that does not result in ending a bracket or in both half-session DFC states being synchronized. That is, ERP PURGE begins when a half session sends an exception response to the following RUs:

- FIC RQE* (not applicable for EB)
- MIC RQE* (not applicable for CD or EB)
- LIC RQE* ¬EB and ¬CD
- OIC RQE* ¬EB and ¬CD

ERP PURGE ends when the purging half session receives an RU that either ends the bracket or causes both half-session DFC states to again be synchronized. That is, ERP PURGE terminates when the purging half session receives the following RUs:

OIC RQD*

This includes the CANCEL and CHASE commands, and is not applicable for CD or EB.

OIC RQE* CD or EB

This includes the CANCEL command.

LIC RQD*

Not applicable for CD or EB.

LIC RQE* CD or EB

ERP PURGE can span several SNA chains before terminating. When more than one chain is purged, an additional exception response is necessary at the point of bracket termination or DFC state synchronization (except for CANCEL or CHASE commands) to prevent an ambiguous or incorrect understanding as to the disposition of the message by the sending half session (response receiver). This sender ERP exception response (X'0867') indicates that the chain was purged because of an error on a previous chain and that the message sender should resend the message at the next possible opportunity. The X'0867' exception response must be sent when a half session ends ERP PURGE after receiving the following RUs in a chain subsequent to the one that resulted in the original exception response:

OIC RQD*

This does not include CANCEL or CHASE, and is not applicable for CD or EB.

OIC RQE* CD or EB

Except for CANCEL

LIC RQD*

Not applicable for CD or EB.

LIC RQE* CD or EB

The DFC protocol that can be received with the FMH7 ERP message must be either change-direction or end-bracket and can include begin-bracket, as appropriate. The DFC protocol (CD or EB) that is sent by IMS with the FMH7 ERP message is determined by the ATTDSP value resulting after ERP backout. When backout results in a between-brackets state, the ATTDSP value is component 1. When backout results in a state other than between-brackets, the ATTDSP value is the last

committed input component. The protocols are set for the resulting component. End-bracket is sent when an ERP backout has reset the DFC and ATTACH states to between-brackets and no component 1 (COMPT1) was defined during IMS system definition.

Related reference:

“Bracket protocol for IMS output” on page 536

Resulting DFC state after sender ERP purge

The following table reflects all valid bracket and send/receive states that result after both half sessions reach a sync point after an exception response that indicates a sender ERP sense code other than selective receiver ERP.

In the table, the following abbreviations are used:

BETB Between-brackets

INB In-brackets

PHS Primary half session

SHS Secondary half session

Table 85. Resulting DFC states after sender ERP purge

Type of chain in error	PHS sends data, SHS sends exception: half-session states		SHS sends data, PHS sends exception: half-session states	
	PHS	SHS	PHS	SHS
BB/— FIC				
RQ** CD OIC,LIC	1	1	INB.SEND	INB.RCV
RQD* OIC,LIC	BETB	BETB	INB.RCV	INB.SEND
RQD* CANCEL	BETB	BETB	INB.RCV	INB.SEND
RQ** CD CANCEL	1	1	INB.SEND	INB.RCV
RQD* EB CANCEL	1	1	BETB	BETB
BB/EB FIC				
RQE* OIC,LIC	BETB	BETB	BETB	BETB
RQD* OIC,LIC	BETB	BETB	BETB	BETB
RQD* CANCEL	BETB	BETB	BETB	BETB
RQ** CD CANCEL ¹				
RQD* EB CANCEL ¹				
--/-- FIC				
RQE* OIC,LIC	INB.SEND	INB.RCV	INB.RCV	INB.SEND
RQ** CD OIC,LIC	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* OIC,LIC	INB.SEND	INB.RCV	INB.RCV	INB.SEND
RQD* CANCEL	INB.SEND	INB.RCV	INB.RCV	INB.SEND
RQ** CD CANCEL	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* EB CANCEL	BETB	BETB	BETB	BETB
--/EB FIC				
RQ** CD OIC,LIC	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* OIC,LIC	INB.SEND	INB.RCV	INB.RCV	INB.SEND

Table 85. Resulting DFC states after sender ERP purge (continued)

Type of chain in error	PHS sends data, SHS sends exception: half-session states		SHS sends data, PHS sends exception: half-session states	
	PHS	SHS	PHS	SHS
RQD* CANCEL	INB.SEND	INB.RCV	INB.RCV	INB.SEND
RQ** CD CANCEL	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* EB CANCEL	BETB	BETB	BETB	BETB

1. CD and EB might not be sent while in DFC BETB state.

Resulting DFC state after selective receiver ERP purge

The following table reflects all valid bracket and send/receive states that result after both half sessions reach a sync point after an exception response that indicates the SNA selective receiver ERP sense code, and before the FMH7 is sent by the half session detecting the error.

Table 86. Resulting DFC states after selective receiver ERP purge

Type of chain in error	PHS sends data, SHS sends exception: half-session states		SHS sends data, PHS sends exception: half-session states	
	PHS	SHS	PHS	SHS
BB/-- FIC				
RQ** CD OIC,LIC	1	1	INB.SEND	INB.RCV
RQD* OIC,LIC	ERP.RCV	ERP.SEND	INB.SEND	INB.RCV
RQD* CANCEL	ERP.RCV	ERP.SEND	INB.SEND	INB.RCV
RQ** CD CANCEL	1	1	INB.SEND	INB.RCV
RQD* EB CANCEL	1	1	ERP.SEND	ERP.RCV
BB/EB FIC				
RQE* OIC,LIC	ERP.RCV	ERP.SEND	ERP.SEND	ERP.SEND
RQD* OIC,LIC	ERP.RCV	ERP.SEND	ERP.SEND	ERP.SEND
RQD* CANCEL	ERP.RCV	ERP.SEND	ERP.SEND	ERP.SEND
RQ** CD CANCEL				
RQD* EB CANCEL				
--/-- FIC				
RQE* OIC,LIC	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQ** CD OIC,LIC	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* OIC,LIC	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* CANCEL	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQ** CD CANCEL	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* EB CANCEL	ERP.RCV	ERP.SEND	ERP.SEND	ERP.RCV
--/EB FIC				
RQ** CD OIC,LIC	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* OIC,LIC	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQD* CANCEL	INB.RCV	INB.SEND	INB.SEND	INB.RCV
RQ** CD CANCEL	INB.RCV	INB.SEND	INB.SEND	INB.RCV

Table 86. Resulting DFC states after selective receiver ERP purge (continued)

Type of chain in error	PHS sends data, SHS sends exception: half-session states		SHS sends data, PHS sends exception: half-session states	
	PHS	SHS	PHS	SHS
RQD* EB CANCEL	ERP.RCV	ERP.SEND	ERP.SEND	ERP.RCV

Note:

1. CD and EB cannot be sent while in between-brackets state

LUSTATUS protocol

IMS sends and receives LUSTATUS as summarized in the following table. The DFC bracket, send/receive, and response requirements illustrated in this figure are subject to the considerations detailed in the balance of this topic.

Any LUSTATUS sense code not listed in the following table causes the ISC session to be terminated. An S in the table indicates the suggested indicator settings for the given LUSTATUS. An X in the table indicates that IMS also supports the specified indicator settings for the given LUSTATUS.

Table 87. VTAM indicators sent with LUSTATUS

LUSTATUS	VTAM indicators with LUSTATUS								
	RQE1		RQD1			RQE2	RQD2		
	EB	CD	--	CD	EB	CD	--	CD	EB
Sense Code Received ¹ : Commit-X'0006'						S	S	X	S
Sense Code Received ¹ : NO-OP-X'0006'	X	S		X	S				
Sense Code Received ¹ : Queue Empty-X'0007'		S		X	S				
Function Abort-X'0864'		S	S	X	S				
Function Abort-X'0865'		S	S	X	S				
Function Abort-X'0866'		X	X	X	X				
Sense Codes Sent: Commit-X'0006'									
Sense Codes Sent: NO-OP-X'0006'		X			X				
Sense Codes Sent: Queue Empty-X'0007'		X			X				
Function Abort -X'0865'		X	X		X				

Notes:

1. While IMS is waiting for conversational input, it accepts LUSTATUS carrying EB only. EB causes the conversational output message to be dequeued,

conversation mode to be terminated, and the Conversational Abnormal Termination exit routine to be scheduled. The session is terminated if an LUSTATUS carrying any other protocol is received.

IMS responds DR1 or DR2 when receiving LUSTATUS RQD1 or RQD2. IMS requests DR1, DR2, or exception DR1 or DR2 as indicated in the previous table. The session is terminated if an exception occurs to LUSTATUS.

- LUSTATUS - Queue Empty.

IMS sends and receives LUSTATUS - queue empty.

IMS sends LUSTATUS - queue empty:

- After receiving an asynchronous (ATTACH SCHEDULER) input message and having no output immediately available.
- After sending all available output on a given queue.

RQD1 and end-bracket are indicated on the LUSTATUS sent when IMS is not in conversation or response mode and has been left in-brackets/SEND, and no output is available to be sent from a component defined as SINGLE1 or MULT1. RQE1 and change-direction are indicated on the LUSTATUS sent when IMS is not in conversation or response mode and has been left in-brackets/SEND, and no output is available to be sent from a component defined as SINGLE2 or MULT2.

- LUSTATUS - Function Abort. The supported sense codes are:

X'0864'

Loop occurs upon re-execution. Sender should not resend the same data.

X'0865'

Data sender is responsible for detecting and preventing loop.

X'0866'

Data receiver is responsible for detecting and preventing loop.

IMS sends LUSTATUS - function abort X'0865' only under these conditions:

- When an IMS /DEQUEUE command has been issued by an authorized IMS terminal operator before the last page of an MFS-demand or autopaged output message has been sent by IMS.

The DFC bracket and send/receive protocol indicated on the LUSTATUS is the same as that which would occur with the last chain (page) had the message been successfully sent.

The following occurs when IMS receives LUSTATUS - function abort:

- Except during conversational mode, any function abort received after an output message (including the last page of MFS demand-paged or autopaged output) sent RQE/CD causes the output message to be committed and does not cause session termination. LUSTATUS that is processed as a normal flow input before it is processed as an LUSTATUS command also causes the message to be committed.
- Any function abort causes an incomplete input MFS-autopaged message to be discarded.
- Function abort X'0864', received before the last page of an output MFS demand-paged message is sent, causes the message to be dequeued prior to continuing normal input/output operations if the output is conversational and demand-paged. The Conversational Abnormal Termination exit routine is invoked, just as if an /EXIT command had been received on the session.
- Function abort X'0865', received before the last page of an active output MFS demand-paged message is sent, causes the message to be returned to the output message queue and the session to be terminated.

- Function abort X'0866', received before the last page of an active output MFS demand-paged message, causes the message to be returned to the message queue and made available for retransmission before normal input/output operations continue.
- Any function abort resets the ATTACH states to those in effect at the last sync point.

LUSTATUS - function abort might never be sent to IMS in reply to any RU indicating RQD2 or RQE2, or the session is terminated.

- LUSTATUS - Commit

IMS does not send LUSTATUS - commit. IMS receives LUSTATUS - commit as an end-of-message indication immediately following the last page of autopaged input and with EB while awaiting conversational input. This creates a "normal" end of the conversation by invoking the Conversational Abnormal Termination exit routine with a new input vector of X'28'. This allows the exit routine to schedule a transaction to commit pending resources reflected in the scratchpad area (SPA).

- LUSTATUS - NO-OP

IMS sends and receives LUSTATUS - NO-OP.

IMS sends LUSTATUS - NO-OP:

- To allow RQE1 and change-direction (CD) to flow after SIGNAL RCD.
- To end the bracket after receiving a synchronous input message that generates no output. In IMS, this only occurs for the FORMAT, /RDISPLAY, and /DISPLAY commands when attached synchronously. The LUSTATUS - NO-OP is sent indicating RQD1 and EB.
- To end the bracket at those times when bound in-brackets/SEND and no synchronous output or restart is possible. That is, IMS is not in conversation or response mode after session restart. The LUSTATUS is sent, indicating RQD1 and EB.
- To allow required input when bound in-brackets/SEND at session restart and a pending conversational message is dequeued using STSN protocol. LUSTATUS is sent indicating RQE1 and CD.
- To end the bracket after receiving exception or ERP FMH7 (without EB) sense code X'0864', following response mode or the last output message of an IMS conversation.

LUSTATUS - NO-OP can be received during asynchronous or synchronous input processing and might indicate either CD or EB, as necessary. Receipt of LUSTATUS EB during an IMS conversation or response mode or while in test mode causes termination of the IMS conversation, response, or test mode and causes the associated output message to be dequeued. For conversation, the Conversational Abnormal Termination exit routine is invoked just as if an /EXIT command had been received.

When IMS is the secondary half session, it receives a special case of LUSTATUS BB/EB RQ*1 after responding with DR1 to a BID command sent by the primary half session. This LUSTATUS removes IMS (secondary) from a DFC in-brackets/pending state that was set after the aforementioned DR1 was sent and the primary half session has no other message available to be sent. This condition might occur when, for example, the application within the primary half session that caused the BID to be sent might have been abnormally terminated during the time between sending the BID and receiving the BID response.

LUSTATUS-CD should not be followed by another LUSTATUS-CD. Doing so creates a back-and-forth effect between the half sessions. IMS always responds LUSTATUS-EB when receiving LUSTATUS-CD and no output is available to be sent.

Related concepts:

“Handling IMS response mode or conversational output errors” on page 523

Paged messages ERP

For sender-detected errors (/DEQ command) detected after transmission of an IMS multichain demand or autopaged output message has been initiated, IMS must send the CANCEL command (as necessary) in conjunction with LUSTATUS - abort.

The CANCEL command is used to terminate a single multi-RU chain, or page, while the LUSTATUS is used to terminate the process receiving the multichain, or paged, message. If IMS receives a CANCEL during an input autopaged message without receiving the subsequent LUSTATUS - abort on the next input RU, IMS returns an exception response indicating selective receiver ERP followed by an FMH7 with the abort sense code X'0865' and an appropriate ERP message. Only the LUSTATUS - abort need be sent or received when the error is detected while between chains.

Errors detected on the first page of an autopaged input message result in IMS either sending a contention sense code or a selective receiver ERP sense code followed by an FMH7 indicating one of the nonfunction abort sense codes and an appropriate error message. Errors detected on nonfirst pages of an autopaged input message or for input page request received during an IMS demand-paged output message result in IMS sending the selective receiver ERP sense code followed by an FMH7 indicating the X'08650000' function abort sense code and an appropriate error message. The function abort sense code is used in these cases to cause the sending process (of autopaged input or demand-paged requests), rather than just a single chain or page, to be terminated.

Ready-to-receive protocol

IMS does not send the SNA ready-to-receive (RTR) command, but receives the command when acting as a primary half session.

When receiving the RTR command, IMS responds either with a DR1, followed by an available output message, or with an exception DR1 response (X'08190000'), if output is either not available or cannot be sent.

RSHUT protocol

IMS does not send, but does receive, the RSHUT command.

When receiving RSHUT, IMS sends a DR1 response and schedules termination of the session at the next convenient point. The session is normally terminated at the end of the current input or output chain.

Selective receiver ERP

Selective receiver error recovery procedure (ERP) is initiated using the X'08460000' sender ERP sense code.

The procedure has these characteristics:

- The procedure is symmetrical to both half sessions.

- ERP messages can be sent by the sender of the exception response indicating selective receiver ERP. An ERP message is an ERP FMH7 header followed by an IMS error message. Each selective receiver ERP message is preceded by an ERP FMH7 function management header.
- Contention for ERP messages does not occur for the exception response sender.
- ERP messages can consist of a full chain of arbitrary length, but must immediately follow the exception response indicating selective receiver ERP.
- The ERP message is recognized by the receiver for special processing.
- The response sender has an opportunity to reset the function management level bracket and send/receive states.

The ERP messages sent by IMS have associated default IMS message output descriptors (MODs) and display output formats (DOFs). The error message is formatted using SNA character string (SCS) controls even when sent to a component defined with MFS.

If necessary, the sender of the exception response enters into an ERP PURGE mode until both half sessions' send/receive states are again synchronized. The receiver of the exception response must cause the resynchronizing event.

Related concepts:

 SNA character string controls (System Programming APIs)

Related reference:

“Error recovery procedure FM header” on page 562

“ERP purging” on page 542

“Sender ERP” on page 554

Selective receiver ERP sense codes supported

During output, IMS recognizes receipt of the SNA-defined selective receiver ERP system sense code when assuming either half-session role, returns the message to the queue, ensures that a synchronizing event occurs to end the other half session's ERP purge cycle, and enters a receive state that can be satisfied only by input (ERP message) from the other half session.

When receiving this ERP message, IMS takes one of three actions based on the FMH7 sense code and the DFC protocols associated with the received ERP message:

- The ERP message (DFS2083) is routed to the IMS master terminal operator as a warning that an error condition was recovered on the ISC session. The ISC output message in error is either dequeued or retransmitted, and normal input or output operations continue. The message is retransmitted when an FMH7 is received with sense code X'0866' and without EB. The message is dequeued when an FMH7 is received with EB during test, response, or conversational mode output.
- The ERP message (DFS2073) is routed to the IMS LTERM associated with the IMS terminal operator who was the source (using an IMS message switch) of the ISC message switch output in error. The ISC output message in error is dequeued and normal input or output operations continue.
- The ERP message (DFS2073) is routed to the IMS master terminal operator and the ISC session is terminated. The ISC output message in error remains on the IMS message queue.

All ERP messages must carry an EB or CD or the session is terminated. Further, EB received with the ERP message forces end of IMS conversation, response, or test mode. When conversation mode is terminated, the Conversation Abnormal Termination exit routine is also invoked as would occur when an IMS /EXIT command is received. IMS processes FMH7 ERP messages received based on the FMH7 sense code and the DFC protocols associated with the message as in these subtopics.

X'0864xxxx': function abort

The loop occurs upon retransmission. The sender should not resend the data.

The following table shows how IMS processes the FMH7 ERP messages with sense code X'0864xxxx'.

Table 88. IMS processing for sense code X'0864xxxx'

Message type	MTO DFS2083	MTO DFS2073	TERMINAL DFS2073
Response mode			
FMH7 W/EB	X		
FMH7 W/CD ¹	X		
Non-last conversation			
FMH7 W/EB ²	X		
FMH7 W/CD		X	
Last conversation			
FMH7 W/EB ²	X		
FMH7 W/CD ^{1, 2}	X		
Message switch			X
Test mode			
FMH7 W/EB	X		
FMH7 W/CD		X	
Other Message Types			
FMH7 W/EB	X		
FMH7 W/CD ¹	X		

Notes:

1. FMH7 received with CD schedules LUSTATUS-NO-OP (X'0006') with EB to be returned next on the session if the output in error was a response mode reply, the last conversational output, or an ATTACH without SCHEDULER.
2. Conversation Abnormal Termination exit routine is invoked as would occur for /EXIT.

X'0865xxxx': function abort

The sender is responsible for detecting the loop.

The following table shows how IMS processes the FMH7 ERP messages with sense code X'0865xxxx'.

Table 89. IMS processing for sense code X'0865xxxx'

Message type	MTO DFS2083	MTO DFS2073	TERMINAL DFS2073
Response mode			
FMH7 W/EB	X		
FMH7 W/CD		X	
Non-last conversation			
FMH7 W/EB ¹	X		
FMH7 W/CD		X	
Last conversation			
FMH7 W/EB ¹	X		
FMH7 W/CD		X	
Message switch			X
Test mode			
FMH7 W/EB	X		
FMH7 W/CD		X	
Other Message Types			
FMH7 W/EB		X	
FMH7 W/CD		X	

Notes:

1. Conversation Abnormal Termination exit routine is invoked as would occur for /EXIT.

Related concepts:

 SNA character string controls (System Programming APIs)

X'0866xxxx': function abort

The receiver is responsible for detecting the loop.

The following table shows how IMS processes the FMH7 ERP messages with sense code X'0866xxxx'.

Table 90. IMS processing for sense code X'0866xxxx'

Message type	MTO DFS2083	MTO DFS2073	TERMINAL DFS2073
Response mode			
FMH7 W/EB	X		
FMH7 W/CD	X		
Non-last conversation			
FMH7 W/EB ²	X		
FMH7 W/CD	X		
Last conversation			
FMH7 W/EB ²	X		
FMH7 W/CD	X		
Message switch	X		

Table 90. IMS processing for sense code X'0866xxxx' (continued)

Message type	MTO DFS2083	MTO DFS2073	TERMINAL DFS2073
Test mode			
FMH7 W/EB	X		
FMH7 W/CD		X	
Other Message Types			
FMH7 W/EB	X		
FMH7 W/CD ¹	X		

Notes:

1. FMH7 received with CD schedules LUSTATUS-NO-OP (X'0006') with EB to be returned next on the session if the output in error was a response mode reply, the last conversational output, or an ATTACH without SCHEDULER.
2. Conversation Abnormal Termination exit routine is invoked as would occur for /EXIT.

During input, IMS can detect many types of input errors and internal processing conditions, such as undefined transaction codes, incorrect transaction formats, security violations, SNA protocol and data structure errors, and sync-point request errors. These errors can occur on almost any SNA request element of the message and might result in IMS sending an exception response indicating a selective receiver ERP sense code; purging the input message until a synchronizing CANCEL command, RQD1 or RQD2, change-direction indicator, or CHASE command is received; and then sending a selective receiver ERP message.

IMS provides the following sense information in the ERP FMH7 sent with the ERP message:

X'08260000'

Used for input and internal processing errors, other than the errors described in the remainder of this list, that are detected where a normal IMS ERP message is sent.

X'10030000'

Sent when the ATTACH ATTPRN value is not known to IMS.

X'08650000'

A function abort sense code for which the data sender is responsible for detecting and preventing loops. Function abort X'08650000' is sent by IMS after receiving an unsolicited CANCEL command (sender-detected error) during an MFS-autopaged input message that does not result in a DFC between-brackets state. Function abort X'08650000' is also sent by IMS for errors detected on input page requests received during an IMS demand-page output message and for errors detected on nonfirst pages of an autopaged input message.

X'084B0000'

Sent when the attach ATTDPN value is not available for the component entering the data. That is, the input DPN value was not ISCEDT, the optional ISC edit alias, or basic edit; or MFS was not available for the component. However, if MFS was available but the DPN value was not a valid MID, sense code X'08260000' occurs rather than X'084B0000'.

X'080F0000'

Sent when the ATTACH data stream profile (DSP) value does not define a valid or defined component number for the half-session name within IMS.

X'1008xxxx'

The user field (xxxx) is defined by SNA as follows:

X'6001': Invalid ATTACH FM header ATTDDBA value

X'1204': Invalid version ID on FMH4

Other types of errors detected during input or output operations that cause immediate session termination are:

- Bracket protocol violations
- Send/receive protocol violations
- Unsupported response requests
- Unsupported or invalid FMH types and formats
- Other major VTAM-detected errors or conditions such as LOSTERM, buffer pool or copy RPL space exhausted, or input RU truncation.

Related reference:

“Paged messages ERP” on page 549

Receiver-detected errors during data flow reset state

The following sense codes can be sent and received when the session flow is not in data flow active state, as, for example, during bind, resynchronization (STSN), or start data traffic (SDT).

X'0845xxxx'

Bind rejected because of invalid user data.

X'0847xxxx'

Restart mode mismatch detected by secondary half session at receipt of BIND1 STSN or SDT.

X'08210000'

Bind rejected because of invalid session parameters.

X'080Dxxxx'

Bind rejected because of bind race.

Related concepts:

“Resolving a bind race” on page 507

Sender ERP

The SNA CANCEL request can be used during sender error recovery processing as necessary to terminate the chain in error and to provide synchronization between half sessions. CANCEL can be either solicited (by the sender of the exception response) or unsolicited (because of a receiver-detected error, which results in an exception response).

If necessary during chain RU or autopaged output, the sender of the exception response enters into an ERP PURGE mode until both half sessions' send/receive states are again synchronized. The receiver of the exception response must cause the resynchronizing event. This event can be one of the following:

- An RQD1 or RQD2
- Receipt of change-direction
- A CANCEL request (caused by either RQD1 or RQE1/CD), if the exception response occurs during an output chain

- A CHASE request (because of RQD1 on CHASE), if the exception response is received while IMS is between RQE1 output chains (pages) of an autopaged output message

Related reading: For more information on sender error recovery procedure (ERP), see *z/OS Communications Server: SNA Programming*.

Sender ERP sense codes

The only sender ERP sense codes sent between two IMS subsystems are X'08130000' and X'0846xxxx'.

The valid contention sense codes received by IMS and additional operations performed are:

- X'08130000'—Bracket reject; no RTR sent
No ready-to-receive (RTR) condition is set. When the pseudo-wait is satisfied, the message is retransmitted from the beginning. This sense code can also be used to reject the SNA BIS command sent by IMS.
- X'08140000'—Bracket reject; RTR sent
The message is returned to the queue, and an RTR pending state is entered. No output is sent while IMS is between-brackets until an RTR is received. However, messages flow, subject to the bracket and send/receive protocol defined for the message type or component, when IMS is again left in an in-brackets/SEND state following receipt of change-direction.

After either of these sense codes, IMS attempts to send the output message.

IMS recognizes receipt of contention system sense code X'08130000' when assuming a primary half-session role, returns the output message to the message queue, and enters a pseudo-receive state that can be satisfied by:

- Receiving input from the secondary half session
- Being posted for output as a result of an IMS master terminal operator command, a message switch from another logical unit, or additional messages inserted by an IMS application program

After either of these actions, IMS attempts to send the output message.

IMS only indicates contention when assuming a secondary half-session role and receiving either input data or normal flow commands (BID and BIS) after having already initiated a bracket to the primary half session. IMS sends the X'08130000' sense code to indicate contention in these cases. The X'08140000' contention sense code is not sent by IMS.

Additional exception response sender ERP sense codes sent and received are:

- X'08190000'—No output available
This code is sent by IMS when assuming the primary half-session role and no output is immediately available following receipt of a ready-to-receive (RTR) indicator. IMS does not receive this sense code because IMS does not send RTR.
- X'0846xxxx'—Selective Receiver ERP
IMS sends and receives the SNA selective receiver ERP sense code. IMS sends this sense code when a response mode, conversational transaction, or application abnormally terminates. xxxx is a user sense field that is ignored by IMS when it is received. However, IMS does include this sense code in the message that is sent to the master terminal or to the terminal operator that was the source of the message. During output, it is normally set to the binary value of the IMS error message number that is sent as the ERP message when sending the response.

However, the user sense code can be set to X'0000' as, for example, when the error, such as an application abend, occurs for a response mode or conversational transaction.

- X'0864xxxx'—Function abort. Loop occurs upon retransmission. Sender should not resend data.

IMS receives, but does not send, this function abort sense code. When receiving this code, IMS dequeues the associated output message (if still active) and then continues with normal input or output operations. xxxx is a user sense field ignored by IMS. An LUSTATUS - abort is sent by IMS if input is received prior to the end of an IMS MFS-paged output message.

The function abort sense code X'0864xxxx' cannot be sent to IMS during nonlast IMS conversational output; otherwise, the output message is returned to the queue. The master terminal operator is notified and the session is terminated. When the function abort sense code X'0864xxxx' is sent following the last conversational output message, the message is dequeued and the Conversational Abnormal Termination exit routine is invoked just as occurs for /EXIT.

- X'0865xxxx'—Function abort. Sender responsible to detect loop.

IMS receives, but does not send, the sender ERP abort sense code. If received, this code causes IMS to return a message to the queue and close the session. IMS retransmits the message from the beginning at the next opportunity after session restart. xxxx is a user sense field ignored by IMS.

- X'0866xxxx'—Function abort. Receiver responsible to detect loop.

This sense code causes IMS to retransmit the message from the beginning at the next opportunity. The session is not terminated. xxxx is a user sense field ignored by IMS.

- X'08670000'—Multichain ERP purge.

IMS sends the multichain purge sense code at the end of ERP purging. IMS receives the multichain ERP purge sense code only after receiving an exception response to a nonfirst or nonlast page of an MFS output autopaged message. Non-MFS autopaged output and the first and last page of an MFS output autopaged message are sent requesting definite response. Exception to these SNA chains can result only in single chain purge.

Other sender ERP sense codes cause IMS to notify the master terminal operator and terminate the session.

Related reference:

“Selective receiver ERP” on page 549

 ERP purging (Communications and Connections)

Sender-detected errors on nonpaged messages

When an error is detected after initiating transmission of a single SNA nonpaged, multi-RU message, the sender can send the SNA unsolicited CANCEL command.

IMS sends unsolicited CANCEL only if a current output message is terminated by an operator /DEQUEUE command entered prior to the last segment of the message (single SNA chain) being transmitted.

If IMS receives an unsolicited CANCEL while data is being received by IMS, the message is discarded or backed out.

Sense codes that IMS receives

A *sense code* is a 2-byte field containing the category and modifier defined for the particular exception condition that has occurred. The 2 bytes following the sense code can contain optional user data, and are not supported for all sense codes.

The following table shows the input sense codes that IMS receives.

Table 91. Sense codes that IMS receives

Input sense code name	Input sense code	Sender ERP	FMH7	LUSTATUS
Commit/NO-OP	- X'00060000'			X
Queue empty	- X'00070000'			X
Bind race	- X'080D0000'	X		
Contention	- X'08130000'	X		
Contention	- X'08140000'	X		
Invalid parms	- X'08210000'	X		
Sel Rcvr ERP	- X'08460000'	X		
Reject restart	- X'08470000'	X		
Abort	- X'08640000'	X	X	X
Abort	- X'08650000'	X	X	X
Abort	- X'08660000'	X	X	X
ERP purge	- X'0867'	X		
Other		Note 1	Note 1	Note 1

Notes:

1. Other sense codes cause IMS to notify the master terminal operator and terminate the session.

Related reading: For more information on sense codes, see *z/OS Communications Server IP and SNA Codes*.

Sense codes that IMS sends

The following table shows the output sense codes that IMS sends.

Table 92. Sense codes that IMS sends

Output sense code name	Output sense code	Sender ERP	FMH7	LUSTATUS
NO-OP	- X'00060000'			X
Queue empty	- X'00070000'			X
Bind race	- X'080D0000'	X		
ATTDSP	- X'080F0000'		X	
Contention	- X'08130000'	X		
No output	- X'08190000'	X		
Invalid parms	- X'08210000'	X		
Not supported	- X'08260000'		X	
Sel Rcvr ERP	- X'08460000'	X		

Table 92. Sense codes that IMS sends (continued)

Output sense code name	Output sense code	Sender ERP	FMH7	LUSTATUS
Reject restart	- X'08470000'	X		
ATTDPN	- X'084B0000'		X	
Abort	- X'08650000'	X		X
ERP purge	- X'0867'	X		
ATTPRN	- X'10030000'		X	
FMH4 Version ID	- X'10081204'		X	
ATTDBA	- X'10086001'		X	

SIGNAL protocol

IMS sends the SIGNAL command to request change-direction (SIGNAL RCD - X'00010000') at the end of an input message chain that does not indicate change-direction and that produces output that must be sent prior to subsequent input.

For example, IMS replies to input commands, response mode transactions, conversational transactions, or input while in test mode. SIGNAL RCD is sent before IMS sends any required DR1 or DR2 response for the input message.

Based on the IMS response protocol and the point at which IMS sends SIGNAL RCD, IMS requires that the next input following SIGNAL RCD be either LUSTATUS or CHASE indicating CD or EB. If neither of these is received by IMS, the session is terminated. Change-direction allows IMS to send the pending output message; end-bracket causes the message to be dequeued and the conversation mode, response mode, or test mode to be terminated. No other input can be processed until the pending output message has been successfully transmitted or dequeued.

After receiving SIGNAL RCD while in bracket state, IMS sends change-direction either at the end of the current output message or immediately on receipt of an LUSTATUS. A SIGNAL RCD received while between brackets prevents IMS from sending further output messages until normal flow input (data or SNA command) occurs.

Related reference:

"Signal protocol example" on page 663

Symmetrical session shutdown for LU 6.1 (SBI and BIS)

Two data flow control commands allow a symmetrical and orderly termination for peer level LU 6.1 half sessions: stop bracket initiation (SBI) and bracket initiation stopped (BIS).

SBI is a VTAM-expedited flow command, and BIS is a normal flow command.

These commands control only the normal flow requests that initiate new brackets. They do not preclude replies that can flow within existing brackets. Either half session can send SBI at any time to request the receiving half session to suppress future initiation of a bracket. After the receiver of the SBI reaches a point between brackets that is acceptable for shutdown, the receiver returns the BIS command to the SBI initiator. The BIS sender then enters a "no begin-bracket" or "NOBB" state, during which that half session can receive bracket initiation requests and may send

requests or replies in the normal flow when in-brackets, but cannot initiate a bracket on its own. The primary half session is responsible for detecting and resolving SBI race conditions by maintaining indications of which half session is in NOBB state and terminating the session if both half sessions reach this state.

A successful completion of a shutdown sequence refers to successful entry into the NOBB state by both half sessions— that is, normal session termination.

IMS sends the SBI command either when the IMS /QUIESCE command is entered or when the quiesce option is specified on the /CHECKPOINT command. When BIS is returned from the other subsystem, IMS continues with the QUIESCE or CHECKPOINT command process. If the /QUIESCE command was entered and a CHECKPOINT QUIESCE is not in progress, IMS sends BIS at the next transition to a between-brackets state. If a CHECKPOINT QUIESCE is in progress when IMS receives BIS, IMS continues to send output according to other /CHECKPOINT parameters (FREEZE, PURGE, and DUMPQ) and then sends BIS when the shutdown process is complete.

After receiving an SBI or BIS, IMS sends BIS at the next between-brackets state when a CHECKPOINT QUIESCE is not in progress, or at the end of the shutdown process when a CHECKPOINT QUIESCE is in progress.

When IMS is the primary half session, IMS checks after either sending or receiving BIS to ensure that both half sessions are in NOBB state. If both are, IMS resets all message counts and terminates the session. Otherwise, IMS continues with normal processing. When IMS is the secondary half session, IMS waits for the primary half session to terminate the session. If an LTERM subpool is allocated to the session, it is deallocated (and associated message counts are reset) prior to the session termination when both half sessions are in NOBB state. This deallocation/termination process allows a subsequent cold start of the session and frees any allocated subpools for use by other sessions. This process also ensures that no resynchronization or recovery is required when the session is restarted.

Because initiating a bracket is not allowed after the NOBB state is entered, a problem occurs when a single-bracket-chain (begin- and end-bracket) input message requires an immediate reply (such as a system or error message) that initiates a bracket. Rather than preclude end-bracket on these input messages while in NOBB state, IMS accepts and continues to process all input. When a reply occurs that would require IMS to initiate a bracket while in NOBB state, the session is terminated.

While IMS is in response mode, in conversational mode, or waiting for a Fast Path response, subpool deallocation cannot occur nor can BIS be sent, because no EB has flowed either on output or input (that is, IMS is not between-brackets).

Related reference:

“SBI/BIS examples” on page 660

Function management headers

In SNA, function management (FM) headers are an optional part of the request unit sent over a link. This topic describes the FM headers supported by IMS on ISC sessions.

The headers addressed in this topic include:

- The ATTACH function management header. This type 5 header is used to attach a process so it can receive session input. It carries the name of the process to be attached synchronously, as well as other parameters to be used by the attached process. All messages have an explicit or implicit ATTACH FM header. The header is implicit when not actually sent with the message. An implicit ATTACH FM header assumes reset parameter values or values of a previous ATTACH as defined later in this topic.
- The error recovery procedure (ERP) function management header. This type 7 header is sent after an error condition requiring an error message to be sent is detected.
- The reset attached process (RAP) function management header. This type 5 header is input to IMS to cause the attached process and the associated data flow control states to be reset.

The following header types are subordinate to the ATTACH FM header and cannot be used unless the ATTACH FM header has previously been sent within that bracket to attach the associated process:

- The SCHEDULER function management header. This type 6 header is sent after the SCHEDULER process has been attached. It carries the name of the process to be asynchronously scheduled within IMS, as well as other parameters required by the SCHEDULER process.
- The SYSMSG function management headers. The type 6 system message headers precede system messages sent by or received by IMS.
- The data descriptor function management header. This type 4 header is used in conjunction with MFS to send a data structure name, version ID, or a current output field tab separator.
- The QMODEL function management headers. These type 6 headers are used for demand-paged messages sent by IMS Message Format Service.

All function management headers must be on a first-in-chain or only-in-chain. The following table summarizes the header types that IMS sends and receives (Y=Supported). These header types are described in the topics that follow.

Table 93. Function management header types

Function management header type	Sent by IMS	Received by IMS
FMH4 - Data Descriptor	Y	Y
FMH5		
ATTACH	Y	Y
Reset Attach Process (RAP)		Y
FMH6 - QMODEL		
QGET		Y
QGETN		Y
QPURGE		Y
QSTAT	Y	
QXFR	Y	
FMH6 - SCHEDULER	Y	Y
FMH6 - SYSMSG		
SYSSTAT (default)		Y
SYSSTAT	Y	Y

Table 93. Function management header types (continued)

Function management header type	Sent by IMS	Received by IMS
SYSEERROR	Y	Y
FMH7 - ERP	Y	Y

When receiving a header, IMS ignores any coded parameters beyond the last parameter to which it is sensitive (that is, that IMS supports). The other LU 6.1 subsystem should do the same when receiving headers that IMS sends.

Using FM headers to invoke ISC edit

In the following situations, IMS can use ISC edit (ISCEDT) to edit transactions, commands, and message switches between LTERMs for an IMS-ISC session

- The attach manager is in the reset state (between-brackets or following a RAP), and either no ATTACH FM header or an ATTACH FM header that does not specify a destination process name (ATTPDN) is received.
- The destination process name (ATTPDN/SCDDPN) within the ATTACH or SCHEDULER FM header equals ISCEDT or indicates the user-defined alias for ISC edit as defined during system definition on the COMM macro statement.
- The attach manager is not in the reset state, the active process is ISC edit, and either no ATTACH FM header or a header that does not specify an ATTPDN is received.

Input that does not specify the ATTPRN/SCDPRN parameter is edited only for the transaction code and password at the beginning of the message. The optional IMS password and leading characters less than X'41' are deleted during editing. The format and requirements for the IMS transaction code and password are the same as for basic edit. No editing of any kind occurs for the remainder of the message.

Input parameters passed from the attach manager or SCHEDULER (if supplied) are return destination process name (ATTRDPN/SCDRDPN), return primary resource name (ATTRPRN/SCDRPRN), and primary resource name (ATTPRN/SCDPRN).

The input ATTPRN/SCDPRN parameter is passed from the attach manager or SCHEDULER as an IMS transaction code or as the LTERM name on a message switch, as defined under "Initiating a process: ATTACH FM header." When the ATTPRN/SCDPRN parameter is supplied, no editing of any kind occurs for the input. ATTPRN/SCDPRN defines an IMS destination, but does not provide an input password to IMS transaction security (if this security is defined for the ISC node).

Related concepts:



Basic edit (Communications and Connections)

Related reference:

Chapter 34, "Examples using ISC edit ATTACH parameters," on page 597

Initiating a process: ATTACH FM header

The ATTACH FM header is used to attach a process to receive session input. It carries the name of the process to be attached synchronously, as well as other parameters used by the attached process.

Because ATTACH is defined for synchronous execution, you must understand the relationships between the process to be attached, the IMS message types and execution modes, and the SCHEDULER process.

IMS sends the ATTACH FM header without a SCHEDULER FM header:

- On the conversational, response mode, or IMS command output message replies resulting from an input ATTACH.
- On asynchronous replies that result from a message received on the same session with ATTACH EB.
- When the other half session is bound without SCHEDULER model support.
- On system messages sent by IMS through SYSMSG.

An IMS message can consist of either single or multiple SNA chains as indicated within the ATTACH ATTIU parameter. The ATTACH FM header can be present only once for each input or output IMS message. For output MFS demand-paged messages, the ATTACH FM header is present only on the first output SNA chain. The first input paging request to the demand-paged output must contain an ATTACH (for DPN=QMODEL); subsequent input page requests for the same demand-paged output message can optionally contain an ATTACH (which must also indicate DPN=QMODEL).

Related concepts:

“Relationship of ISC and IMS execution modes” on page 475

Related reference:

“ATTACH FM header format” on page 577

Error recovery procedure FM header

IMS sends the FMH7 ERP header and its associated message following a selective receiver ERP exception response (X'0846').

An ERP exception response need not be received by IMS before receiving an FMH7 ERP header and associated message. FMH7 messages sent by IMS indicate either RQE1/CD or RQD1/EB, as appropriate. ERP messages that are sent or received by IMS are limited to a single SNA chain of SCS characters and must not contain VLVB format records. FMH7 messages received by IMS must indicate either CD or EB. RQE1 or RQD1 is allowed on either.

Single-segment error messages that are created within IMS are 79 characters or less in length. Multi-segment messages that are created within IMS can be greater than 79 characters in length, but each segment has a 79-character maximum length.

Error messages that are created by input to the ISC error message process for output to the master terminal (or source operator terminal for a message switch) create single- and multi-segment output messages (DFS2073 or DFS2083).

IMS inserts error messages into IMS message queues by using a default MFS MOD name. These messages are formatted as defined by the MOD during output to a component defined with MFS DPM. These MODs and associated DOFs are contained in the MFS format library.

Each selective receiver ERP message is preceded by an ERP message header that includes the specific system and user error sense codes.

Related reference:

“Selective receiver ERP” on page 549

“Error recovery procedure (ERP) FM header” on page 587

Resetting the active process: RAP FM header

The reset attached process (RAP) FM header is used to reset the receiving half session's active process and all session states, except bracket and send/receive states, to the equivalent of a between-brackets state.

IMS receives the RAP FM header, but does not send it. After receiving the RAP request, IMS responds with an ATTACH SCHEDULER, or with LUSTATUS - queue empty if no output is available to be sent. If the RAP request is received during an IMS demand-paged output message, the message is dequeued (committed) before a check for more output is made. The IMS operation performed for the RAP request is the same as for the NEXTMSG operator control request provided with MFS input or the PA2 request from a 3270 device.

The RAP request must be sent without data and must indicate change-direction, RQD2, or RQE2.

The following figure illustrates the use of the RAP FM header.




Figure 69. RAP FM header example

Related reference:

“Reset attached process (RAP) FM header format” on page 594

Requesting asynchronous transaction processing: SCHEDULER FM header

ATTACH is defined for synchronous scheduling and execution; SCHEDULER is defined for asynchronous scheduling and execution.

The SCHEDULER process can be attached to send messages to be processed asynchronously. That is, IMS perceives the messages as being processed without regard to scheduling or execution timing, whether output will result at all, and ignoring any synchronous relationship between the input and any resulting reply. Each half session between the logical units might indicate, by using the session bind parameters, whether the SCHEDULER process is supported and available to receive asynchronous messages. The IMS half session always indicates the

SCHEDULER process is available. The other half session might indicate that the SCHEDULER process is or is not available to receive asynchronous input.

When the bind indicates that the other half session does not support the SCHEDULER process, IMS defaults to sending each output with the ATTACH FM header.

The SCHEDULER process is attached by indicating the SNA name (X'02') as the ATTACH DPN parameter and by concatenating the SCHEDULER FM header to the ATTACH FM header. During the time that the SCHEDULER is attached, each subsequent message to be scheduled asynchronously must be sent with at least the SCHEDULER FM header, but does not require another ATTACH FM header. Variable-length ATTACH parameters passed to the SCHEDULER are ignored on input and are not sent with output by IMS. The SCHEDULER FM header carries the name of the process to be scheduled (SCDDPN), as well as information necessary for scheduling and return-reply routing. SCDDPN is a required parameter for the SCHEDULER and must be supplied on IMS SCHEDULER input. SCDDPN on messages sent by IMS are set from a previous input SCDRDPN or optionally by MFS using the output message format descriptor. When no other value is available, SCDDPN is set to "ISCEDT" as a default.

The SCHEDULER process is always available for IMS input. IMS sends an ATTACH with the ATTDPN parameter set to X'02' (SCHEDULER) concatenated with a SCHEDULER FM header under each of the following conditions:

- When the other half session supports the SCHEDULER model (as defined in the bind parameters).
- When sending an asynchronous reply resulting from input entered previously on the session and scheduled using the SCHEDULER process.
- When sending unsolicited asynchronous messages. Examples of such messages are an IMS message switch or an output reply resulting from an application not scheduled directly as a result of input on the session.

Within IMS, the SCHEDULER supports both single- and multiple-chain input and output messages using the ATTACH ATTIU parameter. IMS does not permit multiple input or output messages to be scheduled using the ATTACH multiple-chain indicator.

The SCHEDULER FM header can be present only once for each input or output IMS message. For output MFS demand-paged messages, the SCHEDULER FM header is present only on the first output SNA chain (first output MFS page). IMS does not send or receive the SCHEDULER SCHEDSTAT, PURGE, or PURGSTAT FM headers.

Related reference:

"Data flow control protocol reference" on page 533

"SCHEDULER FM header format" on page 594

System message process (SYSMSG) and related FM headers

The system message process is indicated by a process name in the ATTACH FM header.

Input system messages are routed to the master terminal operator when no ATTPRN is supplied in the ATTACH or SYSERROR FM header for a SYSMSG. This routing is accomplished by converting the incoming SYSMSG into IMS system message DFS2072. If an ATTPRN value is supplied, it becomes the IMS input

message destination transaction code or LTERM name. The system message process cannot be used to access the IMS command processor.

A reply can result from an input SYSMMSG if the ATTPRN value supplied is for an IMS transaction. This reply is returned through ISC edit or MFS, depending upon the output component definition.

IMS inserts system messages into IMS message queues using a default MFS MOD name. These messages are formatted as defined by the MOD during output to a component defined with MFS DPM. These MODs and associated DOFs are contained in the MFS format library.

IMS sends SYSMMSG only for:

- IMS broadcast messages
- System messages that IMS sends after sending a response to input. These system messages can occur instead of the application sending a reply (for example, when an abnormal termination occurs). IMS system messages directly solicited by an IMS command are sent as normal replies. Other system-generated messages can result from error conditions. These messages are converted to exception responses to the input transaction and are sent as error messages using the error recovery process (ERP).

MFS might edit an output SYSMMSG, depending on the output component definition and whether the reply was inserted with an MFS MOD name.

Messages sent to or received by SYSMMSG can include either the SYSSTAT or SYSERROR FM header. The SYSSTAT FM header is assumed by the receiver if no header is supplied by the sender. IMS can receive either header, and attaches a prefix to each header with IMS message number DFS2072 (if the message is to be sent to a master terminal operator). IMS sends broadcast output using the SYSSTAT header. All other IMS system messages sent as SYSMMSG use the SYSERROR header. If the ATTACH ATTRDPN or ATTRPRN parameters are supplied on an input message that results in SYSMMSG output, these parameters are included in the output SYSERROR FM header to be used by the receiver's SYSMMSG process. The ATTACH ATTRDPN and ATTRPRN parameters are not sent for an IMS output SYSMMSG.

Single-segment system messages created within IMS are 79 characters or less in length. Multisegment messages created within IMS can be greater than 79 characters in length, but each segment has a 79-character maximum length.

System messages created by input to the ISC SYSMMSG process for output to the master terminal create single- and multisegment output messages. System messages created by input to the ISC SYSMMSG process in which ATTPRN and ATTDPN are specified must have a MID capable of handling multisegment messages with a minimum of five segments.

Related reference:

“SYSMMSG process headers” on page 595

Chapter 32. Using MFS with ISC

Message formatting for MFS DPM is specified on the IMS system definition `TERMINAL` macro statement by the parameter `DPM-Xn`, where X might be A or B. For ISC nodes, the form `DPM-Bn` is always used.

This topic addresses the following:

- How MFS `DPM-Bn` can be invoked to format IMS input and output
- The way in which the `ATTACH`, `SCHEDULER`, and reset attached process (RAP) function management headers are used by MFS
- MFS support for the SNA-defined `QMODEL` and data descriptor function management headers (which control MFS processing)

Some input MFS format errors (such as invalid cursor, incorrect output formats, or no output formats) are not detectable within IMS before IMS sends the requested response to the input message. In these cases, IMS sends the error message by using `ATTACH ATTDPN=SYSMSG`. This `ATTACH` results even if the input message indicated change-direction and exception response and is an implicit acknowledgment to the message.

An `ATTACH SYSMSG` received by IMS during demand-paged output is treated as an invalid MFS paging request and causes session termination. The received `SYSMSG` is discarded.

Related reference:

Chapter 33, “FM header format reference,” on page 577

Activating MFS input formatting

When MFS is used, input messages can be processed by the message and format descriptors.

When IMS receives an input message from an ISC logical unit, basic edit or ISC edit is performed unless MFS is defined for the input component (`ATTDSP`) and a MID name accompanies the message. The MID name can be supplied by including it either as the `ATTDPN` parameter on the `ATTACH` function management header or as the `SCDDPN` parameter on the `SCHEDULER` FM header.

The MFS escape characters (//) are not supported for ISC. When the MID name is present, MFS edits the message using the specified MID and its associated device input format (DIF).

Activating MFS output formatting for ISC

MFS output formatting occurs when an output message has an associated message output descriptor (MOD).

The MOD is supplied in one of the following ways:

- The application program supplies a MOD name with the output message.
- The input message is processed by a message input descriptor (MID) whose definition specifies a MOD name for output formatting.

- The output message is a message switch that is created by a MID whose definition specifies a MOD name for output editing.

If no MOD is associated with the output message, no MFS editing occurs. Values for the ATTACH and SCHEDULER parameters are the default values that are defined under these respective header types.

You can use the MOD to specify the next MID name to be used to format input as the result of output. In an ISC session, IMS cannot control the relationship between formatted IMS output messages and subsequent input messages. Therefore, involvement is required by the remote subsystem or user application program to ensure that the proper MID name is used for input MFS formatting. The ATTACH and SCHEDULER RDPN/DPN parameters perform the ISC function of externalizing the next MID on output and invoking MFS formatting on subsequent input.

Related concepts:

 Relationships between MFS control blocks (Application Programming APIs)

Related reference:

“ATTDPN” on page 581

MFS Distributed Presentation Management (DPM) messages

An IMS MFS DPM message can consist of either single or multiple SNA chains as indicated within the ATTACH parameters.

For MFS DPM input or output using a device format defined with the paging OPTIONS=MSG, IMS sends or receives the entire message as a single chain of one or more related transmissions. For MFS DPM output using a device format defined with paging OPTIONS=DPAGE or PPAGE, IMS sends each logical or presentation page as a single chain of one or more related transmissions.

MFS output DPM format definitions also allow options for paged output (OPTIONS=PPAGE or DPAGE) to be sent as demand-paged or autopaged. Demand-paged output requires paging requests to the other half session for each output page. Autopaged output is sent as a series of consecutive pages (chains) with no paging requests to the other half session.

MFS DPM input using a device format defined with the paging OPTIONS=DPAGE allows IMS to optionally receive each logical page as a single chain of one or more related transmissions. All chains of autopaged input (multiple logical pages as multiple related input chains) are received consecutively without paging requests by IMS.

FM headers can occur as part of the first or only transmission of each input or output chain of the message.

MFS page delete function

The MFS page delete function is not available for use by an ISC session.

OPTIONS=NPGDEL is forced on the IMS system definition TERMINAL macro statement or an ETO user descriptor, because paging requests other than QMODEL-architected paging requests are not supported. This includes MFS operator control table requests and other forms of input data that might occur

during output demand-paging. An error (exception response and appropriate ERP message) occurs if data other than a valid QMODEL paging request is received during demand-paged output. The output message is returned to the queue as a result of the error condition and retransmitted at the next opportunity. If the output demand-paged message is conversational or response mode output, it is immediately resent and input is not allowed until the output message is successfully transmitted or dequeued by a valid QMODEL paging request or appropriate ERP action.

MFS online error detection

This topic describes how IMS detects input errors, output errors, and paging errors for MFS.

Output errors

MFS output errors are detected after IMS has already sent the response to any preceding input. If an error is detected during MFS MOD or DOF block selection, an error message is sent as ATTACH SYSMSG, and the IMS message is returned to the message queue for retransmission. If the MFS test mode is in effect, it is reset.

Invalid page requests cause error messages to be sent.

Input errors

An error message is sent and the input message is rejected if one of the following occurs:

- An error is detected during MFS MID or DIF block selection.
- A nonzero version ID is received in the data descriptor FM header and it does not match the version ID in the MFS descriptor.
- An error is detected during DPAGE selection (that is, no condition is satisfied on matching the DPAGE label with the DSN in the DD FM header or the COND= with the data).

If one of the following errors is detected during multiple DPAGE input, an error message is sent to the other subsystem, and the input message is rejected:

- Multiple transmission chains. More data is present in the chain than defined for the DPAGE selected.
- Any mapped-input LPAGE contains no data segments (as a result of segment routines canceling all segments, for example).

If one of the following errors is detected during a single DPAGE input (that is, multiple DPAGE input is not requested in MFS definitions), an error message is sent and the input message is rejected:

- A single transmission chain is received and contains more data than defined for the DPAGE selected.
- Multiple transmission chains are transmitted.
- The mapped-input message contains no data segments (as a result of segment routines canceling all segments, for example).

If the input message is canceled by the User Segment Edit exit routine, or if a User Segment Edit exit routine failure is detected, an error message is sent to the other subsystem. In the latter case, the input message is rejected.

Paging errors

An error message is sent to the other subsystem when an invalid paging request (QMODEL FM header) is detected:

- The QNAME parameter on the QGETN or QGET function management header does not match the ATTDQN/SCDDQN parameter on the ATTACH or SCHEDULER FM header sent by IMS. An attached demand-paged output message is returned to the message queue for retransmission. A scheduled demand-paged output message continues to wait for a QGETN with the proper QNAME specified.
- The QORG parameter on the QGETN or QGET FM header is invalid. The output message is returned to the message queue for retransmission.
- The QGETN FM header is received and no output message is in progress.
- A paging request other than QGETN is received as the first input following scheduled demand-paged output (ATTACH SCHEDULER), or a QSTATUS FM header is sent by IMS because of an invalid cursor on a paging request.
- The QGETN FM header is received for an OLP demand-paged output and current cursor position is at the last page.

The QGET FM header causes a QSTATUS or an error message to be sent in the following cases:

- Between messages.
- Non-OLP demand-paged output message is in progress. An error message is sent and the output message is placed on the queue for retransmission.
- Cursor value does not contain a valid 2-byte binary number. An error message is sent, and the output message is placed on the message queue for retransmission.
- Cursor value is outside the range for the output message.

Related reference:

“QSTATUS FM header” on page 575

“Error recovery procedure FM header” on page 562

The ATTACH and SCHEDULER FM headers under MFS

All input and output messages include an implicit or explicit ATTACH FM header and, optionally, a SCHEDULER FM header.

MFS DPM provides format definition options to permit the insertion of input routing parameters (ATTRDPN/SCDRDPN, ATTRPRN/SCDRPRN, and ATTPRN/SCDPRN) into the input data stream. The input ATTDPN/SCDDPN is the MFS MID used to format the input message. MFS DPM provides output format definition options to permit the insertion of user-defined information into these headers.

The ATTACH FM header and SCHEDULER FM header, if required, are sent without data as the only element of the first transmission chain of demand-paged output. The first input paging request to the demand-paged output must contain an ATTACH (for DPN=QMODEL); subsequent input page requests for the same demand-paged output message can optionally contain an ATTACH (and must also indicate DPN=QMODEL). For all other MFS input and output messages the ATTACH (and SCHEDULER, if required) FM headers are sent with data as the first or only element in the first or only transmission chain of the message.

MFS uses the following ATTACH and SCHEDULER FM header parameters:

- ATTDN and ATTIU. These ATTACH parameters indicate whether the message is demand-paged or whether it is single or multiple chain.
- ATTDN. This ATTACH parameter indicates the output data blocking algorithm or input data deblocking algorithm to be used by IMS. The data entity deblocked on input becomes the input data record to MFS. MFS then produces a standard IMS message consisting of pages and segments. On output, the MFS record becomes the data entity for output VLVB blocking. An exception is MFS DPM stream mode, which is sent as chain (ATTDN) output.
- ATTDN/SCDDN and ATTPN/SCDPN. The ATTDN/SCDDN parameter is used on input to activate MFS input formatting. The ATTPN/SCDPN can be inserted into the input data stream and, therefore, made available to the application. On output message replies from IMS, the ATTDN/SCDDN or ATTPN/SCDPN can be inserted into the ATTACH or SCHEDULER FM headers of the output message processed by MFS as specified by MFS format descriptions. These override the ATTRDN/SCDRDN and ATTRPN/SCDRPN that might have been wrapped from the source input message ATTACH or SCHEDULER FM headers.
- ATTRDN/SCDRDN. The ATTRDN/SCDRDN can optionally be inserted into the input data stream and, therefore, made available to the application for input messages processed by MFS DPM. The return destination process name can be inserted into the output ATTACH or SCHEDULER FM header of the output message processed by MFS. The MFS-suggested return destination process to be associated with a reply to the output message is the MID name specified in the NXT= operand of the MOD.
- ATTRPN/SCDRPN. The return primary resource name (RPN) can be inserted into the ATTACH or SCHEDULER FM header of the output message processed by MFS. It can be specified using MFS definition. If returned to IMS with a reply to the message on which it was sent, it determines the destination — transaction or LTERM — that is to receive the input reply.
- ATTDQ/SCDDQ. The destination queue name is sent by IMS on demand-paged output as a message identifier and should be returned on all requests associated with demand-paged messages.

When the first chain (ATTACH SCHEDULER) of an asynchronous demand-paged output message is successfully sent, the output queue (IMS LTERM) is automatically locked against input or output. This function is also automatic following a QSTATUS FM header sent by IMS because of a cursor error on a demand-paged request. The queue lock state is reset when the page request (ATTACH QMODEL, QGETN with QNAME) is received and the first page of data is sent without resending the ATTACH or SCHEDULER FM header. The queue lock condition is also reset across session failure and subsequent restart. During the time that IMS waits for this page request (QGETN), provided conversation or response mode is not used, IMS can send or receive messages for other LTERMs. During the time that the output LTERM is locked, no input is accepted from the ISC session that produces output to the locked queue because of the input/output component relationships defined for the locked queue using the IMS system definition NAME macro statement or an ETO logon descriptor. Input is accepted that produces output for other LTERMs.

Data descriptor FM headers

The input and output data descriptor FM headers are used to receive and send a data structure name and version ID or a field tab separator character.

Input data descriptor FM header

On input to IMS, this header can follow the ATTACH or SCHEDULER FM header and is used to receive a data structure name (if DPAGE selection is by DPAGE name) and version ID or the input field tab separator character.

The input data descriptor FM header should be sent to IMS if input DPAGE selection is to be performed on the DPAGE name; that is, if you specify OPTIONS=DNM on the DIV statement (TYPE=INPUT).

If you set OPTIONS=DNM, and specify no DPAGE name or the wrong DPAGE name, an error occurs.

The version ID should be sent to IMS in the first or only FM header of the input message if MFS is to verify that the correct definition is used to map the data. If the version ID is sent and is X'0000', no verification occurs. The field tab separator is not required on input, because MFS provides for up to eight separators with the FTAB function. If received on input, the field tab separator is used instead of the MFS FTAB specification (if any) for the current transmission.

Related reference:

"Data descriptor FM header formats" on page 585

Output data descriptor FM header

On output from IMS, the data descriptor FM header is used to send a data structure name and version ID or the output field tab separator character if OFTAB= parameter of the DIV or the DPAGE statement is specified.

It is sent in the only transmission chain of a nonpaged output message or in each transmission chain of a paged output message. For a nondemand-paged message (OPTIONS=MSG), this header follows the ATTACH or SCHEDULER FM header. For an autopaged message, this header follows the ATTACH or SCHEDULER FM header and precedes the data for the first logical or presentation page. The data descriptor FM header precedes the data for each additional page of output and is the only FM header in the transmission chain. For demand-paged output, this header follows the QFXR header. The version ID is sent only once for each message in the first or only FM header for the output message. Additionally, if OFTAB is specified and OPTIONS=DNM is requested, the output field tab separator character used for the current transmission is included in the FM header. If OPTIONS=NODNM is specified, a data descriptor FM header is not sent, regardless of whether OFTAB has been defined or not.

Related reference:

"Data descriptor FM header formats" on page 585

Controlling demand-paged messages: QMODEL FM headers

Queue model (QMODEL) headers are sent and received to control demand-paged messages. All demand-paging requests (except RAP) must be made using the QMODEL-defined FM headers.

Restriction: MFS DPM-Bn does not support operator control table requests or input messages while sending demand-paged output.

The QMODEL headers that IMS supports are listed in the following table.

Table 94. IMS-supported QMODEL headers

QMODEL headers	Sent by IMS ¹	Received by IMS ²
QXFR	X	
QGETN		X
QGET		X
QPURGE		X
QSTATUS	X	

Notes:

1. QMODEL headers are not sent between two IMS subsystems.
2. A QMODEL FM header must precede a QMODEL page request for the first page. These FM headers can be preceded by the QMODEL ATTACH FM header with the DPN parameter containing the value X'03'.

QXFR is a QMODEL reply. All other headers are QMODEL requests. Each half session between the logical units might indicate, by using the session bind parameters, whether the QMODEL process is available to receive QMODEL requests. The bind parameters have no effect on QMODEL replies. IMS always indicates "QMODEL available" and is prepared to receive QMODEL (paging) requests to IMS demand-paged output. The other half session might or might not indicate "QMODEL available." IMS ignores the other half-session bind indication for QMODEL requests, because only QMODEL replies can be sent by IMS.

Queue model uses LU 6.1 protocols to handle messages as follows:

- Although SNA allows multiple messages to be active, IMS messages are serial. That is, only one message is active at any given time. When active message processing is completed, another message can be processed. Synchronous (ATTACH without SCHEDULER) output MFS DPM demand-paged messages are considered active immediately when the first chain (ATTACH) is sent. Demand-paged output sent with ATTACH SCHEDULER does not become active until the first page request is returned to IMS.
- MFS demand-paged output provides for two types of message organization: linear (OPTIONS=DPAGE) and hierarchic (OPTIONS=PPAGE). For both types, pages can be retrieved sequentially (get next) or linearly (get by cursor). However, linear retrieval is permitted only if operator logical paging (OLP) is defined. Hierarchic retrieval is not supported.

Related reference:

"Data flow control protocol reference" on page 533

Request (input) QMODEL FM headers

The following topics describe the QGETN, QGET, and QPURGE FM headers.

QGETN FM header

A QGETN FM header is received by IMS to cause sequential transfer of a single logical or presentation page of MFS DPM demand-paged output. It is the only QMODEL FM header in the chain and is never followed by function management data.

The QGETN FM header is not sent by IMS.

The QGETN FM header can be sent to IMS in the following cases:

- To retrieve first page of demand-paged output. A QGETN must always be the first request following the receipt of a scheduled demand-paged output. It must be preceded by an ATTACH QMODEL FM header.
- Following a QSTATUS FM header that indicates an invalid cursor in scheduled demand-paged output.
- When a non-OLP demand-paged output message is in progress.
- When an OLP demand-paged output message is in progress and the current cursor position is not at the last page of the message.

Each of these cases causes a QXFR FM header to be returned in reply.

The QGETN FM header cannot be sent to IMS in the following cases:

- Between messages. An error message is returned.
- While an OLP demand-paged output message is in progress and the current cursor position is at the last page.

Related reference:

“QSTATUS FM header” on page 575

“QGETN FM header format” on page 589

QGET FM header

Use the QGET FM header as input to IMS to transfer to the next, the last, or any logical page, according to cursor.

A QGET FM header is only valid for demand-paged output with operator logical paging defined. It is the only FM header in the input chain and is never followed by function management data.

The QGET FM header is not sent by IMS.

The QGET FM header cannot be sent to IMS in the following cases:

- Between messages. An error message is returned.
- After receiving the ATTACH for a scheduled demand-paged output message. The first paging request for this type of output must be QGETN. QGET can be used after requesting the first output page by QGETN.
- While a non-OLP demand-paged output message is in progress. An error message is returned. The output message is placed on the queue for retransmission.
- When the cursor value does not contain a valid 2-byte binary number. An error message is returned, and the output message is placed on the message queue for retransmission.
- When the cursor value is outside the range for the output message.

Related reference:

“QSTATUS FM header” on page 575

“QGET FM header format” on page 587

QPURGE FM header

When IMS receives a QPURGE header, it halts processing on the demand-paged output message in progress.

QPURGE can be received by IMS only while an MFS demand-paged output message is active. The QPURGE FM header is the only FM header in the chain and is never followed by function management data. IMS replies to a QPURGE FM header with a QSTATUS. If IMS receives a QPURGE header to delete the demand-paged output message, the message is deleted only after IMS receives a sync-point response in reply to QSTATUS.

The QPURGE FM header is not sent by IMS.

Related reference:

“QPURGE FM header format” on page 590

Reply (output) QMODEL FM headers

The following topics describe the QXFR and QSTATUS FM headers.

QXFR FM header

The QXFR header is sent in each output transmission chain for a demand-paged output message in reply to a valid paging request by a QMODEL header.

A QXFR FM header can be followed by the data descriptor (DD) FM header (based upon MFS format definitions) and carries logical or presentation page data. This header must not be sent to IMS.

Related reference:

“QXFR FM header format” on page 591

QSTATUS FM header

A QSTATUS FM header is sent by IMS in reply to an input QPURGE FM header or because an invalid cursor was detected on a page request for a demand-paged output.

If a QSTATUS results from a QPURGE during demand-paged output, it is sent to request a sync point. The message is dequeued when the sync-point response is received. The message is returned to the queue and made available for retransmission if an exception sync-point response is received.

A QSTATUS resulting from an invalid cursor during scheduled demand-paged output is sent RQE1/CD and allows a subsequent QGETN to again request the first page of the message or a QPURGE to cause the message to be dequeued.

The QSTATUS FM header is the only FM header in the chain and is not followed by function management data. The QSTATUS FM header must not be sent to IMS.

Related reference:

“QSTATUS FM header format” on page 590

The RAP FM header under MFS

When used by MFS, the reset attached process (RAP) FM header can be sent to IMS to delete the demand-paged output message in progress (prevent further processing).

This header is equivalent to the MFS operator control function of NEXTMSG. This header can also be sent to IMS to delete the operator logically paged message in progress. A RAP FM header cannot be followed by function management data. The RAP FM header is not sent by IMS.

Related reference:

“Reset attached process (RAP) FM header format” on page 594

Chapter 33. FM header format reference

Each header's length is defined by a 1-byte length field. The value in this field includes the length field itself. A concatenation flag indicates whether an additional header follows.

Within the header formats, variable- and fixed-length parameters are positional by command code. A 1-byte length field precedes each variable-length positional parameter. The value contained in this length field can be X'00' through X'08' and does not include the length field itself. If the length field contains X'00', the variable parameter is omitted and the next positional variable-length parameter length field occurs followed by its variable-length parameter field.

Trailing positional parameter length fields of X'00's at the end of the header can be eliminated for input to IMS and are not sent on output from IMS. IMS also does not send trailing blanks for any of these names.

When receiving a header, IMS ignores any coded parameters beyond the last parameter to which it is sensitive (that is, that IMS supports). The other LU 6.1 subsystem should do the same when receiving headers that IMS sends.

Related reference:

Chapter 36, "ISC data flow control examples," on page 653

Chapter 34, "Examples using ISC edit ATTACH parameters," on page 597

ATTACH FM header format

The format of the ATTACH FM header is defined in the following table.

Table 95. Attach FM header format

Byte	Bits	Name	Content
0		FMHL	
1	0	FMHC	
	1-7	FMHT	B'0000101'
2-3		FMH5CMD	X'0202'
4	0-3	FMH5MOD	
	4		Reserved
	5	ATTDTP	Session Local Flag
			B'0' Not demand-paging
			B'1' Demand-paging
	6-7	ATTIU	Interchange Unit code
			B'00' Multiple chain
			B'01' Single chain
			B'10' Reserved
			B'11' Reserved
5		FMH5FXCT	Length of fixed-length parameters must be X'02'

Table 95. Attach FM header format (continued)

Byte	Bits	Name	Content
6		ATTDSP	One-byte hexadecimal input value into a data stream profile. (IMS requires a value from X'00' to X'03' to select components 1-4 respectively. Value X'00', or component 1, is assumed if no ATTACH FM header is received while the Attach manager is in reset state.)
7		ATTDBA	Application data handling algorithm X'00' Undefined (Input only to IMS. IMS also assumes this as the default value if no ATTACH FM header is received while the attach manager is in reset state.) X'01' Variable length, variable blocked (input/output for IMS) X'02' Document subset of SCS (not supported by IMS) X'03' Card subset of SCS (not supported by IMS) X'04' A chain of RUs (input to/output from IMS) X'05' An RU (input only to IMS) X'06' to X'FF' Reserved
8-m		ATTDPN ¹	Name of a process to be initiated. For an SNA-defined process, the DPN is a 1-byte, nongraphic hexadecimal character. Supported SNA processes are: X'01' SYSMSG X'02' SCHEDULER X'03' QMODEL The IMS processes are described under the ATTDPN parameter.
m+1-n		ATTPRN ¹	Name of primary resource for the process being initiated
n+1-p		ATTRDPN ¹	Name of suggested return process name
p+1-q		ATTRPRN ¹	Name of suggested primary resource for the return process
q+1-r		ATTDQN ¹	Name of queue to be associated with the attached process
r+1-s		ATTACC	Access code—Ignored by IMS.

Notes:

1. If the ATTDPN is set to SCHEDULER (X'02'), the remaining ATTACH parameters must not be specified. In this case, an equivalent for each of the asterisked parameters might occur in the concatenated SCHEDULER FM header. These SCHEDULER parameters are prefixed by "SCD" rather than by "ATT", but follow the same definition and rules as those for ATTACH.

ATTIU

The interchange unit code (ATTIU) indicates whether a single input or output IMS message consists of one or more SNA chain.

All input messages that are not MFS-autopaged messages might indicate either single or multiple chain, but are restricted to one actual transmission chain indicating end-of-message. MFS-autopaged input messages might indicate multiple chain regardless of whether the message is actually one or more than one SNA chain. MFS also offers an option that permits single chain messages to produce multiple DPAGES. Non-MFS and MFS output that is not autopaged is always sent as a single transmission chain with ATTIU indicating single chain. MFS-autopaged output messages are always sent with ATTIU indicating multiple chain.

Restriction: IMS does not support any relationship between consecutive input or output messages. Therefore, the interchange unit code indicating multiple chain cannot be used to send IMS batched input messages or consecutive messages to be scheduled using the SCHEDULER process.

IMS requires attached (synchronous) input message switches to be one chain with at least EB specified.

Restriction: IMS does not support MFS-autopaged synchronous input message switches of more than one SNA chain.

The ATTIU parameter is not automatically retained from one input or output message to another. This parameter must be explicitly provided with each input message, or by MFS if the output message is multichain and is formatted by MFS. The input reset state for ATTIU is "multiple chain."

ATTDSP

The data stream profile (ATTDSP) parameter is ignored when attaching QMODEL (ATTDPN=X'03'). During other input, the ATTDSP field indicates a specific IMS input component.

The input component is used to identify an input IMS LTERM from a set of LTERMs (LTERM subpool) allocated to the session. The input DSP value must indicate a valid input component (an IMS LTERM defined with the input component value), or the input is rejected.

Several internal IMS functions are associated with the LTERM defined for the input component. For example, the LTERM defines a particular origin (input) and destination (output) path within IMS. This output destination can be the same as, or different from, the input origin LTERM and represents a stable application and operator reference point within IMS.

For statically defined terminals, input terminal security can be defined for the input LTERM. This security authorization can be unique for each LTERM or can represent a security level, or group, within IMS. Input terminal security authorizes access to IMS terminals, transactions, and commands.

The input ATTACH DSP parameter remains in effect until it is:

- Changed by another ATTACH
- Reset by a system failure
- Backed out to its value at the last sync point by the ERP

- Reset at between-brackets state
- Reset by the RAP FM header

The input reset state for the ATTDSP parameter is X'00' or "component 1."

Because IMS does not remember the active input ATTDSP across IMS system failures, the ATTDSP must be provided using an ATTACH FM header to restore the receiving process after a bind or negotiable bind response indicating "in-brackets" (possible restart).

The output LTERM is determined by either the LTERM selected for input or by the message sender, as for message switches and broadcast messages. IMS also allows the receiving message processing program (MPP) to change the destination of resulting output by issuing a CHNG call and inserting the output to a modifiable alternate PCB. For all system messages resulting directly from input (for example, commands), the output LTERM selected is the normal destination LTERM defined for the LTERM from which the input originated.

During output, the DSP field is set to the output component value associated with the output LTERM.

IMS Message Format Service can be defined as available on a component-by-component basis. The output component also specifies the bracket and send/receive protocol to be used for asynchronous output sent using the ATTACH having a DPN parameter indicating "SCHEDULER."

Related concepts:

"Coordinating the restart process" on page 512

ATTDBA

The deblocking algorithm (ATTDBA) determines the data blocking and deblocking algorithm to be used.

This algorithm, along with an indication of whether MFS is being used, determines the amount of data presented to or from a process for a single get or put operation.

When IMS is sending messages and MFS is not being used, each segment of the message on the message queue becomes a data entity within the RU (or chain of RUs). If a segment of the message spans message queue records (LRECLs), then each spanned portion of the segment becomes a data entity. If MFS is used, the segments are blocked and sent as a data entity according to the definition of the MFS record.

When IMS receives messages, each data entity of the input message is inserted in the message queue as a segment of the message. Because ISC does not support spanned queue segments on input, the largest message queue LRECL must be large enough to handle the largest data entity or MFS data record received. If this LRECL definition is exceeded, an error is detected and an appropriate error message is produced (DFS074). Also, the actual space available in the message queue is reduced by any variable prefix items on the message.

IMS implements four of the algorithms defined by SNA:

UNDEFINED

Same as RU, below, for IMS.

RU Input RU is the entity presented to the attached process. IMS receives this DBA value but does not send it.

VARIABLE LENGTH, VARIABLE BLOCKED (VLVB)

Each data entity sent or received is preceded by a 2-byte length field that allows it independence from RU size or boundaries. Several data entities can be blocked into a single RU, or a data entity can span RUs. The length includes the 2 bytes for the length field itself.

CHAIN OF RUs

Each data entity is sent or received as a single SNA chain.

Each message received by IMS is deblocked based on the ATTDDBA field in the ATTACH FM header. The input ATTDDBA field value must indicate a supported algorithm, or the input is rejected.

The input ATTACH DBA parameter remains in effect until it is:

- Changed by another ATTACH
- Backed out to its value at the last sync point by the ERP
- Reset at between-brackets state
- Reset by the RAP FM header
- Reset by a session failure

The input reset state for the ATTDDBA parameter is “UNDEFINED” (equivalent to “RU”).

Because IMS does not remember the active input ATTDDBA across IMS system failures, the ATTDDBA must be provided using an ATTACH FM header to restore the receiving process after a bind or negotiable bind response indicating “in-brackets” (possible restart).

With the exception of output using MFS stream mode, output from IMS is automatically sent indicating the variable-length, variable blocked (VLVB) algorithm in the ATTACH or SCHEDULER FM header. This output is sent indicating chain assembly.

Related concepts:

“Coordinating the restart process” on page 512

ATTDPN

The destination process name (ATTDPN) parameter identifies, explicitly or implicitly, the input process to be attached to the half session.

One responsibility of the attached process is to determine, for the receiving subsystem, the destination of the input message within that subsystem. In some cases, the named process might be the message destination. Within IMS, the receiving process (except basic edit) uses the input primary resource name (ATTPRN), if provided, as the message destination transaction queue or LTERM. If an ATTPRN is not provided, the standard IMS algorithm for determining message destination is used.

Related reading: For more information on the IMS algorithm for determining message destination, see *IMS Version 11 System Administration* under “IMS Messages and Their Scheduling.”

All attached processes execute synchronously with the session. However, some processes can schedule additional work to be done asynchronously within IMS. The process attached remains attached until it is:

- Changed by another ATTACH
- Backed out to its value at the last sync point by the ERP
- Reset by the end-bracket indicator
- Reset by the RAP FM header
- Reset by a system failure

The input reset state for the ATTDPN parameter is "ISCEDT".

Because IMS does not remember the active input ATTDPN across IMS system failures, the ATTDPN must be provided using an ATTACH FM header to restore the receiving process after a bind or negotiable bind response indicating "in-brackets" (possible restart).

The ATTDPN parameter might indicate that the process to be attached is one of the following:

- ISC edit ('ISCEDT' or user-named alias)
- IMS Basic Edit ('BASICEDT')
- MFS formatting (MFS MID name)
- System Message (SYMSG), X'01'
- SCHEDULER model, X'02'
- Queue model (QMODEL), X'03'

With the exception of MFS and QMODEL, these processes are always available. MFS is only available if defined for the input component (ATTDSP) during IMS system definition. If the ATTDPN is an MFS MID name and MFS is not available, the input message is rejected. QMODEL is available only following an MFS demand-paged output message.

ISCEDT (as well as the alias defined by the user for ISCEDT during IMS system definition) and BASICEDT are reserved names, and cannot be used as MFS MID or MOD names. Use of these names results in the use of the named process rather than an MFS process.

ISC edit is selected if no ATTACH FM header (or no ATTDPN) is supplied when the Attach manager is in reset state. The "active" process is used if no ATTACH FM header (or no ATTDPN) is supplied when the attach manager is not in the reset state.

IMS sets the output ATTDPN on a return reply to the value contained in the ATTRDPN parameter optionally provided within the ATTACH of the previous input request. When sending system messages, IMS automatically inserts the SYMSG ATTDPN where necessary. MFS DPM provides a way to optionally specify, override, or delete the output ATTDPN by using the output message format descriptor. When the ATTDPN parameter is not available for output, a 1-byte field containing X'00' is included within the output ATTACH FM header.

Related reference:

"ATTPRN" on page 583

"ATTRDPN and ATTPRN" on page 583

ATTRPN

The ATTACH primary resource name (ATTRPN) parameter represents the destination for an input message in the receiving subsystem.

This parameter is sent on a return reply from a remotely executed message that results from the returning reply message of the ATTACH ATTRPN parameter that was sent on a message request. However, in situations where the message request contains no input ATTRPN, the output ATTRPN parameter can be created using MFS, based on user-defined information.

IMS does not permit specification of IMS command verbs as PRNs.

This destination normally represents a terminal (an LTERM), rather than an application program. If the request ATTRPN/reply ATTRPN represents an application program, consider the synchronous versus asynchronous relationship between the source application program and the remote transaction executions.

Further, you should consider input transaction security requirements for processing and routing the reply. If supplied on input to IMS, the ATTRPN is used as an IMS transaction code or as the LTERM name for a message switch. The data stream is not edited for destination and security information. If the ATTRPN represents a transaction, and if password security (which applies to statically defined terminals only) was defined for the node, an input security error results, because no available source exists for the input password.

Related reading: For more information on password security, see *IMS Version 11 System Administration*.

The ATTRPN is rejected during IMS conversation mode when an application has previously inserted the transaction code into the SPA. The ATTRPN only applies to a single message instance and overrides, but does not destroy, any “preset” destination established using an IMS /SET command. Subsequent messages with no ATTACH FM header or with an ATTACH FM header where the ATTRPN parameter is not supplied, can use any preset destination previously established by the IMS /SET command or by an IMS transaction code, LTERM name, or command verb supplied in the data or through MFS formatting.

IMS sets the output ATTRPN on a return reply to the value contained in the ATTRPN parameter optionally provided within the ATTACH of the previous input request.

MFS DPM provides a means to optionally specify, override, or delete the output ATTRPN using an output message format descriptor. When this parameter is not present, a 1-byte field containing X'00' is sent with the output ATTACH FM header.

The ATTRPN is not automatically retained from one input or output message to another. This parameter must be explicitly provided with each input message, or through the output MFS descriptors if the output message is processed by MFS.

Related reference:

“ATTRDPN and ATTRPN”

ATTRDPN and ATTRPN

The return destination process name (ATTRDPN) and return primary resource name (ATTRPN) parameters define the reply process and the return primary

resource within the source session and should be returned to the source session on resulting replies to facilitate return-reply routing within the source session.

If provided on input to IMS and not changed or deleted by an output MFS format description, these parameters are returned to the source session unmodified in the reply ATTACH FM header as the output ATTDPN and ATTPRN parameters respectively. These parameters are associated with the message to be processed and are recovered with the message across session and subsystem failures.

The input ATTRDPN and ATTRPRN are not automatically retained from one input or output message to another. The ATTRDPN and ATTRPRN must be explicitly provided with each input message or through the MFS output format descriptors when MFS DPM is used for output. The ATTRDPN results from MFS if a next MID was specified in the MOD.

A different procedure is followed when the reply is to be returned on a session different from the session on which the input message originated. If the output ATTRPRN is not set by MFS, and the source of the output message is not associated with the same session, IMS does not wrap the input ATTRDPN and ATTRPRN as the output ATTDPN and ATTPRN respectively. In this case, IMS also automatically inserts the source LTERM name of the terminal entering the input message switch or transaction (for alternate PCB output) as the output ATTRPRN parameter.

When the ATTRDPN parameter is not available for output, a 1-byte length field containing X'00' is included in the output ATTACH FM header.

The following tables summarize IMS actions relative to the DPN, PRN, RDPN, and RPRN fields in the ATTACH and SCHEDULER headers sent with the message.

Table 96. IMS interpretations for the DPN, PRN, RDPN, and RPRN fields

Input FMH	IMS interpretation
DPN	MFS MID name or input message editor name (ISC edit or basic edit). Defaults to ISCEDT if not supplied.
PRN	Input transaction code or LTERM name, overriding data stream and preset mode except during conversation.
RDPN	Saved as default reply DPN field.
RPRN	Saved as default reply PRN field.

Table 97. IMS actions for the DPN, PRN, RDPN, and RPRN fields

Output FMH field	IMS ACTION reply message sent on same session as input	IMS ACTION reply message or asynchronous output sent on another session
ATTDPN	Wrapped input ATTRDPN if provided; optionally overridden by MFS. For ATTACH-only output, if no parameter is available, the DPN is not sent. For SCHEDULER output, the default is ISCEDT.	Provided only using MFS. For ATTACH-only output, if no parameter is available, the DPN is not sent. For SCHEDULER output, the default is ISCEDT.
ATTPRN	Wrapped input ATTRPRN; optionally overridden by MFS	Provided only using MFS

Table 97. IMS actions for the DPN, PRN, RDPN, and RPRN fields (continued)

Output FMH field	IMS ACTION reply message sent on same session as input	IMS ACTION reply message or asynchronous output sent on another session
ATTRDPN	Provided only using MFS	Provided only using MFS
ATTRPRN	Provided only using MFS	Automatically defaulted to source LTERM name; optionally overridden by MFS

ATTDQN and ATTDQ

The destination queue name (ATTDQN) parameter names a specific message instance. This parameter is valid only for output MFS demand-paged messages.

For MFS demand-paged output, ATTDQN is sent by IMS as an output message identifier. This name can then be used by the other half session as the QNAME parameter within paging requests to access the IMS demand-paged output message. Within IMS, only one message can be active for a given session at any one time. This means that after output paging begins on a message and until the paging operation is terminated, all input must be paging requests indicating the same QNAME (from ATTDQN) value. The session-local flag (ATTDQ) in the ATTACH FM header is set to 1. ATTDQN and ATTDQ are not sent or received under other conditions.

ATTACC

IMS does not send the access code (ATTACC) parameter during output and ignores this parameter on input.

Data descriptor FM header formats

The following tables show the formats of the input and output data descriptor FM headers.

Table 98. Input data descriptor FM header format

Byte	Bits	Name	Contents
0	0-7	FMHL	Length
1	0	FMHC	
	1-7	FMHT	B'0000100'
2	0-7	FMH4FXCT	Fixed-length parm (X'03')
3	0-7	FMH4DTYP	
		FMH4UNDF	X'00' Reserved
		FMH4FIX	X'40' Field formatted record (FFR) - Fixed
		FMH4FXSP	X'41' FFR - All fields terminated by Separator (Note 1)
		FMH4MXSP	X'42' FFR - Fields terminated by Separator or Length defined by map X'43' - X'FF' Reserved (Note 1)
4		FMH4SEP	Separator character (Note 2)
5	0-7	FMH4PCTL	Presentation control byte (Note 3)
6-m		FMH4DSN	Data structure name (Note 4)
m+1-n		FMH4BDT	Block data type (Note 3)

Table 98. Input data descriptor FM header format (continued)

Byte	Bits	Name	Contents
n+1-p		FMH4VERS	Version ID (Note 5)

Notes:

1. From DIV or DPAGE statement.
2. Not required, because it is received on MFS format description.
3. Ignored by IMS if sent.
4. DPAGE name if DPAGE selection is to be performed on the data structure name. FMT name if OPTIONS=MSG. Required field.
5. Version ID to be used by MFS to verify that the correct definition is used to map the data. This field should be present only once for each input message if multiple transmission chains are used to create the input message. If present once for each message, it should be sent in the first transmission chain of the message. If the version ID is not sent or it is X'0000', MFS descriptor verification is not performed.

Table 99. Output data descriptor FM header format

Byte	Bits	Name	Contents
0	0-7	FMHL	Length
1	0	FMHC	
	1-7	FMHT	B'0000100'
2	0-7	FMH4FXCT	Fixed-length parm (X'03')
3	0-7	FMH4DTYP	
		FMH4UNDF	X'00' Reserved
		FMH4FIX	
			X'40' Field formatted record (FFR) - Fixed
		FMH4FXSP	
4			X'41' FFR - All fields terminated by separator
		FMH4MXSP	
			X'42' FFR - Fields terminated by separator or length defined by map X'43' - X'FF' Reserved
4		FMH4SEP	Separator character
5	0-7	FMH4PCTL	Presentation control byte (Note 1)
6-m		FMH4DSN	Data structure name (Note 2)
m+1-n		FMH4BDT	Block data type (Note 3)
n+1-p		FMH4VERS	Version ID (Note 4)

Notes:

1. Set to 0 on output.
2. FMH4DSN is:
 - The FMT name if OPTIONS=MSG
 - The DPAGE name if OPTIONS=DPAGE
 - The PPAGE name if OPTIONS=PPAGE

3. Omitted if version ID not present. Set to 0 if version ID present.
4. Version ID specified in FMT definition or calculated by MFS if the default is used. This field is present only in the first or only data descriptor FM header for the message.

Error recovery procedure (ERP) FM header

The following table shows the format of the ERP FM header.

Table 100. Error recovery procedure (ERP) FM header

Byte	Bits	Name	Contents
0		FMHL	Length
1	0	FMHC	
	1-7	FMHTYPE	X'07'
2-5		ERPSENSE	SNA sense code that would appear on error response
6-7		ERPSEQ	Sequence number of chain for which error was detected

QMODEL FM headers

The beginning of each header is fixed-length. This fixed-length area is followed by a variable number of fixed-length subfields, followed by a variable number of variable-length subfields.

QGET FM header format

The following table shows the format of the QGET FM header.

Table 101. QGET FM header format

Byte	Bits	Name	Contents
0		FMHL	Length
1	0	FMHC	
	1-7	FMHT	Type: B'0000110'
2-3		FMH6CMD	Command Code: X'0A10'
4		FMH6MOD	Modifier
	0	FMH6LNSZ	B'0' 1-Byte Length Fields
	1-6		Reserved
	7	QGETLAST	B'1' Coded request, get last ¹
5		FMH6FXCT	Length of Fixed-Length Parameters (=X'01')

Table 101. QGET FM header format (continued)

Byte	Bits	Name	Contents
6		QORG	Type of paging request X'00' Queue organization not specified. Valid on input to IMS. X'01' Sequential. Invalid. X'02' Linear. Valid on input to IMS if it matches the QORG in the QXFER sent by IMS. ² X'03' Hierarchic. Not supported by IMS. X'04' to X'FF' Reserved
7-m		QNAME	Q from which records are retrieved. IMS message ID (ATTDQN/SCDDQN) ³
m+1-n		QCURSOR	Cursor vector in target queue ⁴
n+1-p		QTRNSZ	Size of maximum record to be returned (binary). Ignored by IMS if sent. All fixed-length parameters must be present when you use this header.

Notes:

1. The QGETLAST selects a logical page whose value is equal to the last logical page in the message.
The QGET FM header is valid and causes a QXFR reply only if the OLP demand-paged output message is in progress and the cursor value requested is within the range of the output message.
2. If QORG is invalid, an error message is sent and the output message is returned to the message queue for retransmission.
3. The QNAME (message identifier) is required on each page for scheduled demand-paged output and is optional for synchronous demand-paged output. If the QNAME parameter is present, it must match the ATTDQN/SCDDQN parameter in the ATTACH FM header sent by IMS. If QNAME is not specified for synchronous demand-paged output, IMS assumes it to be the same as the ATTDQN parameter name. If the QNAME parameter does not match the ATTDQN/SCDDQN parameter, an error message is sent and the output message is returned to the message queue for retransmission.
4. The cursor vector consists of one 2-byte binary number representing the logical page number. IMS does not support hierarchic retrieval; therefore, the 2-level cursor (logical page number followed by the presentation page number) is invalid and results in an error message being issued. The length value of the QCURSOR is 2. The next 2 bytes contain the logical page number (absolute) request.

QGETN FM header format

The following table shows the format of the QGETN FM header.

Table 102. QGETN FM header format

Byte	Bits	Name	Contents
0		FMHL	Length
1	0	FMHC	
	1-7	FMHT	Type: B'0000110'
2-3		FMH6CMD	Command Code: X'0A10'
4		FMH6MOD	Modifier
	0	FMH6LNSZ	B'0' 1-Byte Length Fields
	1-7		Reserved
5		FMH6FXCT	Length of Fixed-Length Parameters (=X'01')
6		QORG	Type of paging request X'00' Queue organization not specified. This form is valid for both OLP and non-OLP output. X'01' Sequential. ¹ X'02' Near. ¹ X'03' Hierarchic. Not supported by IMS. X'04' to X'FF' Reserved.
7-m1		QNAME	Q from which pages are retrieved. Specifies a message ID (ATTDQN/SCDDQN) ²
m+1-n		QTRNSZ	Size of maximum record to be returned (binary). Ignored by IMS if sent.

Notes:

1. If QORG is specified, the QORG specified on the resulting QXFR FM header sent by IMS must match it. Otherwise, an error message is sent and the output message is returned to the message queue for retransmission.
2. The QNAME (message identifier) is required on each page request for scheduled demand-paged output and is optional for synchronous demand-paged output. If the QNAME parameter is present, it must match the ATTDQN/SCDDQN parameter in the ATTACH FM header sent by IMS. If QNAME is not specified for synchronous demand-paged output, IMS assumes it to be the same as the ATTDQN parameter name. If the QNAME parameter does not match the ATTDQN/SCDDQN parameter, an error message is sent and the output message is returned to the message queue for retransmission.

QPURGE FM header format

The following table shows the format of the QPURGE FM header.

Table 103. QPURGE FM header format

Byte	Bits	Name	Contents
0		FMHL	Length
1	0	FMHC	
	1-7	FMHT	Type: B'0000110'
2-3		FMH6CMD	Command Code: X'0A06'
4		FMH6MOD	Modifier
	0	FMH6LNSZ	B'0' 1-Byte Length Fields
	1-7		Reserved
5		FMH6FXCT	Length of fixed-length parameters (=X'01')
6		QORG	Type of queue purge function (if multiple DPAGE input) or paging request (if demand-paged output) ¹
7-m		QNAME	Name of IMS message ID (ATTDQN/SCDDQN) ²

Notes:

1. If IMS receives QPURGE during demand-paged output, the QORG, if specified (that is, QORG is not equal to 0), must match the QORG specified in the QFXR FM header sent by IMS.
2. The QNAME (message identifier) is required on each page request for scheduled demand-paged output and is optional for synchronous demand-paged output. If IMS receives QPURGE during demand-paged output, the QNAME, if specified, must match the ATTDQN/SCDDQN parameter (message ID) in the ATTACH FM header sent by IMS. If QNAME is not specified for synchronous demand-paged output, IMS assumes the same name as the ATTDQN parameter name. If the QNAME parameter does not match the ATTDQN/SCDDQN parameter, an error message is sent and the output message is returned to the message queue for retransmission.

QSTATUS FM header format

The following table shows the format of the QSTATUS FM header.

Table 104. QSTATUS FM header format

Byte	Bits	Name	Contents
0		FMHL	Length
1	0	FMHC	
	1-7	FMHT	Type: B'0000110'
2-3		FMH6CMD	Command code: X'0A0A'

Table 104. QSTATUS FM header format (continued)

Byte	Bits	Name	Contents
4		FMH6MOD	Modifier
	0	FMH6LNSZ	B'0' 1-Byte length fields
	1-7		Reserved
5		FMH6XCT	Length of fixed-length parameters (=X'02')
6		QORG	Type of queue add/retrieval requests. ¹
7	0-4		Not supported by IMS. Bit values should be B'0'.
	5	QINVCUR	B'0' if message is a reply to QPURGE. B'1' Invalid Cursor Indicates invalid logical page request received for OLP output.
	6-7		Not supported by IMS.
8			Reserved
9-12		QCURSOR	Current cursor value is 0 ²
13-16		QSENSE	Sense data ³
17-n		QNAME	Message ID (ATTDQN/SCDDQN)

Notes:

1. If QSTATUS is sent by IMS in reply to the QPURGE, QORG is not specified (that is, QORG contains the value of zero). If QSTATUS is sent by IMS in reply to QGET by cursor request, QORG is set to linear (that is, QORG contains the value of two).
2. This parameter is present only if the QINVCUR value is set to B'1'. The cursor value is set to 0. A QGETN request retrieves the first page of the message for which this error (invalid cursor) was detected.
3. This parameter is present only if the QINVCUR value is set to B'1'. It consists of the DFS223 error message as a binary number.

QXFR FM header format

The following table shows the format of the QXFR FM header.

Table 105. QXFR FM header format

Byte	Bits	Name	Contents
0		FMHL	Length
1	0	FMHC	
	1-7	FMHT	Type: B'0000110'
2-3		FMH6CMD	Command Code: X'0A08'
4		FMH6MOD	Modifier
	0	FMH6LNSZ	B'0' 1-Byte Length Fields
	1-7		Reserved

Table 105. QXFR FM header format (continued)

Byte	Bits	Name	Contents
5		FMH6FXCT	Length of fixed-length fields (=X'02')
6		QORG	Types of paging requests valid for this message. X'01' Sequential retrieval of pages. Operator logical paging (OLP) is not defined. Message is dequeued after successful transmission of last page. X'02' Linear retrieval of pages. Operator logical paging (OLP) or browsing is allowed. Message is dequeued only by explicit action from the other subsystem (RAP or QPURGE FM header is received). X'03' Hierarchic retrieval of pages. Not supported by IMS. X'00', X'04'-X'FF' Reserved.
7	0-4	Reserved	
	5	QDISP	B'0' Disposition is save.
	6		Reserved.
	7	QEMSG	B'0' Not end of message B'1' End of message Set to B'1' if last logical page (OPTIONS=DPAGE) or last presentation page (OPTIONS=PPAGE) of a message is transmitted.
8-n		QCURSOR ¹	Cursor vector for current message. For OPTIONS=PPAGE, this field includes a 1-byte length field of value 4 followed by the 2-byte current logical page number and the 2-byte current presentation page number. For OPTIONS=DPAGE, this field includes a length field of value 2, followed by the current logical page number. ²
n+1-n		QCOUNT ¹	Number of occurrences of pages at lowest level of cursor. If present, this field includes a length byte of value 2, followed by the defined number of pages in the current logical page. The QCOUNT field is present only if the defined number of presentation pages in the current logical page is greater than 1.

Table 105. QXFR FM header format (continued)

Byte	Bits	Name	Contents
n+1-p		QRECLNG	Length of record before truncation. If 0, then either record was not truncated or QTRNSZ on QGET(N) was ignored. This field is not supported by IMS and will not be included in the FM header.

Notes:

1. If operator logical paging (OLP) is allowed, the QCURSOR and QCOUNT fields contain information the receiver can use for formulating subsequent QGET by cursor requests.
2. Assume the following output message with OPTIONS=PPAGE:

Logical Page 1	3 Presentation Pages Defined
Logical Page 2	1 Presentation Page Defined
Logical Page 3	2 Presentation Pages Defined

The following table illustrates the values placed in the QCURSOR and QCOUNT fields with sequential retrieval requests.

Table 106. QCURSOR and QCOUNT values with sequential retrieval requests

Transmission	Count/LP	No./PP No.	Qcursor Count/PP No.	Qcount comment
1	4 / 1	1	2 / 3	Logical Page 1
2	4 / 1	2	2 / 3	Logical Page 1
3	4 / 1	3	2 / 3	Logical Page 1
4	4 / 2	1	(omitted)	Logical Page 2
5	4 / 3	1	2 / 2	Logical Page 3
6	4 / 3	2	2 / 2	Logical Page 3

Note:

Output Message with OPTIONS=DPAGE:

Logical Page 1
Logical Page 2
Logical Page 3

The following table illustrates the values placed in the QCURSOR field with sequential retrieval requests.

Table 107. QCURSOR values with sequential retrieval requests

Transmission	Qcursor Count/LP	Qcount	Comment
1	2 / 1	Not Present	Logical Page 1
2	2 / 2	Not Present	Logical Page 2
3	2 / 3	Not Present	Logical Page 3

Reset attached process (RAP) FM header format

The following table shows the format of the Reset Attach Process (RAP) FM header.

Table 108. Reset attached process (RAP) FM header format

Byte	Bits	Name	Contents
0		FMHL	
1	0	FMHC	
	1-7	FMHT	B'0000101'
2-3		FMH5CMD	X'0204'
4		FMH5MOD	
	0	FMH5LNSZ	B'0': 1-byte parameter length field
	1-7		Length of parameter length fields B'0': 1-byte parameter length field B'1': Reserved
5		FMH5FXCT	Length of fixed-length parameters (X'00')

SCHEDULER FM header format

The following table shows the format of the SCHEDULER FM header.

Table 109. SCHEDULER FM header format

Byte	Bits	Name	Contents
0		FMHL	
1	0	FMHC	
	1-7	FMHT	B'0000110'
2-3		FMH6CMD	X'0802'
4		FMH6MOD	
	0	FMH6LNSZ	B'0': 1-byte parameter length field
	1	FMH6RPLY	B'0' No reply B'1' Reply ¹
	2	FMH6PROT	B'0' No protection; ² B'1' Protection
	3	FMH6DELY	B'0' Timer optional B'1' Timer required ¹
	4-7		Reserved
5		FMH6FXCT	Length of fixed-length parameters - X'01'
6			Schedule Request control
	0	SCDTIME	Initiation delay specification B'0' Interval B'1' Time Not supported by IMS
	1-7		Reserved
7-m		SCDDPN ³	Name of process to be initiated. (Required field)

Table 109. SCHEDULER FM header format (continued)

Byte	Bits	Name	Contents
m+1-n		SCDPRN ³	Name of the primary resource for the process to be initiated
n+1-p		SCDRDPN ³	Suggested name of the return process
p+1-q		SCDRPRN ³	Suggested name of the primary resource for return process name
q+1-r		SCDDQN ³	Name of queue associated with process
r+1-s		SCDREQN	Name of this process request instance (Not supported by IMS)
s+1-t		SCDDELY	Time interval to be decremented prior to initiation of destination process. (Not supported by IMS)

Notes:

1. Must be 0 on both input and output.
2. On input, this field is ignored. On output, this field is set to 0 for nonrecoverable output and to 1 for recoverable output.
3. See “ATTPDN” on page 581, “ATTPRN” on page 583, “ATTRDPN and ATTRPRN” on page 583, and “ATTDQN and ATTDPN” on page 585 for a more detailed description of these parameters. In the ATTACH topic, these parameters are prefixed by “ATT” rather than “SCD”. The rules for operation of these SCHEDULER parameters are the same as the rules defined for the ATTACH parameters.

SYSMSG process headers

The following tables show the formats of the SYSMSG process headers: the SYSERROR FM header and the SYSSTAT FM header.

SYSERROR FM header format

The following table shows the format of the SYSERROR FM header.

Table 110. SYSERROR FM header format

Byte	Bits	Name	Contents
0		FMHL	
1	0	FMHC	
	1-7	FMHT	B'0000110'
2-3		FMH6CMD	X'0404'
4		FMH6MOD	
	0	FMH6LNSZ	B'0': 1-byte parameter length field
	1-7		Reserved
5		FMH6FXCT	Length of fixed-length parameters (X'00')
6-m		ATTPDN	RDPN field supplied with an input message resulting in this SYSMSG

Table 110. SYSEERROR FM header format (continued)

Byte	Bits	Name	Contents
m+1-n		ATTPRN	RPRN field supplied with an input message resulting in this SYSMMSG

SYSSTAT FM header

The following table shows the format of the SYSSTAT FM header.

Table 111. SYSSTAT FM header

Byte	Bits	Name	Content
0		FMHL	Length
1	0	FMHC	
	1-7	FMHT	B'0000110'
2-3		FMH6CMD	X'0402'
4		FMH6MOD	
	0	FMH6LNSZ	B'0' 1-Byte length fields
	1-7	Reserved	
5		FMH6FXCT	Length of fixed-length parameters

Chapter 34. Examples using ISC edit ATTACH parameters

The following topics provide examples of using ISC edit ATTACH parameters.

ATTACH and SCHEDULER parameters with ISC edit

The following series of figures illustrate the use of the ATTACH parameters when using ISC edit for IMS input and output messages.

ISC edit is used for normal transactions, commands, and message switches between LTERMs when one of the following occurs:

- An ATTACH FM header is received while in a between-brackets state that does not specify a destination process name (DPN).
- The DPN names IMS ISC edit as defined during IMS system definition.
- No ATTACH FM header is received either when the attach manager is in the between-brackets reset state or when the active process is ISC edit.
- The DPN=ISCEDT.




Figure 70. Example of message switch from another subsystem terminal to IMS terminal

Note: The IMS LTERM name for terminal T2 can be supplied either as the input primary resource name (PRN) or as the first field in the data.




Figure 71. Example of another subsystem accessing an IMS application program with a reply back to the subsystem terminal

Note: The IMS transaction code can be supplied either as the input PRN or within the data. The input RDPN and RPRN parameters become the DPN and PRN parameters respectively for any resulting reply from the application program when not processed by MFS.

In the following figure, messages are routed from IMS Terminal 1 to the ISC edit process through either a message switch or an MPP or Fast Path message routing application. The ISC edit process uses SCHEDULER parameters to send the message to a transaction on another subsystem. The other subsystem then sends the message back to IMS and the ISC edit process using similar SCHEDULER parameters. The ISC edit process directs the message back to Terminal 1.




Figure 72. Example of message routing with ISC edit

Note: The IMS LTERM name can be supplied either as the input SCDPRN or the first data field.

ATTACH parameters with the IMS SYSMMSG process

The system message process (SYSMMSG) is indicated through a special SNA process name in the ATTACH header.

Input system messages for IMS ISC sessions are logged and routed directly to the IMS master terminal operator when no PRN is supplied using the input ATTACH parameters. If a PRN value is supplied for IMS ISC sessions, it becomes the IMS input message destination transaction code or LTERM name (message switch), and the input system message is passed to ISC edit. SYSMMSG does not provide access to the IMS command processor.

A reply can result from the input SYSMMSG if a PRN value supplied is for an IMS transaction. This reply is sent just as is any other (non-SYSMSG) asynchronous output message.

IMS requests attachment of the SYSMMSG process to send IMS broadcast and other system messages that are not directly solicited by an input IMS command.

Exception: When IMS detects an error condition during an input IMS response mode or conversational transaction, an exception response (selective receiver ERP) to the input transaction occurs and the system message is sent as an ERP message.

The following figures illustrate the use of the ATTACH parameters when attaching the SYSMMSG process for IMS input and output messages.




Figure 73. Example of SYSMMSG without PRN from other subsystem terminal to IMS A

Note: The IMS LTERM name for the master terminal operator is assumed when no input ATTACH PRN parameter is supplied.

In the following figure, the input message from Terminal 1 is routed from a non-IMS subsystem to the IMS SYSMMSG process in IMS using ATTACH parameters. The IMS SYSMMSG process passes the message to the IMS edit process, which sends the data to tran1 in the IMS MPP region. After tran1 processes, the reply is passed to the IMS edit process again, which uses SCHEDULER parameters to send the reply back to Process 1 and then Terminal 1 on the other subsystem.




Figure 74. Example of SYSMSG with PRN from another subsystem terminal to IMS

Note: The PRN parameter becomes the IMS destination. The input RDPN and RPRN parameters become the output DPN and PRN parameters respectively for any resulting reply from the MPP.

ATTACH and SCHEDULER parameters with IMS MFS

The following series of figures illustrate the use of the ATTACH and SCHEDULER parameters when using IMS MFS for IMS input and output messages.




Figure 75. Example of message switch from other subsystem terminal to IMS terminal

Note: The IMS LTERM name for terminal T2 can be supplied as the input primary resource name (PRN), as a data field, or as an MFS-defined literal. RPRN and RDPN, if supplied, can be optionally included in the data presented to the IMS terminal by MFS. The DPN parameter must be supplied in the input ATTACH FM header to invoke MFS.

In the following figure, a transaction message originates from Terminal 1 and is sent to Process P1 on a non-IMS subsystem. The ATTACH parameters are then used to send the message to the MFS process P2 on the IMS subsystem. The MFS process P2 sends the message to Tran2 in the MPP region. The reply from the MPP region is sent back to the MFS process, which uses the ATTACH parameters to send the reply to the other subsystem. The other subsystem allocates Terminal 1 as a resource of Transaction Tran1 and the message is sent back to the terminal.




Figure 76. Example of other subsystem terminal accesses IMS application program with reply

Note: The IMS transaction code can be supplied as the input PRN, as a data field (first field for basic edit), or as an MFS-defined literal. The DPN parameter must be supplied to invoke IMS MFS DPM. If a reply inserted by the MPP is processed by MFS, the DPM process might create or override any of the output ATTACH parameters. The input RDPN and RPRN parameters become the output DPN and PRN parameters respectively for any resulting reply from the application program if they are not overridden by the DPM process.

In the following figure, a transaction message originates from Terminal 1 and is sent to Process P1 on a non-IMS subsystem. The ATTACH parameters are then used to send the message to the MFS DPM process P2 on the IMS subsystem. The MFS DPM process P2 sends the message data to Tran2 in the MPP region. The reply from the MPP region is sent back to the MFS DPM process, which uses the ATTACH parameters to send the reply to the non-IMS subsystem as Tran3, which is then stored for later retrieval by Terminal 1.




Figure 77. Example of IMS an MPP accessed from other subsystem terminal with reply to a temporary storage file of the other subsystem

Note: The IMS transaction code can be supplied as the input PRN, as a data field, or as an MFS-defined literal. The DPN parameter must be supplied to invoke IMS MFS DPM. If a reply is inserted by the MPP, the DPM process might override or create any of the output ATTACH parameters. The input RDPN and RPRN parameters become the output DPN and PRN parameters respectively for any resulting reply from the application program if they are not overridden by the DPM process.




Figure 78. Example of a message switch from an IMS terminal to a terminal on another subsystem

MFS DPM is necessary to create the required output SCHEDULER parameters. The SCDDPN and SCDPRN can be entered as data from the terminal or application program. If a reply is returned to another IMS terminal or application program, the default SCDRPRN parameter can be overridden within the IMS output message's SCHEDULER FM header to specify the suggested IMS terminal or application to receive that reply. The reply is returned with this value in the SCDPRN field.

In the following figure, messages are routed from IMS Terminal 1 to the IMS MFS process either through a message switch or through an MPP or Fast Path message routing application. The MFS process P2 sends the message with SCHEDULER parameters to a transaction on another subsystem. The other subsystem then sends the reply message back using similar SCHEDULER parameters to IMS and the MFS process P2. The MFS edit process directs the message back to Terminal 1.




Figure 79. Example of message routing with IMS MFS

Note: The IMS LTERM name can be supplied as the input SCDPRN, as the first data field, or as an MFS-defined literal. The SCDDPN parameter must be supplied to invoke IMS MFS. MFS can override or create any of the output SCHEDULER parameters.

Chapter 35. How IMS and CICS use the ISC interface

When designing and implementing an ISC network that contains both IMS and CICS nodes you should be aware of the issues involved.

You should also be familiar with the basic system design concepts of both IMS and CICS running in a VTAM environment. The CICS information provided here is for planning purposes.

The CICS information provided here is for planning purposes and should be verified in the CICS documentation or with a CICS system programmer.

Related reading: For more information on the design and coding of a CICS application that communicates with IMS through LU 6.1 ISC, see *CICS Transaction Server for z/OS CICS Intercommunication Guide*.

Functions available to the ISC session

The functions available to an ISC session between IMS and CICS differ depending on whether the processing is synchronous or asynchronous.

Related concepts:

“Relationship of ISC and IMS execution modes” on page 475

Overview of CICS synchronous and asynchronous processing for ISC

CICS implements ISC using its command level (EXEC) application programming interfaces (APIs). The CICS APIs can be synchronous or asynchronous.

The synchronous API uses the SEND and RECEIVE EXEC commands. In *synchronous mode* in an ISC session, the session is held between the moment a request is entered and the moment a reply is returned. Thus, a direct correlation can be made between the input request and the output reply.

The asynchronous API uses the START and RETRIEVE EXEC commands. In *asynchronous mode* in an ISC session, no correlation can be made between an input request and an output reply. No assumptions can be made as to the timing of the output reply or the availability of any other output for that session.

When used within the context of a CICS application program, the EXEC commands create the data flow control protocols for synchronous or asynchronous processing that are associated with messages on the session.

The CICS implementation is quite different from IMS, in which the data flow control protocols sent with a message are created by the system and based on:

- Whether the attributes (such as recoverability and segmentation) are defined for the message on the TRANSACT, TERMINAL, and SUBPOOL macro statements during IMS system definition or on an ETO logon descriptor
- Whether the characteristics are defined in the bind parameters
- Whether the message is a reply that is returned on the same session as an associated request or on a different session than the associated request

- Whether the message is unsolicited asynchronous output

When an asynchronous message is sent from CICS to IMS, the sending transaction might terminate. The returned reply can be processed by a newly initiated CICS transaction according to parameters in the function management headers returned on the message by IMS.

The following figure illustrates the concept of synchronous SEND and RECEIVE processing.




Figure 80. Synchronous processing with SEND and RECEIVE

The following figure illustrates the concept of asynchronous START and RETRIEVE processing.




Figure 81. Asynchronous processing with START and RETRIEVE

Related concepts:

 CICS: CICS-to-IMS applications

Functions available to an ISC VTAM session

ISC sessions with CICS on VTAM connections support both the synchronous and asynchronous CICS APIs.

This topic identifies synchronous and asynchronous command sequences available to a CICS application on ISC VTAM sessions.

The support for front-end/back-end system utilization that the IMS Front-End Switch exit routine provides uses ISC in asynchronous mode.

CICS SEND and RECEIVE commands create the predefined SNA protocols associated with synchronous processing when SEND is used with the INVITE option. SEND INVITE causes an SNA ATTACH function management header and change-direction to be issued with the message. The message is sent from CICS to the remote IMS subsystem for processing. When the message reply is returned, it is processed by a RECEIVE command issued within the same CICS transaction that issued the first message.

Exception: Restarted transactions are handled differently.

When the SEND command is used with the LAST option, it causes the message to be sent on the session with an ATTACH function management header and carries both begin-bracket and end-bracket (BB/EB) status. IMS treats this special case of synchronous processing in the same way as an IMS asynchronous transaction is treated. If IMS generates an output reply as a result of an input generated with SEND LAST, the returned reply can be processed by a newly initiated CICS transaction. The IMS output reply is also sent with ATTACH BB/EB. CICS uses RECEIVE to process that reply.

The START and RETRIEVE commands create the SNA protocols associated with asynchronous processing. A START command causes an SNA ATTACH function management header and a concatenated SCHEDULER function management header to be issued with the message. Messages sent to IMS with the START command can be only one chain in length and are sent with BB/EB. CICS uses the RETRIEVE command to obtain asynchronous messages from the session for CICS transaction processing. The other SNA data flow control protocols issued with an asynchronous message depend upon the type of message issued.

The following series of tables summarize the functions available to an ISC VTAM session between IMS and CICS. You must know whether the flow on a session is to be synchronous or asynchronous in order to determine what functions are available to that session. See the notes following the tables for further explanations.

The tables do not take into consideration the recovery aspects of any of the listed combinations. A “Yes” indicates only that the approach is possible. However, a system designer should take into account the need for ease of recovery and restart.

Table 112. Functions available to an ISC VTAM session for CICS EXEC commands: CICS front end

Functions available	SEND(INVITE)/ RECEIVE (Synchronous)	SEND(LAST) (Synchronous)	START/ RETRIEVE ¹ (Asynchronous)
Nonresponse or nonconversational mode IMS input transaction	No	Yes	Yes
Nonresponse or nonconversational mode IMS output transaction	Yes ²	N/A	Yes
Response mode IMS transaction (including Fast Path)	Yes ³	No	No
Conversational mode nonlast IMS input transaction	yes	No ⁴	No
Conversational mode last IMS input transaction	No	No ⁴	Yes
IMS message switch	No	Yes	Yes
IMS operator command	Yes ⁵	Yes ⁶	Yes ⁶
Recoverable ⁷ IMS transaction	N/A	N/A	N/A
Nonrecoverable ⁷ IMS transaction	N/A	N/A	N/A
Single-segment input/output IMS transaction	Yes	Yes	Yes
Multi-segment input/output IMS transaction	Yes	Yes	No

Table 113. Functions available to an ISC VTAM session for CICS EXEC commands: IMS front end

Functions available	SEND(INVITE)/ RECEIVE (Synchronous)	SEND(LAST) (Synchronous)	START/ RETRIEVE ¹ (Asynchronous)
CICS transaction	No	No	Yes
CICS operator command (system transaction)	No	No	No

Table 114. Functions available to an ISC VTAM session for CICS EXEC commands: MFS and BMS support

Functions Available ⁸	SEND(INVITE)/ RECEIVE (Synchronous)	SEND(LAST) (Synchronous)	START/ RETRIEVE ¹ (Asynchronous)
Without IMS MFS	Yes	Yes	Yes
MFS without paging	Yes	Yes	Yes
MFS autopaged input	Yes ⁹	Yes ⁹	No
MFS autopaged output	Yes ^{10, 2}	N/A	No
MFS demand-paged output	Yes	N/A	No ¹¹
BMS support ⁸	N/A	N/A	N/A

Notes for CICS EXEC tables:

1. Single-chain messages only.
2. The nonresponse-mode reply to a transaction sent with SEND LAST is sent ATTACH BB/EB. CICS uses RECEIVE (without SEND) to obtain the message from the session. (See CICS to IMS using the SEND LAST EXEC command (Communications and Connections).)
3. IMS forces response mode regardless of system definition if a transaction is specified externally (using the function management header) as synchronous. See Relationship of ISC and IMS execution modes (Communications and Connections).
4. IMS conversational mode requires that IMS terminate the conversation by sending the last reply with EB. The exception to this is explained in "CICS versus IMS conversation mode" on page 634.
5. The replies that result from the IMS /DISPLAY, /FORMAT, and /RDISPLAY commands are sent asynchronously after IMS replies synchronously to the input command.
6. Supported only for /DISPLAY, /RDISPLAY, and /FORMAT.
7. A transaction is defined as recoverable or nonrecoverable during IMS system definition and is acceptable to IMS only when the sync point protocols and transaction definitions are consistent. See Relationship of ISC and IMS execution modes (Communications and Connections) and Keeping half sessions synchronized (Communications and Connections) for more information.
8. MFS maps data between the session and the IMS application program. BMS maps data at the request of the CICS application program. BMS does not interact with the session, but can be used to handle mapping of data to and from IMS. IMS's MFS and CICS's BMS do not communicate.

9. Both single- and multiple-message chains can be used, but only one chain can occur per SEND.
10. To avoid tying up the session, when processing synchronous autopaged output from IMS, CICS should read all pages of the IMS output before processing it.
11. Although demand paging is possible between IMS and CICS, it is not recommended, because it requires complex CICS application coding and a complex terminal user interface. A preferable approach is to create autopaged output from IMS to CICS and use CICS page retrieval in the local system for paging.

Related concepts:

“Keeping half sessions synchronized” on page 525

“CICS to IMS using the SEND LAST EXEC command” on page 614

Related reference:

 Front-End Switch exit routine (DFSFEJ0) (Exit Routines)

General flow of CICS EXEC commands within a CICS application

The design of a CICS application program using the CICS EXEC command level application programming interfaces is determined by whether the transaction being sent on the ISC session is to be processed using SEND/RECEIVE, SEND LAST, or START/RETRIEVE.

The sequence of EXEC commands issued is determined by whether the transaction is defined as recoverable or nonrecoverable. Program design can also depend on whether CICS is the front-end system (initiating a transaction) or back-end system (replying to an IMS transaction).

This topic presents an overview of synchronous and asynchronous transaction processing flow within CICS. Understanding the content and functions of the ATTACH and SCHEDULER FM headers is helpful in understanding this topic.

Related concepts:

 Functions available to the ISC session (Communications and Connections)

Related reference:

“Coding function management headers for CICS” on page 638

CICS to IMS using SEND/RECEIVE EXEC commands

In a system in which CICS is the front-end subsystem and IMS the back-end subsystem, the SEND/RECEIVE commands are used to process IMS response mode (including Fast Path) and conversational transactions, and IMS commands.

The general CICS application program flow for a synchronous message from CICS to IMS is shown in the following figure.




Figure 82. Application program flow using SEND/RECEIVE from CICS to IMS

In this example, the CICS application reads the input message from the terminal and establishes a CICS-to-IMS session using the ALLOCATE command.

The CICS application program builds the required ATTACH function management header (the only header required, because this is a synchronous application) by using the BUILD ATTACH EXEC command. The RPROCESS and RRESOURCE fields should be specified in the BUILD ATTACH to provide for restart. Issuing SEND causes the output message to be constructed. SEND INVITE causes change-direction (CD) to be appended to the output message.

The output message is not sent across the session until the next command is executed. The next EXEC command that is issued depends upon the definition of the transaction within IMS, because this command must both append the sync-point request to the message as required and send the message across the ISC session. If this transaction is defined to IMS as recoverable, the next command must be SYNCPOINT. SYNCPOINT causes the outbound message to be sent with an RQD2. If the transaction is defined as nonrecoverable, a SYNCPOINT or a RECEIVE command can be issued next. A RECEIVE command immediately

following a SEND command (without an intervening SYNCPOINT command) causes the outbound message to be issued with RQE1.

RECEIVE issued after either the SYNCPOINT or SEND commands reads the IMS reply. The reply from IMS is returned synchronously on the session (with the ATTACH FM header). If the transaction reply is to a response mode (including Fast Path) transaction, a command (except /DISPLAY, /RDISPLAY, and /FORMAT), or the last conversational reply, it is returned to CICS carrying RQD2 EB if recoverable, or RQD1 EB if nonrecoverable. If the transaction is a nonlast conversational reply, it is returned to CICS carrying RQE2 and CD (because the conversation continues).

If the transaction is defined within IMS as a nonrecoverable-inquiry, the CONVERSE command can be used. The CONVERSE command acts as if a SEND command were issued, immediately followed by a RECEIVE command. Transactions sent with a CONVERSE command carry BB/CD and a request for an exception response (RQE1). If transactions other than those defined as nonrecoverable are sent using CONVERSE, an error results.

After issuing RECEIVE, the application program must save and check the values in the EXEC interface block (EIB) before performing any additional processing. If the transaction reply is received successfully as determined by checking the EIB, the application now issues a sync point (SYNCPOINT). Issuing this sync point causes a DR2 response to be returned to the IMS transaction reply. If CICS wants to perform any application processing based upon the contents of the IMS reply message, including sending an output message to the terminal, that processing is performed before the sync point is issued. This ensures that CICS resources and IMS resources are committed in synchronism.

If the transaction completes successfully, the application program must free the session using either a FREE or RETURN EXEC command.

Related concepts:

“Recovery and restart concepts” on page 645

Related reference:

“Coding function management headers for CICS” on page 638

CICS to IMS using the SEND LAST EXEC command

In a system in which CICS is the front-end subsystem and IMS the back-end subsystem, IMS supports the use of the CICS synchronous API to allow access to IMS asynchronous transactions. SEND LAST is used to process IMS nonresponse mode and nonconversational transactions, message switches, and the /DISPLAY, /RDISPLAY and /FORMAT commands.

The general CICS application program flow for CICS to IMS using SEND LAST is shown in the following figure.




Figure 83. Application program flow using *SEND LAST* from CICS to IMS

In this example, the CICS application reads the input message from the terminal and establishes a CICS-to-IMS session using the *ALLOCATE* command.

The CICS application program builds the required *ATTACH* function management header (the only header required, because this is a synchronous application) by using the *BUILD ATTACH EXEC* command. The *RPROCESS* and *RRESOURCE* fields should be specified in the *BUILD ATTACH* to identify the terminal and transaction to be used to process the reply. Issuing a *SEND* command causes the output message to be constructed. A *SEND LAST* command causes end-bracket to be appended to the message.

The output message is not sent on the session until the next command is executed. The next *EXEC* command that is issued depends upon the definition of the transaction within IMS, because this command both appends the sync-point request to the message and issues the message on the ISC session. If this transaction is defined to IMS as recoverable, the next command must be *SYNCPOINT*. *SYNCPOINT* causes the outbound message to be issued with an *RQD2*. If the transaction is defined as nonrecoverable, either *SYNCPOINT*, *FREE*, or *RETURN* can be issued next. *FREE* or *RETURN* issued immediately after the *SEND* and without an intervening *SYNCPOINT* causes the outbound message to be issued with *RQE1*.

If IMS generates a reply message as a result of receiving this transaction, that reply from IMS is returned with the *ATTACH FM* header. In order to get the IMS reply message from the session, CICS must issue a *RECEIVE* (because the reply message carries only *ATTACH*).

Related reference:

“Coding function management headers for CICS” on page 638

IMS to CICS using the RECEIVE EXEC command

When IMS has a reply to send to CICS that results from a message sent by CICS with SEND LAST (ATTACH BB/EB), IMS sends that message using the ATTACH function management header and BB/EB. In order to receive the message from the session, CICS must initiate a new transaction (or a new instance of the original transaction) that uses the RECEIVE EXEC command for this purpose.

The general CICS application program flow for this is shown in the following figure.




Figure 84. Application program flow using RECEIVE from IMS to CICS

The reply from IMS is returned synchronously on the session (with the ATTACH FM header) and carries BB/EB. A recoverable reply carries RQD2; a nonrecoverable reply carries RQD1. The RPROCESS (RDPN) and RRESOURCE (RPRN) fields sent to IMS are automatically wrapped by IMS into the DPN and PRN fields of the outbound ATTACH FMH.

CICS examines the returned reply and uses the DPN, or the first four characters of the data, to determine the transaction that can process the returned reply. RECEIVE is used to obtain the IMS reply from the session. The PRN field is passed to the initiated transaction, indicating the terminal to which the returned reply is to be sent.

After issuing RECEIVE, the application program must save and check the values in the EXEC interface block (EIB) before performing any additional processing.

CICS must schedule an asynchronous transaction to send the returned reply to the terminal by issuing a START command. The transaction that is initiated is brought up, and owns the terminal to which the reply is to be sent. This asynchronous transaction uses the START interface to communicate with the terminal, and all processing within this transaction occurs asynchronously to the transaction that initiated it. If the transaction that sends output to the terminal is scheduled

successfully, the originating application now issues a sync point (SYNCPOINT), if required. Issuing this sync point causes an appropriate response to be returned to the IMS transaction reply. If the transaction is not scheduled successfully, issuing the sync point causes an appropriate exception response to be returned to the IMS transaction reply. If CICS wants to perform any application processing based upon the contents of the IMS reply message, that processing is performed before the sync point is issued. This ensures that CICS resources and IMS resources are committed in synchronism.

If the processing completes successfully, the application program must free the session using either a FREE or RETURN EXEC command.

Coding asynchronous messages

When receiving asynchronous messages (sent with ATTACH and SCHEDULER function management headers) from IMS, CICS invokes a CICS-supplied *mirror transaction* to obtain these incoming messages from the session.

The mirror transaction is CICS's name for the SCHEDULER process. This mirror transaction examines the incoming data stream and schedules (using START) the transaction that is to be initiated as a result of the incoming DPN or the first four characters of the data. The receiving CICS transaction is assumed to be one of the following:

- ISC edit (ISCE), the default DPN set by IMS
- The transaction whose identifier was placed in the DPN field of the ATTACH FM header by IMS's wrapping of the incoming RDPN (the RTRANSID specified on the CICS START)
- A transaction whose identifier has been supplied by an MFS MOD

CICS uses the PRN in the incoming message as the TERMID parameter of the START command so that the transaction is initiated and owns the CICS terminal (if applicable) to which this message should be written.

A SYNCPOINT command is now issued to return the appropriate response to IMS and to complete the scheduling of the asynchronous transaction. Following the sync point, a RETURN command is issued to terminate the transaction and free the session. The response returned to IMS causes the output asynchronous message to be dequeued from the IMS output queue.

CICS to IMS using the START/RETRIEVE EXEC commands

CICS uses START/RETRIEVE commands to process IMS nonresponse mode and nonconversational transactions, message switches, and the /DISPLAY, /RDISPLAY, and /FORMAT commands.

The general CICS application program flow for an asynchronous message from CICS to IMS is shown in the following figure.




Figure 85. Application program flow using START/RETRIEVE from CICS to IMS

After the input message is received from the terminal, the transaction is assembled using the START command. The values to be contained in the ATTACH and SCHEDULER function management header DPN, PRN, RDPN, and RPRN fields are specified as parameters on the START command. The DPN and PRN are used by IMS to determine the editing process to be initiated and the destination of the incoming message. The RDPN and RPRN are used to identify the CICS terminal and transaction to which the reply is to be returned. The START command must also specify NOCHECK if this is an IMS nonrecoverable transaction, and NOCHECK PROTECT if this is an IMS recoverable transaction. (PROTECT can optionally be used for nonrecoverable transactions.)

NOCHECK is a mandatory parameter for an ISC session with IMS. In a CICS-CICS ISC session, the use of START causes a return code reply to be sent to the initiating CICS from the receiving CICS. This return code indicates that the receiving CICS has scheduled the transaction to be executed as a result of the message it has received. This return code does not occur on an IMS-CICS session. NOCHECK informs the sending CICS that no such response is to be expected. START NOCHECK creates a message carrying BB/EB RQE1.

PROTECT, specified for recoverable transactions, delays the sending of the transaction on the session until CICS successfully takes a SYNCPOINT. Specifying START causes begin- and end-bracket (BB/EB) and the ATTACH and SCHEDULER function management headers to be appended to the outbound message. NOCHECK PROTECT requests a definite response (RQD2).

A SYNCPOINT command is now issued to append the appropriate sync point protocols and to send the message on the session. If this transaction elects to wait for a reply, it issues a RETRIEVE command. If this transaction does not choose to wait for a reply, it can terminate by issuing RETURN following either the START

command or the SYNCPOINT command. The CICS mirror transaction schedules a new instance of this transaction when a reply is received.

If the transaction chooses to wait for the reply, it can issue RETRIEVE with or without the WAIT parameter. Issuing RETRIEVE without the WAIT parameter causes CICS to check for any queued asynchronous input for this transaction and terminal. If such input is immediately available, it satisfies the RETRIEVE command. If no such input is immediately available, an appropriate indication is returned to the RETRIEVE command, and the application can then perform other processing or issue RETURN to terminate.

If RETRIEVE WAIT is issued, CICS does not return control to the application program until an asynchronous message is sent from IMS destined for this transaction and this terminal. Using WAIT causes the terminal be held until an IMS reply is received for it.

In both cases (RETRIEVE with or without WAIT), when a message is sent from CICS to IMS requesting asynchronous processing, no assumptions can be made by the originating application as to the timing of the output reply or the availability of any other output. That is, if unsolicited asynchronous output is pending for CICS, this output can be sent to CICS before the reply to this asynchronous transaction is returned. Therefore, requests and replies are not correlated within CICS in an asynchronous environment. Issuing RETRIEVE obtains any reply from IMS on any session that has the indicated transaction and terminal identification.

IMS to CICS using the RETRIEVE EXEC command

The following figure illustrates the general CICS application program flow in asynchronous mode from IMS to CICS.




Figure 86. Application program flow using RETRIEVE from IMS to CICS

Replies resulting from previous asynchronous input transactions and unsolicited asynchronous output from an IMS front-end to a CICS back-end are sent to CICS asynchronously (that is, with the ATTACH and SCHEDULER FM headers). When IMS is the front-end subsystem, asynchronous mode is the only flow supported.

The CICS mirror transaction obtains the input, analyzes the incoming message, and schedules (using START) the transaction that is to be initiated as a result of the incoming DPN or the first four characters of the data. If this transaction is sending output to a terminal, the mirror transaction issues the START with the TERMID parameter (from the PRN parameter of the incoming FMH), so that the asynchronous transaction that is scheduled owns the terminal to which output is to be sent.

If a reply to IMS is required, the CICS application must wrap the incoming RTERMID and RTRANSID into the TERMID and TRANSID fields to be used in a subsequent START. These fields indicate, respectively, the destination LTERM or transaction within IMS and the IMS editor or MFS MID that is to receive the message. The START command to send this reply to IMS is processed.

If no reply to IMS is required, RETURN can be issued to terminate the session.

Related concepts:

“CICS to IMS using the START/RETRIEVE EXEC commands” on page 617

Related reference:

“Coding function management headers for CICS” on page 638

Commands that should not be used on an ISC session

Do not use the WAIT SIGNAL or WAIT TERMINAL commands on an ISC session.

The WAIT SIGNAL command cannot be used in a CICS-IMS session.

The WAIT TERMINAL command is not normally used in a CICS-IMS ISC session.

Selecting appropriate CICS installation options for ISC

This topic describes only those system definition options and resource definition options that have unique considerations for an IMS-CICS ISC environment.

Some of the parameters coded during the CICS installation process must be compatible with some of the IMS system definition parameters coded on the IMS system definition macros. The topic “Defining compatible IMS and CICS nodes” on page 622 brings together the definitions provided here with the IMS system definition information provided in “Statically defining an ISC node to IMS” on page 496.

Coding CICS system definition options

The CICS standard pregenerated system as supplied on the distribution tape includes a pregenerated version for each of the CICS management modules required to support an Intersystem Communication environment that includes all of the required system definition options.

Related reading:

- For information on system definition, see *CICS Transaction Server for z/OS CICS Intercommunication Guide* for the specifications necessary to generate a CICS subsystem that participates in IMS-CICS Intersystem Communication.
- For a checklist of recovery options, see *CICS Transaction Server for z/OS Recovery and Restart Guide*.

Preparing CICS resource definition

CICS user-created tables and Resource Definition Online (RDO) describe the database and data communications environment, and the treatment of the elements in that environment.

They contain information about terminals, files, databases, programs, transactions, transient data destinations, and temporary storage data identifiers. They are created independently of system definition, but must be present for the system to be operational. The CICS System Log and Dynamic Log, which are optional in some CICS environments, are required for IMS-CICS ISC.

Related reading:

- For more information on table preparation and RDO options, see the *CICS Transaction Server for z/OS CICS Resource Definition Guide*.
- For a summary of the parameters required for IMS-CICS Intersystem Communication, see *CICS Transaction Server for z/OS CICS Intercommunication Guide*.
- For a checklist of recovery-related options, see *CICS Transaction Server for z/OS Recovery and Restart Guide*.

Defining IMS-CICS LU 6.1 links

A CICS link to an IMS system requires a definition of the connection (or system) and a separate definition of each of the sessions.

Resource Definition Online or Macro-Level Resource Definition can be used to define an ISC link to IMS. IMS-CICS ISC links are implemented with LU 6.1 protocol.

The following table shows the RDO form of definition for individual LU 6.1 sessions, and shows the macro form for how the operands are related.

The TRMTYPE, TRMIDNT, SYSIDNT, NETNAMQ, and SESTYPE operands must be coded for each session that you define. The remaining operands of TYPE=TERMINAL can optionally be coded on the TYPE=SYSTEM macro to provide defaults for all the defined sessions. Also, the CONNECT, DATASTR, and RECFM operands of TYPE=SYSTEM can be coded for individual sessions, if required.

Table 115. Defining an LU 6.1 link with individual sessions

RDO definition	Macro-level definition
DEFINE CONNECTION(sysidnt) GROUP(groupname) NETNAME(name) ACCESSMETHOD(VTAM) PROTOCOL(LU61) DATASTREAM(USER 3270 SCS STRFIELD LMS) RECORDFORMAT(U VB) AUTOCONNECT(NO YES) SECURITYNAME(name) INSERVICE(YES)	DFHTCT TYPE=SYSTEM ,SYSIDNT=sysidnt) ,NETNAME=name ,ACCMETH=VTAM) ,DATASTR=({USER 3270 SCS STRFIELD LMS}) ,RECFM={U VB} [,CONNECT=AUTO ALL] ,XSNAME=name

Table 115. Defining an LU 6.1 link with individual sessions (continued)

RDO definition	Macro-level definition
Each individual session is then defined as follows:	Each individual session is then defined as follows:
<pre> DEFINE SESSIONS(csdname) GROUP(groupname) SESSNAME(name) CONNECTION(sysidnt) NETNAMEQ(name) PROTOCOL(LU61) SENDCOUNT(0 1) RECEIVECOUNT(1 0) SENDSIZE(size) RECEIVESIZE(size) BUILDCHAIN(N Y) OPERID(operator-id) OPERPRIORITY(number) OPERRSL(number) OPERSECURITY(number) IOAREALEN(value) SESSPRIORITY(number) </pre>	<pre> DFHTCT TYPE=TERMINAL ,TERMIDNT=name ,SYSIDNT=sysidnt ,NETNAMQ=name ,TRMTYPE=LUTYPE6 ,SESTYPE= SEND RECEIVE ,BUFFER=size ,RUSIZE=size ,CHNASSY={NO YES} ,OPERID=operator-id ,OPERPRI=number ,OPERRSL=number ,OPERSEC=number ,TIOAL=value ,TRMPRTY=number ,TRMSTAT=TRANSCIVE </pre>

Defining compatible IMS and CICS nodes

To define IMS-CICS ISC links, you should understand the relationship between the way remote systems and sessions are defined in CICS and the way they are defined in IMS.

Resource Definition Online (RDO) terms are used in the following explanation of the compatibility requirements. For the equivalent macro-level operands, see “Defining IMS-CICS LU 6.1 links” on page 621.

System names

The network name of the CICS system is specified in the APPLID operand of the DFHSIT macro.

The name can be provided as an override during CICS startup or in the APPLID operand of the DFHTCT TYPE=INITIAL macro. This name must be specified in the NAME operand of the TERMINAL macro or on the ETO logon descriptor that defines the CICS system. This name must be known to the IMS master terminal operator as it is required to be specified within those IMS master terminal operator commands (for example, /OPNDST) that reference sessions between IMS and CICS.

The network name of the IMS system is specified in the APPLID operand of the IMS COMM macro. You must specify this name in the NETNAME operand of the DEFINE CONNECTION command that defines the IMS system.

Number of sessions

For statically defined terminals in IMS, the number of parallel sessions that are required between the CICS and IMS system must be specified in the SESSION operand of the IMS TERMINAL macro.

Each session is then represented by a SUBPOOL entry in the IMS VTAMPOOL. In CICS, each of these sessions is represented by an individual session definition.

Session names

Each CICS-IMS session is uniquely identified by a *session-qualifier pair*, formed from both the CICS and IMS session names.

The CICS session name is specified in the SESSNAME operand of the DEFINE SESSIONS command. For sessions that are to be initiated by IMS, this name must correspond to the ID parameter of the IMS /OPNDST command for the session. For sessions initiated by CICS, the name is supplied on the CICS /OPNDST command and is saved by IMS.

The IMS session name is specified in the NAME operand of the IMS SUBPOOL macro. You must make the relationship between the session names explicit by coding this name in the NETNAMEQ operand of the corresponding DEFINE SESSIONS command.

Recommendation: For operational convenience, use the same CICS and IMS names for a session.

Other session parameters

The following additional operands of the DEFINE CONNECTION and DEFINE SESSIONS commands are also significant for CICS-IMS sessions.

SENDSIZE

- If CICS is the *primary* half session, specifies the maximum request unit (RU) that CICS sends to the remote IMS system. This value must be less than or equal to the value specified by the RECANY parameter in the IMS COMM macro.
- If CICS is the *secondary* half session, specifies the maximum RU that CICS receives from the remote IMS system. This value must be less than or equal to the value specified by the OUTBUF parameter in the IMS TERMINAL macro.

RECEIVESIZE

- If CICS is the *primary* half session, specifies the maximum request unit (RU) that CICS receives from the remote IMS system. This value must be less than or equal to the value specified by the OUTBUF parameter in the IMS TERMINAL macro.
- If CICS is the *secondary* half session, specifies the maximum RU that CICS sends to the remote IMS system. This value must be less than or equal to the value specified by the RECANY parameter in the IMS COMM macro.

BUILDCHAIN(N|Y)

Specifies whether multiple RU chains are to be assembled before being passed to the application program. If Y is specified, a complete chain is passed to the application program in response to each RECEIVE command, and the application must perform any required deblocking. If N is specified, a single RU is passed to the application program in response to each RECEIVE command.

Recommendation: BUILDCHAIN(Y) is recommended, even when IMS record mode (VLVB) is used, because the logical records produced as IMS output might not coincide with RU boundaries.

DATASTREAM(USER)

Specifies a data stream profile when CICS is communicating with IMS using the START command (asynchronous processing). CICS messages generated by the START command always cause IMS to interpret the data stream profile as input for component 1. This parameter is required.

The data stream profile for distributed transaction processing can be specified by the application program using the DATASTR option of the BUILD ATTACH command.

RECORDFORMAT(U|VB)

Specifies the type of chaining that CICS is to use for transmissions that are initiated by START commands (asynchronous processing) on a particular session.

Two types of data handling algorithms are supported between CICS and IMS:

- **Chained**

Messages are sent as SNA chains. The user can use private blocking and deblocking algorithms. This format corresponds to RECORDFORMAT(U). When IMS is the receiver, a chain of RUs is interpreted as the complete single-segment message.

- **Variable length variable blocked records (VLVB)**

Messages are sent in variable-length variable-blocked format with a halfword length field before each record. This format corresponds to RECORDFORMAT(VB). When IMS is the receiver, the input deblocked element is treated as the input segment.

The data stream format for distributed transaction processing can be specified by the application program using the RECFM option of the BUILD ATTACH command.

SENDCOUNT and RECEIVECOUNT

These operands are used to specify whether the session is a SEND session or a RECEIVE session. (In macro-level definition, this is specified in the SESTYPE=SEND|RECEIVE operand.)

A SEND session is one in which the local CICS is secondary and is the contention winner. It is specified by:

```
SENDCOUNT(1)  
RECEIVECOUNT(0)
```

A RECEIVE session is one in which the local CICS is primary and is the contention loser. It is specified by:

```
SENDCOUNT(0)  
RECEIVECOUNT(1)
```

Recommendation: SEND sessions are recommended for all CICS-IMS sessions.

You need not specify a SENDPFX or a RECEIVEPFX; the name of the session is taken from the SESSNAME operand.

The following table shows the relationship between the CICS Resource Definition Online and IMS definitions of an ISC link. Related operands are shown by numbers. If CICS is communicating with an XRF IMS, NETNAME(SYSIMS) should be the USERVAR or MNPS ACB associated with the XRF IMS.

Table 116. Defining compatible CICS and IMS nodes: RDO

CICS	IMS
DFHIST TYPE=CSECT ,SYSIDNT=CICL ,APPLID=SYSCICS 1	COMM APPLID=SYSIMS 2 RECANY= <i>mmm</i> + 22 7 EDTNAME=ISCEDT
DEFINE CONNECTION(IMSR) 3 GROUP(<i>group_name</i>) NETNAME(SYSIMS) 2 ACCESSMETHOD(VTAM) PROTOCOL(LU61) DATASTREAM(USER)	TYPE UNITYPE=LUTYPE6 4 TERMINAL NAME=SYSCICS 1 SESSION=2 COMPT1= COMPT2= OUTBUF= <i>nnn</i> 8
DEFINE SESSIONS(<i>csdname</i>) GROUP(<i>group_name</i>) PROTOCOL(LU61) 4 SESSNAME(IMS1) CONNECTION(IMSR) 3 NETNAMEQ(CIC1) 5 SENDCOUNT(1) RECEIVECOUNT(0) SENDSIZE(<i>mmm</i>) 7 RECEIVESIZE(<i>nnn</i>) 8 IOAREALEN(<i>nnn</i> ,16364)	VTAMPOOL SUBPOOL NAME=CIC1 5 NAME CICLT1 COMPT=1 NAME CICLT1A
DEFINE SESSIONS(<i>csdname</i>) GROUP(<i>group_name</i>) PROTOCOL(LU61) 4 SESSNAME(IMS2) CONNECTION(IMSR) 3 NETNAMEQ(CIC2) 9 SENDCOUNT(1) RECEIVECOUNT(0) SENDSIZE(<i>mmm</i>) 7 RECEIVESIZE(<i>nnn</i>) 8 IOAREALEN(<i>nnn</i> ,16364)	SUBPOOL NAME=CIC2 9 NAME CICLT2 COMPT=2

The following table shows the relationship between the CICS macro-level definitions and IMS definitions of an ISC link. Related operands are shown by numbers. NETNAME=SYSIMS should be the USERVAR or MNPS ACB associated with the IMS system if you are defining a CICS to XRF IMS session.

Table 117. Defining compatible CICS and IMS nodes: macro level

CICS	IMS
DFHIST TYPE=CSECT ,SYSIDNT=CICL ,APPLID=SYSCICS 1	COMM APPLID=SYSIMS 2 RECANY= <i>mmm</i> + 22 7 EDTNAME=ISCEDT
DFHTCT TYPE=SYSTEM ,ACCMETH=VTAM ,SYSIDNT=IMSR 3 ,NETNAME=SYSIMS 2	TYPE UNITYPE=LUTYPE6 4 TERMINAL NAME=SYSCICS 1 SESSION=2 COMPT1= COMPT2= OUTBUF= <i>nnn</i> 8
DFHTCT TYPE=TERMINAL ,TRMTYPE=LUTYPE6 4 ,TRMIDNT=IMS1 ,SYSIDNT=IMSR 3 ,NETNAMEQ=CIC1 5 ,SESTYPE=SEND ,BUFFER= <i>mmm</i> 7 ,RUSIZE= <i>nnn</i> 8 ,TIOAL=(<i>nnn</i> ,16364) ,DATASTR=USER	VTAMPOOL SUBPOOL NAME=CIC1 5 NAME CICLT1 COMPT=1 NAME CICLT1A SUBPOOL NAME=CIC2 9 NAME CICLT2 COMPT=2
DFHTCT TYPE=TERMINAL ,TRMTYPE=LUTYPE6 4 ,TRMIDNT=IMS2 ,SYSIDNT=IMSR 3 ,NETNAMEQ=CIC2 9 ,SESTYPE=SEND ,BUFFER= <i>mmm</i> 7 ,RUSIZE= <i>nnn</i> 8 ,TIOAL=(<i>nnn</i> ,16364) ,DATASTR=USER	

Related concepts:

“LTERM users (subpools) and components” on page 478

Defining multiple links to an IMS system

You can define more than one intersystem link between a CICS and an IMS system by defining two or more connections (systems), with their associated session definitions, having the same NETNAME but different SYSIDs.

Although all the system definitions resolve to the same NETNAME, and therefore to the same IMS system, using a SYSID name in CICS causes CICS to allocate a session from the link with the specified SYSIDNT.

Recommendation: Define up to three links (that is, groups of sessions) between a CICS and an IMS system, depending upon the application requirements of your installation:

- A group of sessions for CICS-initiated distributed transaction processing (synchronous processing). CICS applications that use the SEND/RECEIVE interface can use the SYSIDNT of this group to allocate a session to the remote system. The session is held (“busy”) until the conversation is terminated.
- A group of sessions for CICS-initiated asynchronous processing. CICS applications that use the START command can name the SYSIDNT of this group. CICS uses the first “nonbusy” session to ship the start request. IMS sends a positive response to CICS as soon as it has queued the start request, so that the

session is in use for a relatively short period. Consequently, the first session in the group shows the heaviest usage, and the frequency of usage decreases towards the last session in the group.

- A group of sessions for IMS-initiated asynchronous processing. This group is also useful as part of the solution to a performance problem that can arise with CICS-initiated asynchronous processing. An IMS transaction that is initiated as a result of a START command shipped on a particular session uses the same session to ship its “reply” START command to CICS. For the reasons given in (2), the CICS START command is probably shipped on the busiest session, and, because CICS is the contention winner, the replies from IMS can create a backlog while waiting for a chance to use the session. However, facilities exist in IMS for a transaction to alter its default output session, and a switch to a session in this third group can reduce backlog problems.

The following series of tables provide an example of defining three groups of sessions on an IMS node. The tables contain examples of defining the session groups in both CICS Resource Definition Online (RDO) format and CICS macro format.

Table 118. CICS system initialization parameters

RDO definition	Macro-level definition
DFHSIT TYPE=CSECT ,SYSIDNT=CICL ,APPLID=SYSCICS	DFHSIT TYPE=CSECT ,SYSIDNT=CICL ,APPLID=SYSCICS

Table 119. CICS-initiated distributed transaction processing

RDO definition	Macro-level definition
DEFINE CONNECTION(IMSA) ACCESSMETHOD(VTAM) NETNAME(SYSIMS)	DFHTCT TYPE=SYSTEM ,ACCMETH=VTAM ,SYSIDNT=IMSA ,NETNAME=SYSIMS
DEFINE SESSIONS(csdname) PROTOCOL(LU61) SESSNAME(IMS1) CONNECTION(IMSA) NETNAMEQ(DTP1)	DFHTCT TYPE=TERMINAL ,TRMTYPE=LUTYPE6 ,TRMIDNT=IMSA ,SYSIDNT=IMSA ,NETNAMQ=DTP1
DEFINE SESSIONS(csdname)	DFHTCT TYPE=TERMINAL
⋮	⋮

Table 120. CICS-initiated asynchronous processing

RDO definition	Macro-level definition
<pre>DEFINE CONNECTION(IMSB) ACCESSMETHOD(VTAM) NETNAME(SYSIMS)</pre>	<pre>DFHTCT TYPE=SYSTEM ,ACCMETH=VTAM ,SYSIDNT=IMSB ,NETNAME=SYSIMS</pre>
<pre>DEFINE SESSIONS(csdname) PROTOCOL(LU61) SESSNAME(IMS1) CONNECTION(IMSB) NETNAMEQ(ASP1)</pre>	<pre>DFHTCT TYPE=TERMINAL ,TRMTYPE=LUTYPE6 ,TRMIDNT=IMSA ,SYSIDNT=IMSA ,NETNAMQ=DTP1</pre>
<pre>DEFINE SESSIONS(csdname) : :</pre>	<pre>DFHTCT TYPE=TERMINAL : :</pre>

Table 121. IMS-initiated asynchronous processing

RDO definition	Macro-level definition
<pre>DEFINE CONNECTION(IMSC) ACCESSMETHOD(VTAM) NETNAME(SYSIMS)</pre>	<pre>DFHTCT TYPE=SYSTEM ,ACCMETH=VTAM ,SYSIDNT=IMSC ,NETNAME=SYSIMS</pre>
<pre>DEFINE SESSIONS(csdname) PROTOCOL(LU61) SESSNAME(IMS1) CONNECTION(IMSC) NETNAMEQ(IST1)</pre>	<pre>DFHTCT TYPE=TERMINAL ,TRMTYPE=LUTYPE6 ,TRMIDNT=IMS1 ,SYSIDNT=IMSC ,NETNAMQ=IST1</pre>
<pre>DEFINE SESSIONS(csdname) : :</pre>	<pre>DFHTCT TYPE=TERMINAL : :</pre>

Defining CICS transactions for IMS-CICS ISC

This topic describes the unique considerations for transactions that participate in IMS-CICS ISC.

Related reading: For complete details on the definition of CICS transactions, see the *CICS Transaction Server for z/OS CICS Resource Definition Guide*.

Defining CICS backout in-doubt processing

During the period between the sending of the syncpoint request to IMS and the receipt of the positive response, CICS does not know whether the remote system has committed. This period is known as the indoubt period.

CICS processing during the indoubt period is controlled by the `DEFINE TRANSACTION IN-DOUBT` parameter (or the `DFHPCT DTB=` parameter).

`DEFINE TRANSACTION IN-DOUBT (WAIT)` or `DTB=(YES,WAIT)` must be specified in order to ensure consistency between the IMS and CICS subsystems within an ISC network.

Defining CICS transactions for asynchronous communication to IMS

A DEFINE TRANSACTION REMOTENAME (or DFHPCT TYPE=REMOTE RMTNAME) is needed for every IMS edit name or MID name that is to be referenced by a transaction created as the result of a CICS START command.

The 4-character CICS transaction ID used on the START command is converted to the name of the IMS editor or to an MFS MID name, each of which can be up to 8 characters.

Initiating and allocating a session from CICS

CICS can initiate a session in one of several ways.

- The session can be initiated explicitly by use of AUTOCONNECT YES on the DEFINE CONNECTION (or CONNECT=AUTO on the DFHTCT TYPE=TERMINAL) macro. When CICS is initiated, it attempts to establish all sessions for which AUTOCONNECT is specified. In order for session initiation to occur, IMS must be active when CICS is initiated.
- The master terminal operator can initiate a session by entering the command:
`CEMT SET TERMINAL(tttt) ACQUIRED|COLDACQ`

where *tttt* is the SESSNAME on the DEFINE SESSIONS or the TRMIDNT on the DFHTCT TYPE=TERMINAL.

If ACQUIRED is specified, normal resynchronization with the IMS is attempted.

If COLDACQ is specified, no resynchronization is performed.

The status of the CICS half session is typed over in the display area of the CEMT INQUIRE|SET TERMINAL command. If the session is initiated successfully, the status is changed from REL to ACQ. If the attempt to initiate the session is unsuccessful, an error message is written to the transient data destination CSMT.

- A session can be initiated implicitly by an application program using the ALLOCATE command. The ALLOCATE command is used on the SEND/RECEIVE interface only.

Recommendation: Use the ALLOCATE SYSID form of this command, because it allows CICS to select an available session.

If a session is not immediately available, control is returned to the application program in the following situations:

- A HANDLE CONDITION for this condition is issued by the application program.
- NOQUEUE is specified on the ALLOCATE command.

Otherwise, the command is queued.

The following conditions cause a session to be “not immediately available”:

- All sessions to the specified system (or the specified session) are in use.
 - The only available sessions are not bound.
 - The only available sessions are contention losers.
- A session can be initiated by CICS automatic task initiation (ATI).

The bind parameters for CICS are built from a hard-coded model. The RECEIVESIZE and SENDSIZE parameters on the DEFINE CONNECTION/SESSION (or the RUSIZE and BUFFER parameters on the DFHTCT TYPE=TERMINAL/SYSTEM) are used to determine the primary and secondary RU sizes.

Related reading: For more information on the ALLOCATE command, and subsequent processing, see *CICS Transaction Server for z/OS CICS Intercommunication Guide*.

Other ways of initiating a session

Sessions that are not initiated by CICS can be initiated either by IMS, by the VTAM operator, or Tivoli NetView for z/OS.

Terminating a session from CICS

Sessions can be terminated by CICS only by using master terminal operator commands.

The master terminal operator can use the CEMT command to release the session by entering:

```
CEMT SET TERMINAL(tttt) [RELEASED|OUTSERVICE]
```

In the example, *tttt* is the SESSNAME on the DEFINE SESSIONS or the TRMIDNT on the DFHTCT TYPE=TERMINAL.

If the master terminal operator specifies RELEASED, the session terminates when any active transaction completes and the session is in a between-brackets state.

If the master terminal operator specifies OUTSERVICE, the session also terminates when any active transactions complete and the session is in a between-brackets state. However, in this case, the session is taken out of service and cannot be used until the master terminal operator places it back into service.

If RELEASED is specified, it initiates an orderly termination between IMS and CICS. Although, from the CICS view, the session might appear to be in warm-start state, it is actually in cold-start state as a result of the SBI/BIS flow.

If OUTSERVICE is specified, the session is left in a warm-start state; that is, CICS requests resynchronization upon reinitiation. In order to initiate a session in cold-start mode, the CICS master terminal operator must specify:

```
CEMT SET TERMINAL (tttt) COLDACQ.
```

You can specify RELEASED and OUTSERVICE together.

Recommendation: Although orderly session termination can occur as a result of the CICS application program issuing EXEC CICS DISCONNECT, this approach is not recommended in a normal application. However, an operator-control type of application can be written to use this function.

Any messages relative to the session's termination are sent to transient data destination CSMT.

Related tasks:

"Initiating and allocating a session from CICS" on page 629

Related reference:

"Symmetrical session shutdown for LU 6.1 (SBI and BIS)" on page 558

Designing CICS applications for ISC

CICS differs from IMS in that many of the flows that must occur on an ISC session occur under the control of the CICS application program.

The CICS application must do the following:

- Build outbound ATTACH function management headers
- Supply fields for SCHEDULER function management headers (if required)
- Examine inbound ATTACH function management headers and process their contents
- Issue sync points at appropriate intervals during the application program's processing.

The SNA bracket, SEND/RECEIVE, and sync-point protocols associated with CICS-generated messages depend upon the EXEC commands (The EXEC commands are SEND/RECEIVE for synchronous and START/RETRIEVE for asynchronous.) used by the application program, and when the sync points are issued during the application flow.

Information is inserted into the ATTACH and SCHEDULER function management header DPN, RDPN, PRN, and RPRN fields for message editing and routing by IMS system code, and can be overridden only by using the IMS Message Format Service (MFS). For CICS, insertion of information into the ATTACH and SCHEDULER function management headers is performed under the control of the CICS application.

Application-related concepts

This topic describes some application-related concepts to establish a common frame of reference for both CICS and IMS users.

Subsystem design: direct-control versus queued

CICS is a direct-control subsystem, while IMS is a queued subsystem. That is, CICS accepts data entered from a terminal, and as a result, invokes the appropriate application program to process that data.

The terminal and other system resources are owned by the invoked application until the application task completes its processing. Information that results from processing is held in main storage rather than being queued. The implication for ISC is that the CICS application program is directly involved in generating the appropriate SNA data flow control, sync point, and response protocols, and in controlling most system services.

In contrast, IMS is a queued subsystem. In this case, all input and output transactions and message switches are queued by the IMS control region on behalf of the related IMS applications and terminals. Thus, the input and output of a message to or from the terminal are asynchronous from the processing of the message. (However, the processing might appear to be synchronous to the terminal because of the way in which the message is defined to IMS; for example, a response-mode message or a conversation might appear to be synchronous.) The implication for ISC is that the SNA protocols and many system services are handled under the control of IMS system code; the application program does not need to provide them.

Synchronous and asynchronous processing

Message transmission can be synchronous or asynchronous between the terminal entering the message and the receiving subsystem. Messages can also be either synchronous or asynchronous with respect to the ISC session.

From the point of view of the CICS application, using the CICS SEND/RECEIVE EXEC commands produces synchronous processing on the session, while the CICS START/RETRIEVE EXEC commands result in asynchronous processing. When CICS is the front-end subsystem, both processing types are supported. When IMS is the front-end subsystem, only asynchronous processing can occur, unless the special support for front-end/back-end system utilization provided by IMS Front-End Switch exit routine is used.

The following table shows how the CICS START/RETRIEVE and SEND/RECEIVE EXEC commands are supported from a CICS point of view.

Table 122. The SEND/RECEIVE and START/RETRIEVE commands

Message type	CICS to IMS	IMS to CICS
SEND/RECEIVE (synchronous)	YES	YES
START/RETRIEVE (asynchronous)	YES	YES

The following table shows the differences in the CICS application program interface when using SEND/RECEIVE or START/RETRIEVE.

Table 123. CICS API for SEND/RECEIVE and START/RETRIEVE

Attribute	SEND/RECEIVE	START/RETRIEVE
Mode	Synchronous with (SEND INVITE) or without (SEND LAST) reply	Asynchronous
FMH	ATTACH	ATTACH SCHEDULER
Name length	Supports 8-byte names	Supports 4-byte names ¹
Change Direction (CD)	Supported	CICS front-end cannot send CD. IMS sends CD to CICS if COMPT=SINGLE2 or MULT2.
Multiple chains per message (IUTYPE)	Supported ²	Not supported
IMS component selection (DATASTREAM)	Supported	Not supported for input. Input component=1. Output component=any.
RECORDFORMAT: ³ Undefined (U)	Supported ⁴	Not supported
RECORDFORMAT: ³ RU	Supported ⁴	Not supported
RECORDFORMAT: ³ VLVB	Supported	Supported
RECORDFORMAT: ³ CHAIN ⁵	Supported	Supported
FMH built by	Application program (BUILD ATTACH)	System: DPN, PRN, RDPN, RPRN specified in EXEC START command ⁶
Access to FMH	Yes (EXTRACT Attach)	Any supplied parameters available on EXEC RETRIEVE command ⁶

Table 123. CICS API for SEND/RECEIVE and START/RECEIVE (continued)

Attribute	SEND/RECEIVE	START/RETRIEVE
Relation to ISC session	Directly related (ALLOCATED)	Independent
IMS transaction types supported	Response mode Conversation mode Message switches ⁷ Commands Nonresponse, nonconversational ⁷	Nonresponse mode Nonconversation Message switches Commands ⁸
Recoverable I/O	Supported	Supported
Nonrecoverable I/O	Supported	Supported
Multisegment input	Supported	Single segment only
Multisegment output	Supported	Single segment only
IMS edits	Basic edit	Basic edit
Available	ISC Edit	ISC Edit
MFS: Without paging	Supported	Supported
MFS: Autopaged input: Single chain	Supported	Supported
MFS: Autopaged input: Single chain: Multichain	Supported	Not supported
MFS: Autopaged output	Supported	Not supported
MFS: Demand-paged output	Supported	Not recommended

Notes:

- Names can be embedded in the message text. The DPN field can be converted from a 4-character to an 8-character name by using DEFINE TRANSACTION REMOTENAME (or DFHPCT TYPE=REMOTE).
- Only used to send or receive MFS autopaged input or output messages.
- The ATTACH parameters are defined in the CONNECTION definition. However, the application using START/RETRIEVE must understand and provide the correct data format within the application.
- Marked as reserved by CICS, but can be specified on the BUILD ATTACH.
- Chained output from IMS requires the use of MFS.
- The RDPN and RPRN are limited to 4-character names.
- Nonresponse, nonconversational, and message switch input use SEND LAST only. CICS acquires replies to nonresponse and nonconversational transactions using RECEIVE.
- /DIS, /RDIS, /FOR only.

Related concepts:

“Relationship of ISC and IMS execution modes” on page 475

Related reference:

“ATTACH FM header format” on page 577

 Front-End Switch exit routine (DFSFEJ0) (Exit Routines)

Principal and alternate facility

The functions of CICS synchronous EXEC terminal control commands (such as SEND, RECEIVE, CONVERSE) are exercised against a principal or an alternate facility.

In CICS, the principal facility for a task is the terminal or ISC session that is made available to the application program when the task is initiated. The alternate facility is the ISC session that is allocated as needed by the application program. Commands issued against the principal facility are issued without the use of the SESSION (name) option. Commands issued against the alternate facility are issued with this option.

When CICS is the front-end subsystem, the user terminal is the principal facility. The session allocated to IMS is the alternate facility. However, when a restart transaction is attached from the ISC session, the session is the principal facility. The restart transaction does not have direct access to the user terminal.

Related concepts:

“Recovery and restart concepts” on page 645

CICS versus IMS conversation mode

In IMS, conversation mode is an attribute of a transaction in which the conversation is carried on through the transfer of a scratchpad area (SPA) and an associated message between the terminal and transaction message queues.

Each transfer of information to the transaction message queue and back to the terminal queue results in a sync point, and causes the commitment of resources and the release of locked resources. Also, all of the steps of an IMS conversation must occur entirely within a bracket. EB cannot be sent at any intermediate point in the conversation, because the receipt of EB causes IMS to discard the output message, terminate the conversation, and invoke the conversational abnormal termination exit. During a conversation, if a primary resource name (PRN) parameter is supplied in the input function management header, it is ignored, because the transaction code is carried in the SPA.

In CICS, a conversation is a series of interactions within a bracket between a terminal and an application. End-bracket occurs upon termination of the application. CICS also permits a “pseudo-conversation,” which is a sequence of transactions between a terminal and an application. Each transaction has one terminal input and one terminal reply within its own bracket. Either of these types of conversations can interact with IMS on the ISC session between the requests and replies by using the EXEC CONVERSE command. A message with end-bracket being sent to the terminal by CICS also causes an end-bracket message to be returned on the ISC session to IMS, if IMS has not previously sent an end-bracket status on the reply.

The differences in the way in which conversations are defined in CICS and IMS have the following implications for ISC:

- A terminal connected to CICS can only be held synchronously with the ISC session and IMS transaction processing if the CICS transaction is conversational in the CICS sense.
- If more than one input and output is to occur on the ISC session within a single bracket, the CICS transaction can interact with an IMS transaction that is defined as conversational. However, in the event of a session failure, the conversation might not be restartable.

- If only one input and output is to occur on the ISC session within a single bracket, the CICS transaction can interact with an IMS transaction defined as response mode. A series of transactions that make up a CICS pseudo-conversation can interact on an ISC session with a series of IMS transactions defined as response mode.

IMS conversational mode requires that IMS terminate the conversation by sending the last reply with RQD2,EB. Intermediate conversational messages are sent RQE2,CD. The CICS application can decide to terminate the conversation normally by issuing SYNCPOINT or RETURN. CICS can then terminate the conversation by sending LUSTATUS RQD2,EB in response to the last conversational output from IMS. In this case, the LUSTATUS acts as the acknowledgment for the RQE2,CD sent by IMS and closes that logical unit of work. The RQD2,EB causes IMS to respond with DR2 to close the bracket, dequeue the output message, and notify the Conversational Abnormal Termination exit routine that the conversation has been terminated normally by the remote program.

Related concepts:

“Recovery and restart concepts” on page 645

Sending IMS commands from CICS

IMS commands can be entered from any terminal authorized to use them.

For terminals that are defined statically, the Resource Access Control Facility (RACF) (or an equivalent product) is used for authorization. For terminals that are defined dynamically using ETO, this authorization is validated using RACF.

However, permitting a CICS application to issue IMS commands through the ISC session introduces IMS dependencies into the CICS application program. The consequences of this decision should be weighed before this facility is used.

A CICS application program can issue IMS commands by using the SEND/RECEIVE application program interface.

Exception: A CICS application program cannot issue the IMS commands /DISPLAY, /RDISPLAY, /FORMAT, and /TEST.

The IMS command verb is embedded in the message. It cannot be specified in the PROCESS parameter of the BUILD ATTACH command.

/DISPLAY, /RDISPLAY, and /FORMAT can be sent using START/RETRIEVE or SEND/RECEIVE. These commands differ from other IMS commands in that replies to these commands are queued by IMS rather than being immediately processed.

If these commands are issued with SEND (without LAST), IMS returns an LUSTATUS NO-OP as the reply to force end-bracket. IMS replies to these commands by sending the reply ATTACH BB/EB. CICS must use RECEIVE to obtain the reply. If a command is issued with SEND LAST, the reply message, when it is returned, is processed by a CICS transaction whose name has been previously specified in the RPROCESS field of the BUILD ATTACH.

If a command is issued using START/RETRIEVE, the reply is returned with both the ATTACH and SCHEDULER FM headers.

You can use the /TEST command for testing communication protocols and editing facilities on an ISC session when IMS is a back end to a CICS front end. Test mode requires the SEND/RECEIVE application program interface to be used.

Related concepts:

“Issuing IMS commands from an ISC session” on page 474

Chapter 5, “Overview of the Extended Terminal Option,” on page 79

 IMS security (System Administration)

Related tasks:

“Using IMS test mode for ISC VTAM sessions” on page 474

Sync points

ISC session protocols define RQ*2 requests and their associated responses (DR2 or exception DR2) as sync point requests and responses.

These requests and responses are functionally independent from RQ*1 requests and their associated responses (DR1 and exception DR1), because these latter requests and responses do not cause a sync point. The sync-point responses (DR2) are used between ISC session partners to ensure that both partners' subsystem sync point managers can commit or back out recoverable resources in synchronism.

All messages sent or received on an ISC session are defined as either recoverable or nonrecoverable, depending on the message type. The session response protocols are used to ensure that both partners of an ISC session mutually understand and agree with the recoverability attributes associated with each message. The response protocols used must be consistent with the message type and are enforced by the sending and receiving subsystem sync point managers.

When one session partner commits a message, the other session partner is notified by positive session sync-point responses. When a session partner backs out a message as the result of errors detected during input or output processing, the other session partner is notified either by session termination, by an exception response sent by the receiving subsystem, or by an LUSTATUS-function abort sense code sent by the sending subsystem. Backout results in discarding the currently active message and resetting of the associated ISC data flow control (DFC) and ATTACH states to those of the last sync point.

Because CICS is a direct-control subsystem, synchronous transactions directly control the ISC session. Therefore, backout during synchronous transaction processing also results in backing out application updates made since the last sync point.

Queued subsystems and asynchronous application output do not provide direct control of the ISC session to the transaction; therefore commit or backout does not affect application updates that are made asynchronously to the message on the ISC session. Because IMS is a queued subsystem, session sync points are separate from application sync points, and always occur outside control of the executing application.

Sync points on IMS input

For any type of input, IMS does not schedule the intended transaction until the complete input message is successfully received. Processing errors and session failure prior to the receipt of the complete input message cause the entire message to be discarded or backed out and have no effect on other recoverable IMS

resources, such as databases. However, the input message cannot be canceled by ISC session failures or protocols after the complete message is received, enqueued, and made available for scheduling. Only the current input message is backed out, even when several consecutive input (nonrecoverable) messages are received and enqueued because of a previous input sync point was requested. The definition and relationship of ISC input sync points to the application sync points depend upon whether the internal IMS execution mode for the input is synchronous or asynchronous.

Synchronous transactions from CICS are sent using SEND/RECEIVE, which generates ATTACH with CD. The following IMS transaction types are synchronous:

- Response mode transaction
- Conversational mode transaction
- IMS commands
- Test mode input

Asynchronous transactions from CICS are sent using SEND LAST, which generates ATTACH with EB, or START/RETRIEVE, which generates ATTACH SCHEDULER with EB. The following transaction types are asynchronous:

- Nonresponse mode transaction
- Nonconversational mode transaction
- IMS message switch

For IMS, the sync-point response returned to CICS can indicate successful receipt of the message or can indicate the results of IMS processing. For asynchronous input to IMS, the sync-point response returned to CICS indicates only that the input message has been successfully enqueued and that the responsibility for message recovery is assumed by IMS. For synchronous input to IMS, the sync-point response returned to CICS indicates that the transaction has executed successfully and a sync point has been taken, even though transaction execution is independent of the session. IMS reflects this by returning the input sync-point response only after the application-inserted reply message is made available for output as the result of a successful application sync point.

Sync points on IMS output

For output, IMS commits the output message when the requested sync-point response is returned by CICS. This means that the message has been successfully sent and dequeued and the session sync point information has been updated as appropriate. IMS backs out (depending upon the sense code used) the output message when an exception response is returned by CICS. In this case, the message is returned to the message queue for retransmission, and the associated DFC and Attach states are reset to those of the last ISC sync point. Backout on the ISC session does not affect the IMS application program or other recoverable resources, such as database updates. Backout of IMS conversational-mode output, in most cases, causes termination of the IMS conversation and the Conversational Abnormal Termination exit routine to be invoked. Based on the contents of the conversational SPA, user-provided exit logic might schedule another IMS transaction to back out database changes.

CICS sync points

For CICS, sync points occur under the control of the application program and can be issued at any time. For synchronous processing, when a sync-point request is

issued by the application program, CICS logs the completion of the logical unit of work (that is, the resources are committed). In addition, when the SYNCPOINT command is issued within a CICS application, CICS is responsible for appending the ISC message that was created as a result of previous processing to the SNA response request (RQD1 or RQD2) that IMS expects. However, when a CICS asynchronous transaction is scheduled as a result of IMS input, the mirror transaction has already issued a sync point, that causes the appropriate response to be returned to IMS. Any sync points issued subsequently by the application have no effect on the ISC session.

Related concepts:

“Keeping half sessions synchronized” on page 525

“Relationship of ISC and IMS execution modes” on page 475

“Logical unit of work” on page 645

Coding function management headers for CICS

CICS uses some of the same SNA-defined function management header fields that are used by IMS.

Unlike IMS, the CICS application program must prepare these function management headers and send them to IMS. CICS, itself, does not send them.

Both ends of an ISC session have an “attach manager” that performs the functions requested in the FMH.

Related reference:

“Function management headers” on page 559

ATTACH function management header

This topic describes the ATTACH function management header fields. The primary fields that CICS uses are ATTDPN, ATTPRN, ATTRDPN, and ATTRPRN.

ATTDPN

The CICS transaction specifies the outbound destination process name field (ATTDPN) in the PROCESS field of the BUILD ATTACH command.

This field, when sent by CICS to IMS, contains the name of an IMS editor (basic edit or ISC edit) or an MFS MID name that is to receive the inbound message. When received by CICS in a reply from IMS, this field contains the value that CICS sent to IMS as the return destination process name field (ATTRDPN) on the originating outbound message. If this is a reply to a nonresponse or nonconversational message sent to IMS with SEND LAST, this field identifies the transaction to be initiated to process this reply. In other cases, this value can be a CICS transaction code that identifies a restart transaction. If this value is not available, CICS uses the first four characters of the incoming data stream as the transaction code. If the reply is returned on the same session as that of the incoming message, IMS system code automatically wraps the incoming RDPN field into the DPN field of the reply message.

CICS ignores the ATTDPN if it is received in a reply returned from IMS.

Exception: After a session restart or a reply to a message sent with SEND LAST, ATTDPN is used to attach a CICS transaction.

When used with the SCHEDULER, the ATTDPN contains the DPN of the SCHEDULER model. In this case, the remaining ATTACH parameters are not sent by IMS, but rather, the information is supplied on the associated SCHEDULER model parameters.

Related reference:

“SCHEDULER function management header” on page 641

ATTPRN

The CICS transaction specifies the primary resource name field (ATTPRN) in the RESOURCE field of the BUILD ATTACH command.

This field, when sent by CICS to IMS, names an IMS LTERM or transaction code that is the message destination. When this parameter is omitted, IMS uses the first eight characters of the data stream to identify the message destination.

Recommendation: To maintain consistency between IMS and CICS, use this field to name an LTERM for IMS message switches and place transaction codes in the first eight characters of the data stream.

When received by CICS in a reply from IMS, this field contains the value that CICS sent to IMS as the return primary resource name (ATTRPRN) on the outbound message. CICS ignores the ATTPRN if it is received in a reply returned from IMS.

Exceptions:

- After session restart, ATTPRN is used to determine the identity of the terminal that originated this transaction and must be acquired during the restart.
- When the reply is to a nonresponse or nonconversational transaction issued by CICS with SEND LAST, ATTPRN is used.

If the reply is to be returned on the same session as that of the incoming message, IMS automatically wraps the value received on the incoming RPRN field into the PRN field. MFS can be used to set or override this value on IMS output.

ATTRDPN

The CICS transaction specifies the return destination process name (ATTRDPN) in the RPROCESS field of the BUILD ATTACH command.

When CICS sends this field to IMS, it contains the code of a transaction to be executed in the event that a session restart is required or a reply is to be returned to a message that CICS sent to IMS using SEND LAST. IMS sends this to CICS when a next MID is specified within an MFS message output descriptor (MOD). When CICS receives the ATTRDPN field on an input message, it can be examined by the EXTRACT ATTACH command. This value should be saved and returned to IMS as the DPN on subsequent replies.

ATTRPRN

The CICS transaction can place the identification of the originating terminal in the ATTRPRN field.

This identification is wrapped into the PRN field by IMS if a session restart is required or a reply is required to a message sent to IMS by CICS with SEND LAST. The CICS transaction places this value in the field by using the RRESOURCE field of the BUILD ATTACH and examines this field by using the EXTRACT ATTACH command.

For asynchronous unsolicited output, IMS uses the source LTERM name as the default, if a reply is necessary. MFS can be used to set or override this value on any type of output. When a CICS application receives this value, the value should be saved and returned as the PRN field for subsequent replies.

ATTDQN and ATTDQ

When sent to CICS from IMS, on the ATTACH for a demand-paged output message, ATTDQN contains a unique message identifier.

This field is only used for IMS demand paging. The CICS transaction specifies the destination message ID (ATTDQN) for subsequent demand-paged requests as the QNAME value in the returned QMODEL FM header. When CICS receives ATTDQN from IMS, the ATTDQ field is set to 1. These fields are not set by IMS on any other conditions nor are they accepted by IMS on input.

ATTIU

The CICS transaction specifies the interchange unit code (ATTIU) in the IUTYPE field of the BUILD ATTACH command.

Two values can be specified: single chain, which applies to all non-MFS input or output, and multichain, which applies only to MFS-autopaged input and output. Single chain should be used unless the input to IMS is multiple-page MFS-autopaged input.

ATTDSP

The CICS transaction specifies the data stream profile (ATTDSP) in the DATASTR field of the BUILD ATTACH command to determine the IMS component.

This field can contain the following values:

- X'00' Identifies IMS component 1
- X'01' Identifies IMS component 2
- X'02' Identifies IMS component 3
- X'03' Identifies IMS component 4

When received by IMS, this input component determines the default output component, the input security requirements, and whether MFS can be used for input. The output component determines whether MFS is used for output, and the protocols that are to be sent on asynchronous output are sent ATTACH SCHEDULER. On output, IMS sets this value to the value of the output LTERM.

ATTDBA

The CICS transaction specifies the deblocking algorithm (ATTDBA) in the RECFM field of the BUILD ATTACH command.

ATTACC

The ATTACC parameter is not supported on an ISC session and causes session termination if sent.

The following table summarizes the source of the values placed in the ATTACH FM header fields.

Table 124. Source of values placed in the ATTACH FM header fields

Major field	A ¹	B ²	C ³
ATTDPN ⁴	PROCESS	DPN	

Table 124. Source of values placed in the ATTACH FM header fields (continued)

Major field	A ¹	B ²	C ³
ATTPRN	RESOURCE	PRN	
ATTRDPN	RPROCESS	RDPN	
ATTRPRN	RRESOURCE	RPRN	
ATTDQN	QUEUE	Message ID ⁵	
ATTIU	IUTUPE ⁶	Paging ⁶	
ATTDSP	DATASTR ⁷		COMPTn ⁸
ATTDBA	RECFM ⁷		VLVB
			DPM-Bn ⁹

Notes:

1. CICS EXEC BUILD/EXTRACT ATTACH.
2. Set by IMS MFS; or, if MFS is not used, and the output is asynchronous unsolicited output, the RPRN defaults to the source LTERM name.
3. IMS COMPTn in TERMINAL/NAME macros.
4. For asynchronous processing, this field indicates “SCHEDULER follows” by containing the value X'02'. In this case, the remaining header fields are not contained in the ATTACH FM header, but are contained instead in the concatenated SCHEDULER FM header.
5. Message ID for MFS demand paging.
6. Value is set depending on MFS autopause.
7. Second byte of these halfword binary fields is used.
8. Components numbered 1 through 4.
9. DPM-Bn permits MFS to be used. Using MFS MODE=STREAM on output changes the ATTDBA to chain mode; otherwise, IMS uses the default of VLVB.

SCHEDULER function management header

The SCHEDULER function management header can be used to send asynchronous messages between IMS and CICS.

On input, receipt of the SCHEDULER FMH causes CICS to invoke the mirror transaction to schedule the asynchronous transaction that is to process the input. On output, it is generated as a result of CICS using the START command with its associated fields as follows:

SCDDPN

The CICS transaction specifies the value to be placed in the outbound destination process name field (SCDDPN) in the TRANSID field of the START command.

This field, when sent by CICS to IMS, contains the name of an IMS editor (MFS, basic edit, or ISC edit) that is to receive the inbound message. If SYSID is not coded on the START command, this TRANSID can be modified (to contain an 8-character name, for example) as a result of the RMTNAME option in the program control table (PCT).

When received by CICS in a reply from IMS, this field contains the value that CICS sent to IMS as the return destination process name field (SCDRDPN) on the outbound message. IMS automatically wraps the incoming RDPN field into the

DPN field of the reply message. Alternatively, this field might have been set by the IMS MFS. This value is a CICS transaction code. If this field is not supplied on input to CICS, CICS uses the first four characters of the input data stream as the transaction code.

SCDPRN

The CICS transaction specifies the primary resource name field (SCDPRN) in the TERMID field of the START command.

This field, when sent by CICS to IMS, names an IMS LTERM or transaction code that is the message destination. If this field is omitted when CICS sends a message to IMS, IMS examines the first eight bytes of the incoming message to extract the message destination.

Recommendation: To maintain consistency between IMS and CICS, use this field to name an LTERM for IMS message switches and place transaction codes in the first eight characters of the data stream.

When received by CICS in a reply from IMS, this field contains the value that CICS sent to IMS as the return primary resource name (SCDRPRN) on the outbound message or a value placed into the field by MFS. This field contains a value that identifies the terminal that is to be used as a principal facility and to which the reply should be sent.

SCDRDPN

The CICS transaction specifies the return destination process name (SCDRDPN) in the RTRANSID field of the START command.

When CICS sends this field to IMS, it contains the name of the asynchronous transaction to be scheduled by the CICS mirror transaction when the reply is received.

For unsolicited asynchronous output, the RDPN can be set by MFS to indicate the next message input descriptor (MID) to be used for subsequent replies. If this parameter is received as input by CICS, this parameter should be saved and returned to IMS in the TRANSID (DPN) of the START command.

SCDRPRN

The CICS transaction specifies the return primary resource name (SCDRPRN) in the RTERMID field of the START command.

When CICS sends SCDRPRN to IMS, it contains the identification of the terminal to which the return reply should be routed. On output from IMS to CICS, IMS sets SCDRPRN to the name of the source LTERM. MFS can also set SCDRPRN. The SCDRPRN value should be saved by CICS and returned to IMS as the TERMID (PRN) on any subsequent replies.

SCDDQN and SCDDP

Because asynchronous demand paging is not recommended, the SCDDQN and SCDDP fields are not used in a CICS-IMS ISC session.

The SCHEDULER header is preceded by an ATTACH header. Some mandatory ATTACH FM header fields also apply to asynchronous messages. These fields are not carried on the SCHEDULER header, but are retained on the ATTACH header. These fields are the ATTIU, ATTDSP, and ATTDDBA fields.

Their values in the asynchronous environment are as follows:

- **ATTIU**
CICS only supports single-chain messages between itself and IMS.
- **ATTDSP**
When sent using the START command, the ATTDSP value is determined by the terminal-control table generation, and should be X'00' (component 1). For input and output, the definition of the ATTDSP field for IMS is the same as that described in “ATTACH function management header” on page 638.
- **ATTDBA**
The CICS transaction specifies the deblocking algorithm (ATTDBA) in the RECFM field of the BUILD ATTACH command.

The following table summarizes the source of the values placed in the SCHEDULER FM header fields.

Table 125. Source of values placed in the SCHEDULER FM header fields

Major field	A ¹	B ²	C ³	D ⁴
SCDDPN	TRANSID ⁵		DPN	
SCDPRN	TERMID ⁵		PRN	
SCDRDPN	RTRANSID ⁵		RDPN	
SCDRPR	RTERMID ⁵		RPRN	
SCDDQN	QUEUE		Note 6	
ATTIU ⁷			Note 6	
ATTDSP ⁷		DATASTR ⁸		COMPTn ⁸
ATTDBA ⁷		RECFM		VLVB
				DPM-Bn ⁹

Notes:

1. CICS EXEC START/RETRIEVE.
2. CICS DEFINE CONNECTION or DFHTCT TYPE=SYSTEM.
3. Set by IMS MFS; or, if MFS is not used, for output that is unsolicited asynchronous output, the RPRN defaults to the source LTERM name.
4. IMS COMPTn in TERMINAL/NAME macros or on an ETO logon descriptor.
5. Four-byte name; MFS supports an eight-byte name.
6. MFS paging is not recommended on an asynchronous session between IMS and CICS.
7. ATTACH FMH fields preceding SCHEDULER FMH.
8. DATASTR must be specified as USER. This causes X'00', specifying IMS component 1, to be indicated as the input component. The IMS output component can be components 1 through 4 (X'00'- X'03').
9. DPM-Bn permits MFS to be used. Using MFS MODE=STREAM on output changes the ATTDBA to chain mode; otherwise, IMS uses the default of VLVB.

Related reference:

“ATTACH function management header” on page 638

Queue model function management headers

CICS uses the QMODEL headers to access IMS MFS demand-paged messages.

QMODEL headers are not sent by IMS to CICS. For SEND/RECEIVE, the CICS application program must prepare the QMODEL function management headers in order to receive demand-paged output from IMS. Temporary storage function shipping can be used with START/RETRIEVE. No EXEC commands exist to support the coding of QMODEL FM headers. The following QMODEL headers are used for MFS demand paging:

QGET Requests a page directly

QGETN
Requests the next sequential page

QPURGE
Requests termination of demand paging

QXFR Sent with the requested page

QSTAT
Sent when QPURGE is received, or when the requested page number is invalid

Data descriptor function management header

IMS uses the data descriptor function management header to specify a device page (DPAGE).

This header is not used by CICS, but can be passed to a CICS transaction. A CICS synchronous transaction can build and send data descriptor FM headers to IMS as appropriate. However, no EXEC commands exist to support the coding of data descriptor FM headers.

System message process (SYSMSG) function management header

IMS uses the system message process (SYSMSG) function management header to send system messages.

System messages are appended to an ATTACH carrying DPN=SYSMSG. CICS routes these messages to a synchronous application when the session is allocated to that transaction or to destination CSMT at any other time. If routed to a transaction, that transaction is responsible for processing the SYSMSG function management header. CICS does not send SYSMSG, but a synchronous application program can cause a SYSMSG to be sent by setting the appropriate process name in the BUILD ATTACH. The application program is responsible for building the appropriate SYSMSG FMH to be appended to the ATTACH.

SYSMSG is sent by IMS for broadcast messages or for the case in which a response to input has been sent, but has had a subsequent processing error. After a response to asynchronous input, SYSMSG is sent in lieu of the reply message. During synchronous processing, SYSMSG is sent:

- After a session restart, if the transaction in progress abends during the session outage
- During MFS output, if MFS output format blocks are not available or are not valid

Error recovery procedure function management header

CICS sends and receives the error recovery procedure function management header to transmit error information from one process to another.

The header is sent after an exception response (X'0846") and carries sense codes, which are used to inform the half-session partner of the nature of the error. It is followed by an error message. CICS writes the received IMS error message to the transient data destination CSMT; IMS writes the received CICS error message to the master terminal or the input source node.

Recovery and restart concepts

This topic describes the system and user functions that must be performed to recover an ISC session after a session, system, or application failure and assumes the reader understands the role of the STSN command in session resynchronization.

The concepts of sync point, commit, backout, and logical unit of work are common to IMS and CICS. However, their meanings differ slightly within IMS and CICS.

Related concepts:

"Determining session synchronism using STSN" on page 513

"Sync points" on page 636

Logical unit of work

A logical unit of work is any processing that occurs between sync points.

Within subsystems where the application directly controls the session, such as CICS, the application sync point and the ISC session sync point occur simultaneously as the result of a single explicit or implicit application command. Within queued subsystems, such as IMS, the application processing and sync points are independent of those of the ISC session. IMS must map the application sync point to the ISC session sync-point request.

In the synchronous case (SEND/RECEIVE), CICS is the front-end subsystem. One logical unit of work includes all of the processing performed in both the IMS and CICS subsystems, from the point at which the CICS application last issued a sync point to the point at which the next sync point is issued, and a sync-point response is returned.

If the transaction is recoverable, only one message can be sent to IMS before the sync point must be requested by the CICS application. However, a series of nonrecoverable requests and replies can occur between a CICS sync-point request and an IMS response.

For recoverable messages, IMS returns this response only after the IMS application has inserted a response or conversational mode reply message and caused an application sync point. IMS then begins the next logical unit of work by sending the reply message requesting a sync point. The reply message satisfies the RECEIVE issued by the CICS transaction. The CICS transaction should complete all necessary processing prior to issuing a SYNCPOINT or RETURN (implicit sync point) command.

The application sync point causes CICS to return a sync-point response to IMS. IMS then completes the logical unit of work by dequeuing the reply message from the output message queue.

In the asynchronous case, where CICS is the front-end subsystem, for recoverable input to IMS, the input request and the returned reply message are separate logical units of work. When the CICS application issues a SEND LAST or START and requests a sync point, IMS returns the sync-point response as soon as the transaction is enqueued (made available for scheduling). IMS sends any recoverable

asynchronous reply messages as they are made available by the transaction and requests a sync point on each. When the reply message is read by the CICS mirror transaction, an appropriate response is automatically returned to IMS. If the reply message is read using RECEIVE as a result of a previous SEND LAST, the sync point response is not returned to IMS until a subsequent SYNCPOINT, FREE, or RETURN command is issued.

One or more nonrecoverable transactions can be sent to IMS between CICS sync points using multiple START commands, each followed by a RETRIEVE. When the reply message is read by the CICS mirror transaction, an appropriate response is automatically returned to IMS.

The concept of logical unit of work (or work unit) is important in that session resynchronization must be performed when initiating a session and up to one work unit is considered to be indoubt on a flow from either the primary to the secondary half session or from the secondary to the primary. An indoubt work unit is one that is waiting to be committed or backed out based on the results of the session BIND and STSN flow.

Recovering outstanding message traffic after a failure

A CICS-IMS session can fail or terminate in any of several situations.

The situations in which a CICS-IMS session can fail or terminate include:

- A communication component (for example, VTAM or NCP) fails.
- The CICS or IMS subsystem fails.
- A direct VTAM or subsystem command (for example, an IMS /CLSDST or /STOP command) is entered.
- The CICS transaction or IMS transaction fails (indirectly as a result of subsequent error processing).

In each of these situations, the session failure appears similar to the remaining operative subsystems. However, the resulting recovery and resynchronization processes you must follow differ. This topic describes what happens in the event of a communication component failure, an IMS or CICS subsystem failure, or the issuing of a direct command against the session between the two subsystems.

Related concepts:

"Handling transaction abends" on page 649

Reestablishing the session

The failed session can be restarted by either half-session partner.

When it is restarted, the relationship of the half sessions is the same as when the session failed; that is, the former primary half session continues to be primary and the former secondary half session to be secondary. Both IMS and CICS remember this orientation by logging the session qualifier pair (SQP) used to initiate the session with an indicator of the session polarity.

The CICS master terminal operator can reinitiate the session using the command:
CEMT SET TERMINAL (*tttt*) ACQUIRED

where *tttt* is the TERMDNT of the CICS session as defined in the terminal control table. CICS then brings up the session, maintaining the session polarity that has been fixed by CICS system definition parameters. If the session is being initiated by an authorized IMS terminal operator's issuing the /OPNDST command, that

command must specify the same subpool name (IMS local session qualifier name) and session ID (CICS local session qualifier name) that was allocated to the session at the point of failure. These are available to the operator as necessary using an IMS /DISPLAY command. IMS automatically reestablishes the session, while maintaining the same session polarity as was in effect at session failure.

If the session is being restarted with IMS in cold-start mode, and if CICS is defined to be the primary half session, CICS must initiate the restart.

When the session is reestablished, IMS attempts to reset the data flow control bracket state to its status at the last sync point before session failure. CICS allows the session to be reestablished as between-brackets or as in-brackets with IMS in SEND state.

Resynchronizing the session

When a session is active, both IMS and CICS maintain a set of SNA message sequence numbers for that session. When a failure occurs and restart is attempted, the total of sync points on the session is checked between the subsystems by the STSN command.

The CICS/IMS session resumes if both of the following are true:

- The session is being cold-started.
- The half sessions' sequence numbers agree, or are within acceptable limits.

If a session fails, a CICS transaction having pending activity for that session also fails. CICS uses dynamic transaction backout (DTB) to ensure integrity of recoverable resources. The DTB indoubt parameter must be specified as IN-DOUBT(WAIT) or DTB=(YES, WAIT) in order to ensure consistency between the IMS and CICS subsystems within an ISC network.

WAIT holds locks on recoverable resources until resynchronization occurs.

If this situation is undesirable, you can specify IN-DOUBT(BACKOUT) or DTB=YES on an IMS-CICS ISC transaction. However, doing so might cause duplicate messages to be sent to IMS after session resynchronization, and can require logic in the IMS user transaction to resolve them.

If the DTB indoubt parameter does not specify WAIT, you must specify FORCSESS on the IMS TERMINAL macro system definition OPTIONS parameter or use an authorized IMS terminal operator /CHANGE command.

Related concepts:

“Determining session synchronism using STSN” on page 513

“Set-and-Test-Sequence-Numbers (STSN)” on page 930

Related tasks:

“Defining CICS backout in-doubt processing” on page 628

Processing outstanding traffic

In IMS, transaction processing is independent of session status; that is, if a message is received and successfully queued, it is always processed.

However, the way in which IMS handles subsequent replies and continues processing conversations depends on the type of message, the type of error that occurs, and whether the IMS subsystem fails.

In general, when only the session fails (and IMS does not fail), and no synchronous processing is occurring between IMS and CICS, the following is true:

- The session is bound between-brackets when it is reinitiated.
- IMS always sends or resends asynchronous replies or unsolicited asynchronous output after successful session restart.
- IMS sends or resends queued asynchronous replies resulting from /DISPLAY, /RDISPLAY, and /FORMAT commands.
- If the IMS asynchronous transaction abends during the session outage, IMS sends an error message using ATTACH SYSMMSG, because exception responses are no longer possible.

For replies processed synchronously by IMS:

- If response mode or conversational output is pending, the session is bound with IMS in-brackets/SEND to permit the pending reply to be sent or re-sent.
- When in conversation mode and input is required, IMS attempts to bind in-brackets/RECEIVE. CICS negotiates this bind to a between-brackets state, causing IMS to terminate the conversation, discard the output message, and invoke the Conversational Abnormal Termination exit routine. This exit routine can invoke user processing to schedule an IMS transaction to back out any database changes resulting from previous conversational processing based on the contents of the SPA.

When a session is reestablished and sent to between-brackets as described in the preceding paragraph, if the IMS transaction has not yet completed processing the conversational or response mode input (that is, output is not yet available), IMS terminates the session with a message to the master terminal operator requesting that the session be reinitiated (when output is available). This occurs because IMS cannot terminate the in-process transaction as required by the between-brackets bind.

- Except as mentioned, IMS cancels any commands received at the point of failure and discards pending synchronous reply messages that result from an IMS command.
- If the IMS synchronous transaction abends during the session outage, IMS sends an error message using ATTACH SYSMMSG, because exception responses are no longer possible.

IMS handles these transactions in the following manner:

- After an emergency restart, IMS discards a pending reply message resulting from a nonrecoverable transaction.
- After IMS receives and enqueues an input transaction, session protocols or failures cannot cancel that transaction.
- After a reply message to a CICS synchronous transaction (response mode or conversation) is enqueued, it indicates that the results of all previous IMS processing have been committed. These updates are not backed out even if the reply message is discarded.

If a CICS transaction must be re-sent to IMS as determined by STSN processing, the CICS application, rather than the CICS subsystem itself, must make this determination and either reconstruct the transaction or request that it be reentered.

After a session fails, CICS cannot reestablish the environment in which the original transaction was executing. That is, the connection to the terminal that originally entered the transaction no longer exists. Further, the transaction has abnormally terminated; therefore, it is necessary for the CICS application programmer to define

a “restart transaction” to be invoked by CICS to handle any IMS output that CICS might receive on the restarted session. When this transaction is invoked, the session, rather than the terminal, is now the primary facility. This transaction must in turn invoke an asynchronous transaction to acquire the terminal if output exists that needs to be sent to it.

The transaction code of the restart transaction to be invoked is carried in the DPN field of the function management header that is sent to CICS with the input message. This transaction code is the one specified in the RPROCESS field of the BUILD ATTACH command sent by CICS with the outbound message. This becomes the RDPN parameter automatically wrapped by IMS to the DPN field, which IMS sends on the outbound FMH unless modified by the IMS MFS.

The ID of the terminal to which CICS is to send any output reply is found in the PRN field of the incoming function management header. This is the value specified in the RRESOURCE field of the BUILD ATTACH sent by CICS with the outbound message. This becomes the RPRN parameter automatically wrapped by IMS to the PRN field of the outbound FMH unless modified by MFS. This value is used as the TERMID of the asynchronous transaction scheduled (START) by the restart transaction to acquire the terminal.

The restart transaction issues a RETRIEVE carrying both TRANSID and TERMID fields. When the transaction specified by the TRANSID is initiated, it owns the terminal specified by the TERMID as its principal facility.

Related concepts:

“Handling transaction abends”

“Coding CICS applications for restart” on page 650

Handling transaction abends

In addition to the considerations for session and subsystem failure, the design of CICS recovery transactions must also take into account the actions to be taken in the case of transaction termination as described in the topics that follow.

IMS transaction abend

When an IMS transaction terminates abnormally, all changes to IMS resources that were performed by this transaction are backed out by IMS. If the transaction is synchronous (CICS SEND/SYNCPOINT), IMS sends a negative response (exception DR2) to CICS. This causes the CICS transaction to also terminate abnormally. As a result of the negative response and subsequent CICS transaction abend, CICS backs out any changes made to recoverable resources after the previous CICS sync point. An error message is also sent to CICS; CICS routes this message to the master terminal destination CSMT. The exception occurs when MFS output formats do not exist in the format library or incur an I/O error. In this case, IMS has already sent a positive sync-point response and must indicate the MFS error to CICS by using an ATTACH SYSMSG.

If the transaction is asynchronous (CICS START or SEND LAST), IMS sends an error message to CICS; CICS routes this message to its master terminal destination. The error message is preceded by either a SYSMSG or an ERP header, depending upon the type of error that occurs. If the error occurs before IMS sends a sync-point response to input, an exception response and ERP are used; if after the sync-point response, ATTACH SYSMSG is used.

CICS transaction abend

If a transaction is initiated by a CICS subsystem acting as a front end for IMS, and that CICS transaction terminates abnormally before reaching sync point, any processing performed is handled in accordance with the specification on the DEFINE TRANSACTION IN-DOUBT parameter (or DFHPCT DTB= parameter).

If an asynchronous transaction is initiated by an IMS front end and the CICS transaction abends upon receipt of the transaction, IMS is not notified. This is because the receipt of the message causes the CICS mirror transaction to return an immediate DR2 to the asynchronous input and to schedule a transaction to process it. If this scheduled transaction fails, the ISC link to IMS is no longer active and no notification is possible. CICS does, however, notify the CICS destination CSMT of the transaction failure.

During synchronous processing, if CICS is the secondary half session, the ALLOCATE command has been successfully issued, and no message is available to be sent to IMS (due, for example, to DTB causing backout), CICS automatically deallocates the session by sending LUSTATUS BB/EB to IMS.

If the synchronous transaction terminates abnormally any time after it has sent a message to IMS using the SEND/SYNCPPOINT command, and that message has been enqueued on an IMS input queue, IMS processes it. When IMS attempts to send an output message to CICS, that message receives a negative response from CICS. As a result, the message can remain on the IMS output queue until it can be re-sent or until it is dequeued by the IMS master terminal operator.

Related reference:

“Coding function management headers for CICS” on page 638

“Error recovery procedure function management header” on page 644

Coding CICS applications for restart

When a failure occurs on an ISC session as a result of session or subsystem failure, session restart and resynchronization can be attempted.

If session resynchronization (STSN processing) is unsuccessful and the session cannot be restarted without intervention, the master terminal operators of both subsystems are notified. This situation can require that the terminal user also be notified of the session's status and of any actions that must be taken, such as reentering the failing transaction or waiting until corrective action is taken by the MTO.

After an ISC transaction defined as recoverable is received and logged by IMS, it can be recovered across session and subsystem failures. After restart, when the session has been recovered, if IMS has pending output, it is sent as follows:

- When IMS sends output on the same session as that on which the input message is received, that output message is sent with the same type of headers as those originally sent with the input. Thus, a reply to a message sent with SEND/RECEIVE or SEND LAST is returned with the ATTACH FM header and a reply to a message sent with START/RETRIEVE is returned with the ATTACH and SCHEDULER FM headers.
- When IMS sends output on a session other than that on which the input message is received, that output message is considered as unsolicited output and is sent asynchronously (with ATTACH SCHEDULER).

After a restart, when IMS sends an output message, CICS reads the input message and initiates a transaction.

- Within the application, the EXEC ASSIGN STARTCODE parameter is examined to determine whether this transaction had been initiated synchronously or asynchronously.
- If the restart transaction is initiated asynchronously:
 - A RETRIEVE is issued to obtain the ISC reply. The EIB indicators should be saved. CICS attaches the end-user's terminal to this CICS program as the principal facility. This permits the IMS output to be issued directly to the terminal using a CICS EXEC SEND command.
 - The EIB indicators, saved after RETRIEVE execution, must be tested to determine whether SYNCPOINT or RETURN can be issued to end this transaction.
- If the transaction is initiated synchronously by SEND/RECEIVE or asynchronously by SEND LAST:
 - The input message is obtained using RECEIVE. In this case, the session is now the principal facility. Therefore, the RECEIVE command does not require specification of the SESSION parameter.
 - After the RECEIVE, the CICS application saves and then checks the EIB fields EIBSYNC, EIBFREE, and EIBRECV, in that order.
 - The application now issues an EXTRACT ATTACH to determine the format of the input data and the ID of the terminal that originally submitted the transaction.
 - The restart logic must now START an additional CICS transaction to acquire the end user's terminal as the principal facility. The START request is made with the terminal ID found in the PRN field by using the EXTRACT ATTACH as the TERMID parameter on the START.
 - The EIB indicators are now tested to determine whether a sync point is required and whether the session can be freed. If this can be done, the application can issue a RETURN command to implicitly cause a sync point to be issued and the session to be freed.

Within the context of the restart transaction, data can be retrieved for inquiry purposes or databases and files can be updated.

In addition to the foregoing logic, a CICS restart transaction can contain logic to:

- Store input received from the originating terminal on temporary storage. If it is necessary to re-create an input transaction, this information can be obtained and re-sent to IMS.
- Create logic to inform the terminal user of any special processing to be performed, such as reentering a transaction or waiting for some master terminal operator action before proceeding.

Chapter 36. ISC data flow control examples

The following topics provide examples of ISC data flow control.

Non-MFS bracket and half-duplex protocol examples

The following series of examples illustrate ISC data flow control protocols.

In the examples, the items in parentheses are optional on the flow.

For simplicity, the examples in this topic assume only-in-chain SCHEDULER messages for flows in both directions, and the response, ATTACH, and SCHEDULER protocols have been excluded. Also, the bracket protocol is considered symmetrical in that either the primary half session (PHS) or the secondary half session (SHS) can initiate the bracket as illustrated.




Figure 87. Example of bracket protocol for a PHS component defined to IMS as SINGLE1

Sample flow 1 shows the PHS ending a bracket on the first output following receipt of change-direction. When no output exists and the input cannot guarantee output, the bracket is ended through a stand-alone LUSTATUS.

Sample flow 2 shows the PHS sending output that occurs while in a between-brackets state.




Figure 88. Example of bracket protocol for a PHS component defined to IMS as SINGLE2

Sample flow 1 shows the PHS returning the flow to a bracket initiator after receipt of a change-direction. The bracket initiator can optionally continue input or end

the bracket. When no output exists and the input cannot guarantee output, the flow is returned to the bracket initiator using a stand-alone LUSTATUS.

Sample flow 2 shows the PHS sending asynchronous output. Also, IMS detects a potential LUSTATUS CD loop and forces end-bracket using LUSTATUS.




Figure 89. Example of bracket protocol for a PHS component defined to IMS as MULT1

Sample flow 1 shows the PHS ending a bracket using an LUSTATUS. The bracket ends following the last output from a queue after receipt of change-direction. If no output exists and the input cannot guarantee output, the bracket is ended immediately by using a stand-alone LUSTATUS.

Sample flow 2 shows the PHS both beginning and ending a bracket by using an LUSTATUS following the last output from a queue. Additional output available from other queues causes subsequent brackets to be initiated.




Figure 90. Example of bracket protocol for a PHS component defined to IMS as MULT2

Sample flow 1 shows the PHS returning the flow to the bracket initiator using an LUSTATUS. The bracket ends following the last output from a queue after receipt of a change-direction. The bracket initiator can optionally continue input or end the bracket. If no output exists and the input cannot guarantee output (that is, the input is not in conversational or response mode), the flow is returned by using a stand-alone LUSTATUS.

Sample flow 2 shows the PHS sending output that occurs in a between-brackets state.

MFS bracket and half-duplex protocol examples

The following topics provide examples for both MFS output and MFS input.

MFS output examples

The following figures provide examples of MFS output.

In the examples shown in the following figures, the IMS message consists of 3 presentation pages:

- Logical page 1 that makes up 2 presentation pages
- Logical page 2 that makes up 1 presentation page

In all of the examples, the ATTACH FM header showing DPN=X'.03' is optional.

Exception: For the first paging request, the ATTACH FM header showing DPN=X'03' is required.




Figure 91. Example of demand-paged output, operator logical paging (OLP) not defined, sequential retrieval using QGETN FM header

IMS action: Remove message from the message queue

Note:

1. QCURSOR=0400010001, QCOUNT=020002
2. QCURSOR=0400010002, QCOUNT=020002
3. QCURSOR=0400020001, QCOUNT not present
4. DD FM header can precede the data

The QNAME returned on paging requests must be the same value as sent on the output ATTACH DQN.





Figure 92. Example of demand-paged output, OLP defined, QGET (last page request and by cursor) used. QPURGE used to dequeue the message


Note:

1. QCURSOR=0400010001, QCOUNT=020002
2. CURSOR=0400020001, QCOUNT not present
3. Data descriptor FM header can precede the data

These same conditions for non-SCHEDULER demand-paged output result in an exception response and in subsequent ERP FMH7 being sent.



Note:



Note:

2. Data descriptor FM header can precede the data in each chain. The ATTACH indicates a multichain message.




Figure 95. Example of nonpaged output message

Note: Data descriptor FM header can precede the data.

MFS input examples

The following figures provide examples of MFS input.

In the following examples, three DPAGEs are defined:

- DPAGE 1 defines data for 2 segments
- DPAGE 2 defines data for 1 segment
- DPAGE 3 defines data for 3 segments

Three LPAGEs are created for the message. A data descriptor function management header with a DSN supplied with each chain is used to select each DPAGE. The first DSN name selects DPAGE 1, the second DSN name selects DPAGE 2, and the third DSN name selects DPAGE 3.

*ATTACH and data for first DPAGE:




Figure 96. Example of autopaged input, three chains

A data descriptor FM header can precede the data.

Three LPAGEs are created for the message. A conditional test on the data is performed to select each DPAGE. Results of first test selects DPAGE 1, results of the second test selects DPAGE 2, and results of the third test selects DPAGE 3.

*ATTACH and data for first DPAGE:




Figure 97. Example of autopaged input, single transmission chain

The data descriptor FM header can be sent to select one DPAGE only. The ATTACH indicates a multichain message.

SBI/BIS examples

IMS can be either the primary half session (PHS) or the secondary half session (SHS) for all of the examples given. Therefore, all of the functions and commands shown for PHS and SHS are IMS functions and commands.

Although all SBI and BIS sequences are valid for any session using LU 6.1 protocols, the following examples assume IMS to be both the primary and the secondary half session. No assumptions are made here as to when or how a subsystem other than IMS might send or receive these sequences.




Figure 98. Example of PHS shutdown using SBI/BIS




Figure 99. Example of nonsimultaneous shutdown using SBI/BIS by both half sessions




Figure 100. Example of simultaneous PHS and SHS shutdown using SBI/BIS fails because of operator override

Signal protocol example

The following example illustrates the use of SIGNAL RCD.




Figure 101. Signal protocol example

Chapter 37. ISC error recovery procedure examples

The following topics provide examples of the ISC error recovery procedure.

Sender-detected error examples

The following figures show examples of sender-detected errors during paged output (or input). IMS is the sender in both figures.




Figure 102. Example of sender-detected error while sending demand-paged output message

Flow 1 in Figure 96 shows an IMS /DEQUEUE command occurring during the second output demand page. The CANCEL command terminates the chain (page) in progress, and the LUSTATUS terminates the receiving process. If the paged output message follows another input or output message that began a bracket, the BB is not sent on the first page OIC.

Sample flow 2 is the same, except that the IMS /DEQUEUE command occurs between output pages so that the CANCEL command is unnecessary.




Figure 103. Example of sender-detected error while sending autopaged output message

Again, the CANCEL command terminates the chain (page) in progress and the LUSTATUS terminates the receiving process. The initial begin-bracket and the resulting bracket send/receive protocol are the same as for demand-paged output in Figure 102 on page 665.

Flow 2 is the same as flow 1 in Figure 97, except that the IMS /DEQUEUE command occurs while between output pages, so that the CANCEL command is unnecessary.

Receiver-detected error examples

For simplicity, in the following series of figures, all examples assume only-in-chain (OIC) messages for flows in both directions. Also the bracket and ERP protocol are considered symmetrical in that either half session can initiate the bracket or ERP procedures as illustrated.

For all session terminations not involving system failures, ATTACH information is recovered and backed up to the last sync point. In the examples that follow, IMS is the receiver.




Figure 104. Example of receiver ERP for an IMS component defined as SINGLE1




Figure 105. Example of receiver ERP for an IMS component defined as *SINGLE2*




Figure 106. Example of receiver ERP for an IMS component defined as *MULT1*




Figure 107. Example of receiver ERP for an IMS component defined as *MULT2*

In this case, the receiver is the other subsystem. The IMS action follows selective receiver ERP sequence.




Figure 108. Example of demand-paged output with receiver-detected error

IMS action: Return message to queue and continue output.

In this example, flow 1 shows the receiver using receiver ERP to reject a demand-paged message. Flow 2 shows alternate methods with the same end result, but not involving an ERP message from the receiver.

Chapter 38. Sample program for IMS-CICS ISC

The following topics provide a sample program that illustrates the use of ISC between IMS and CICS.

By combining the sample program with the corresponding sample program in the *CICS Transaction Server for z/OS CICS Intercommunication Guide*, you can use the ISC function in both products.

The sample program is shipped as object and source code data set DFSISC00. The load library containing the object code is IMS.ADFSLOAD. The source library containing the source code is IMS.ADFSSRC.

Installation procedure

Before the sample program can be used to perform ISC functions between IMS and CICS , you must complete several steps.

The steps you must complete include:

1. Compile and bind the sample COBOL program.
2. Define the transaction codes to IMS using the system definition procedure.
3. Define an ISC session between IMS and CICS using the system definition procedure.
4. Compile the MFS utility statements to place the necessary formats into the online library.
5. Perform the PSB generation and then the ACB generation.
6. Initialize the IMS system.
7. Establish a session between IMS and CICS.
8. Run the corresponding CICS sample program.

In related topics, sample code is provided for performing steps 1 through 5.

IMS sample program (DFSISC00)

The following IMS sample program, written in COBOL, accepts an input message from an ISC session with CICS and returns the same message to CICS.

The program has two transaction codes: SAMPLA1 and SAMPLA2. When invoked using SAMPLA1, the output message is not formatted by MFS. When invoked using SAMPLA2, the output message is formatted by MFS using distributed presentation management (DPM).

```
CBL APOST
  IDENTIFICATION DIVISION.
    PROGRAM-ID. SAMPLA.
    AUTHOR:
  ENVIRONMENT DIVISION.
    CONFIGURATION SECTION.
    SOURCE-COMPUTER. IBM-370.
    OBJECT-COMPUTER. IBM-370.
  INPUT-OUTPUT SECTION.
    FILE-CONTROL.
  DATA DIVISION.
```

```

FILE SECTION.
  EJECT
WORKING-STORAGE SECTION.
77 BEGIN-LIT          PIC X(16) VALUE 'BEGIN 77 ENTRIES'.
77 FILLER             PIC X(45) VALUE
  '*****DATE LAST COMPILED: xxx xx, 19xx*****'.
77 CALL-FUNCTION      PIC XXXX.
77 MOD-NAME           PIC X(08).
  EJECT
01 INPUT-AREA.
  02 IN-LL            PIC S9999 COMP.
  02 IN-ZZ            PIC S999 COMP.
  02 TRAN-CODE        PIC X(08).
  02 INPUT-DATA       PIC X(79).
  02 INPUT-DATA-D     REDEFINES INPUT-DATA.
    05 INPUT-DATA-1   OCCURS 3 TIMES
                      PIC X(20).
    05 INPUT-DATA-2   PIC X(19).
01 OUTPUT-AREA.
  02 OUT-LL           PIC S9999 COMP.
  02 OUT-ZZ           PIC S999 COMP.
  02 OUTPUT-DATA      PIC X(79).
01 OUTPUT-AREA-1.
  02 OUT-LL-1         PIC S9999 COMP.
  02 OUT-ZZ-1         PIC S999 COMP.
  02 OUTPUT-DATA-1    PIC X(20).
  EJECT
01 DLI-FUNCTIONS.
  02 GET-UNIQ         PIC XXXX VALUE 'GU '.
  02 GET-NEXT         PIC XXXX VALUE 'GN '.
  02 ISRT             PIC XXXX VALUE 'ISRT'.
  EJECT
01 STATUS-ERROR-SEG.
  02 FILLER           PIC S999 COMP VALUE +83.
  02 FILLER           PIC S999 COMP VALUE +0.
  02 FILLER           PIC X(28) VALUE
    '***** STATUS ERROR *****'.
  02 FILLER           PIC X(10) VALUE ' TRANCODE:'.
  02 ERROR-TRAN       PIC X(8).
  02 FILLER           PIC X(17) VALUE
    ' STATUS RECEIVED:'.
  02 ERROR-STATUS     PIC XX.
  02 FILLER           PIC X(10) VALUE ' FUNCTION:'.
  02 ERROR-FUNCTION   PIC XXXX.
  EJECT
LINKAGE SECTION.
01 IOTP-PCB.
  02 IOTP-LTERM       PIC X(8).
  02 FILLER           PIC XX.
  02 IOTP-STATUS      PIC XX.
  02 IOTP-PREFIX.
    03 IOTP-DATE      PIC S9(7) COMP-3.
    03 IOTP-TIME      PIC S9(7) COMP-3.
    03 IOTP-MSG-NUMBER PIC S999 COMP.
    03 FILLER         PIC XX.
  02 IOTP-MOD-NAME    PIC X(8).
  EJECT

PROCEDURE DIVISION.
  ENTRY 'DLITCBL' USING IOTP-PCB.
100-RETRIEVE-MESSAGE-SEGMENT.
  MOVE GET-UNIQ TO CALL-FUNCTION.
  CALL 'CBLTDLI' USING CALL-FUNCTION IOTP-PCB INPUT-AREA.
  IF IOTP-STATUS = 'QC'
    GO TO 800-GOBACK-ROUTINE.
  IF IOTP-STATUS NOT = SPACES
    GO TO 700-INVALID-STATUS-CODE.

```

```

200-CHECK-TRAN-CODE.
*****
*   THE ONLY DIFFERENCE BETWEEN THESE TWO TRANSACTIONS   *
*   IS THE ABSENCE OR PRESENCE OF MFS.  SAMPLA1 DOES NOT  *
*   CONTAIN MFS.  SAMPLA2 CONTAINS MFS.                  *
*****
      IF TRAN-CODE = 'SAMPLA2'
          PERFORM 400-SAMPLA2-ROUTINE
              THRU 450-SAMPLA2-ROUTINE-EXIT,
      ELSE
          PERFORM 300-SAMPLA1-ROUTINE
              THRU 350-SAMPLA1-ROUTINE-EXIT.
      GO TO 100-RETRIEVE-MESSAGE-SEGMENT.
      EJECT

300-SAMPLA1-ROUTINE.
      MOVE ISRT TO CALL-FUNCTION.
      SUBTRACT 8 FROM IN-LL GIVING OUT-LL.
      MOVE IN-ZZ TO OUT-ZZ.
      MOVE INPUT-DATA TO OUTPUT-DATA.
      CALL 'CBLTDLI' USING CALL-FUNCTION IOTP-PCB OUTPUT-AREA.
      IF IOTP-STATUS NOT = SPACES
          GO TO 700-INVALID-STATUS-CODE.
350-SAMPLA1-ROUTINE-EXIT.
      EXIT.

400-SAMPLA2-ROUTINE.
      MOVE 'MODA' TO MOD-NAME,
      MOVE +0 TO OUT-ZZ-1.
      MOVE 24 TO OUT-LL-1.
      MOVE INPUT-DATA-1 (1) TO OUTPUT-DATA-1.
      PERFORM 500-INSERT-ROUTINE
          THRU 550-INSERT-ROUTINE-EXIT.
      SUBTRACT 20 FROM IN-LL.
      IF IN-LL IS LESS THAN 13 GO TO 450-SAMPLA2-ROUTINE-EXIT.
      MOVE INPUT-DATA-1 (2) TO OUTPUT-DATA-1.
      PERFORM 600-INSERT-ROUTINE
          THRU 650-INSERT-ROUTINE-EXIT.
      SUBTRACT 20 FROM IN-LL.
      IF IN-LL IS LESS THAN 13 GO TO 450-SAMPLA2-ROUTINE-EXIT.
      MOVE INPUT-DATA-1 (3) TO OUTPUT-DATA-1.
      PERFORM 600-INSERT-ROUTINE
          THRU 650-INSERT-ROUTINE-EXIT.
      SUBTRACT 20 FROM IN-LL.
      IF IN-LL IS LESS THAN 13 GO TO 450-SAMPLA2-ROUTINE-EXIT.
      MOVE INPUT-DATA-2 TO OUTPUT-DATA-1.
      PERFORM 600-INSERT-ROUTINE
          THRU 650-INSERT-ROUTINE-EXIT.
450-SAMPLA2-ROUTINE-EXIT.
      EXIT.
      EJECT

500-INSERT-ROUTINE.
      MOVE ISRT TO CALL-FUNCTION.
      CALL 'CBLTDLI' USING CALL-FUNCTION IOTP-PCB OUTPUT-AREA-1
          MOD-NAME.
      IF IOTP-STATUS NOT = SPACES
          GO TO 700-INVALID-STATUS-CODE.
550-INSERT-ROUTINE-EXIT.
      EXIT.

600-INSERT-ROUTINE.
      MOVE ISRT TO CALL-FUNCTION.
      CALL 'CBLTDLI' USING CALL-FUNCTION IOTP-PCB OUTPUT-AREA-1.
      IF IOTP-STATUS NOT = SPACES
          GO TO 700-INVALID-STATUS-CODE.
650-INSERT-ROUTINE-EXIT.
      EXIT.

700-INVALID-STATUS-CODE.
      MOVE CALL-FUNCTION TO ERROR-FUNCTION.

```



```

        MOVE IOTP-STATUS TO ERROR-STATUS.
        MOVE TRAN-CODE TO ERROR-TRAN.
        MOVE ISRT TO CALL-FUNCTION.
        CALL 'CBLTDLI' USING CALL-FUNCTION IOTP-PCB STATUS-ERROR-SEG.
300-GOBACK-ROUTINE.
        GOBACK.

```

Job control statements for the sample program

Use the following JCL to compile and bind the sample program.

This JCL uses the z/OS COBOL 2.4 compiler and places the resulting load module in IMS.PGMLIB.

```

//IMSCOBOL JOB ACCT,NAME,CLASS=A,MSGLEVEL=(1,1),MSGCLASS=A
//          PROC MBR=,PAGES=60,RGN=512K,
//              SOUT=A
//C          EXEC PGM=IKFCBL00,REGION=&RGN,
//              PARM=' SIZE=130K,BUF=10K,LINECNT=50,APOST,BATCH'
//SYSLIN DD   DSN=&&LIN,DISP=(MOD,PASS),UNIT=SYSDA,
//              DCB=(USER.PROCLIB),
//              SPACE=(3520,(40,10),RLSE,,ROUND)
//SYSPRINT DD SYSOUT=&SOUT,DCB=(LRECL=121,BLKSIZE=605,RECFM=FBA),
//              SPACE=(605,(&PAGES.0,&PAGES),RLSE,,ROUND)
//SYSUT1 DD   UNIT=SYSDA,DISP=(,DELETE),
//              SPACE=(3520,(100,10),RLSE,,ROUND)
//SYSUT2 DD   UNIT=SYSDA,DISP=(,DELETE),
//              SPACE=(3520,(100,10),RLSE,,ROUND)
//SYSUT3 DD   UNIT=SYSDA,DISP=(,DELETE),
//              SPACE=(3520,(100,10),RLSE,,ROUND)
//SYSUT4 DD   UNIT=SYSDA,DISP=(,DELETE),
//              SPACE=(3520,(100,10),RLSE,,ROUND)
//L          EXEC PGM=IEWL,REGION=&RGN,PARM='XREF,LET,LIST',
//              COND=(4,LT,C)
//*STEPLIB DD   DSN=IMS.SDFSRESL,DISP=SHR
//SYSLIB DD   DSN=SYS1.COBLIB,DISP=SHR
//SDFSRESL DD   DSN=IMS.SDFSRESL,DISP=SHR
//SYSLIN DD   DSN=&&LIN,DISP=(OLD,DELETE),VOL=REF=*.C.SYSLIN
//          DD   DSN=IMS.PROCLIB(CBLTDLI),DISP=SHR
//          DD   DDNAME=SYSIN
//SYSLMOD DD   DSN=IMS.PGMLIB(&MBR),DISP=SHR
//SYSPRINT DD SYSOUT=&SOUT,DCB=(RECFM=FBA,LRECL=121,BLKSIZE=605),
//              SPACE=(605,(&PAGES.0,&PAGES),RLSE,,ROUND)
//SYSUT1 DD   UNIT=(SYSDA,SEP=(SYSLMOD,SYSLIN)),DISP=(,DELETE),
//              SPACE=(3520,(100,10),RLSE,,ROUND)
//          PEND
//IMSCOBOL EXEC IMSCOBOL,SOUT=A,MBR=SAMPLA
//C.SYSIN DD *

```

You can also bind the sample program directly from IMS.ADFSLOAD to IMS.PGMLIB.

IMS system definition statements

Use the following statements to define to IMS the two transaction codes supported by the sample program.

Restriction: The DPM-Bn specified in the following code must match the DPM-Bn specified on the MFS format DEV statement.

```

APPLCTN PSB=SAMPLA
        TRANSACT CODE=SAMPLA1,INQ=(YES,NORECOV)
        TRANSACT CODE=SAMPLA2,INQ=NO
TYPE    UNITYPE=LUTYPE6,OPTIONS=(TRANRESP,OPNDST,NOMTOMSG,      X
                                SYNCSESS),                        X

```

```

MSGDEL=SYSINFO,
OUTBUF=256,
SEGSIZE=256
*
**CICSA1 PARALLEL SESSION NODE
*
      TERMINAL  NAME=CICSA1,
      SESSION=5,
      COMPT1=(SINGLE1,DPM-B1,IGNORE),
      COMPT2=(SINGLE2,DPM-B1,IGNORE),
      COMPT3=(MULT1,DPM-B1,IGNORE),
      COMPT4=(MULT2,DPM-B1,IGNORE)
VTAMPOOL
*
      SUBPOOL  NAME=PS01
      NAME  PSLT01,COMPT=1,ICOMPT=1
      NAME  PSLT02,COMPT=2,ICOMPT=2
      NAME  PSLT03,COMPT=3,ICOMPT=3
      NAME  PSLT04,COMPT=4,ICOMPT=4

```

MFS formats

The following statements comprise the JCL and MFS statements needed to produce the MFS formats used by the sample program.

This is shipped as copy code in data set MFSSISC0 in IMS.ADFSMAC.

```

//MFSUTL JOB ACCT,NAME,CLASS=A,MSGLEVEL=(1,1)
//JOB LIB DD DSN=IMS.SDFSRESL,DISP=SHR
//      PROC RGN=360K,SOUT=A,SNODE=IMSVS,
//      SOR=NOLIB,MBR=NOMBR,PXREF=NOXREF,
//      PCOMP=NOCOMP,PSUBS=NOSUBS,PDIAG=NODIAG,
//      COMPR=NOCOMPRESS,COMPR2=COMPRESS,
//      LN=55,SN=8,DEVCHAR=0
//S1      EXEC PGM=DFSUPAA0,REGION=&RGN,
//      PARM=(&PXREF,&PCOMP,&PSUBS,&PDIAG,&COMPR,
//      'LINECNT=&LN,STOPRC=&SN,DEVCHAR=&DEVCHAR')
//SYSLIB DD DSN=IMS.SDFSMAAC,DISP=SHR
//SYSIN DD DSN=&SNODE..&SOR.(&MBR),DISP=SHR
//REFIN DD DSN=IMS.REFERAL,DISP=OLD
//REFOUT DD DSN=IMS.REFERAL,DISP=OLD
//REFRD DD DSN=IMS.REFERAL,DISP=OLD
//SYSTEM DD DSN=&TXTPASS,UNIT=SYSDA,
//      SPACE=(CYL,(1,1)),DCB=BLKSIZE=800
//SYSUT3 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT4 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//UTPRINT DD SYSOUT=&SOUT
//SYSPRINT DD SYSOUT=&SOUT,DCB=(RECFM=FBA,LRECL=133,BLKSIZE=1330)
//SYSUDUMP DD SYSOUT=&SOUT
//SEQBLKS DD DSN=&&BLKS,DISP=(NEW,PASS),
//      UNIT=SYSDA,SPACE=(CYL,(1,1))
//S2      EXEC PGM=DFSUNUB0,REGION=&RGN,
//      PARM='&COMPR2,DEVCHAR=&DEVCHAR',
//      COND=(8,LT,S1)
//SEQBLKS DD DSN=&&BLKS,DISP=(OLD,DELETE)
//UTPRINT DD SYSOUT=&SOUT,DCB=(RECFM=FBA,LRECL=133,BLKSIZE=1330)
//SYSUDUMP DD SYSOUT=&SOUT
//FORMAT DD DSN=IMS.FORMAT,DISP=OLD
//SYSPRINT DD SYSOUT=&SOUT
//SYSUT3 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT4 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
// PEND
//MFSUTL EXEC MFSUTL,SOUT=A
//S1.SYSIN DD *
      PRINT ON,NOGEN
FM TA      FMT

```

```

SPACE 3
DEV  TYPE=DPM-B1,FEAT=IGNORE,DSCA=X'00A0'
SPACE 2
DIV  TYPE=OUTPUT,OPTIONS=(DPAGE,NODNM)
SPACE 3
DPAGE01 DPAGE FILL=NULL
PPAGE01 PPAGE
      DFLD 'PAGE'
LPGNO  DFLD LTH=04
      DFLD ' '
SPACE 2
DATA01 DFLD LTH=20
SPACE 3
      FMTEND
      EJECT
MODA   MSG  TYPE=OUTPUT,SOR=(FMTA,IGNORE),PAGE=YES
SPACE 3
      SEG
SPACE 2
      MFLD (LPGNO,LPAGENO)
      MFLD DATA01,LTH=20
SPACE 3
      MSGEND
      END
/*

```

Program specification block (PSB) generation for the sample program

The following JCL and program specification block (PSB) generation statements produce the IMS control blocks necessary to execute the sample program.

```

//PSBGEN EXEC PSBGEN,SOUT=A,MBR=SAMPLA
//C.SYSIN DD *
          PSBGEN  LANG=COBOL,PSBNAME=SAMPLA
          END
/*

```

Application control block (ACB) generation

Use the following JCL and application control block (ACB) utility statements to place the PSB into the appropriate IMS online library.

```

//ACBGEN JOB ACCT,NAME,CLASS=A,MSGLEVEL=(1,1)
//ACBGEN EXEC ACBGEN,SOUT=A
//G.SYSIN DD *
      BUILD PSB=SAMPLA
      BUILD PSB=SAMPLB
/*

```

Part 8. Multiple Systems Coupling (MSC)

These topics introduce multiple systems coupling (MSC). You can use MSC to connect multiple IMS subsystems. These topics include an overview of MSC and the information you need to design, implement, and administer an MSC network.

Chapter 39. Overview of Multiple Systems Coupling

Multiple Systems Coupling (MSC) makes it possible for transactions to be entered in one IMS and processed in another IMS.

The responses can be returned to the terminals that entered the transactions or to other terminals. IMS uses MSC to route and control message traffic between connected IMS systems.

Multiple Systems Coupling concepts and terminology

MSC provides the ability to connect geographically dispersed IMS systems in such a way as to allow programs and operators of one IMS to access programs and operators of the connected IMS systems.

Communication can occur between two or more (up to 2036) IMS systems running on any supported combination of operating systems.

MSC also provides a way to extend the throughput of an IMS beyond the capacity of a single CPU. This extension is possible if the IMS applications can be partitioned among IMS systems two ways:

- *Horizontal partitioning*: Applications execute in more than one IMS with database contents split between IMS systems.
- *Vertical partitioning*: Applications execute in one IMS with the complete database that they reference attached to that IMS. Transactions can originate in any IMS.

A *link* is a connection between two IMS systems. All links must be defined during the IMS system definitions for each IMS. There are two types of links: physical links and logical links.

- A *physical link* is the access method connection or hardware connection between two IMS systems.
- A *logical link* is the mechanism through which a physical link is related to the transactions and terminals that make use of that physical link.

You can assign a logical link to a physical link during system definition, or the operator can assign it dynamically during system execution.

MSC physical links

A physical link is how the IMS systems connect to one other through access methods or hardware.

You define a physical link with the MSPLINK macro.

A maximum of 999 physical links are allowed in each IMS in a Multiple Systems Coupling (MSC) network, except in IMS systems that are defined as large systems by the LGEN parameter in the IMSCtrl macro, where the maximum is 676.

Multiple Systems Coupling supports the following types of physical links:

Channel-to-channel (CTC) adapter

Usually used only when the IMS systems are in the same data center.

Memory-to-memory (MTM)

Used when the IMS systems are in the same logical partition.

VTAM

Usually used when the IMS systems are in different data centers.

VTAM is an access method that usually uses a teleprocessing media connection. The CTC adapter is a channel-to-channel hardware connection. The memory to memory (MTM) link is a software link between IMS subsystems that are running in the same operating system, and it is used primarily for backup and testing purposes. Physical link buffer sizes must be equal between the two IMS systems.

Depending on the type of physical link, one or multiple logical links can be assigned to a physical link.

The following figure illustrates the types of physical links.




Figure 109. MSC physical link types

Related reference:

- ➡ UPDATE MSPLINK command (Commands)
- ➡ QUERY MSPLINK command (Commands)
- ➡ MSPLINK macro (System Definition)

MSC logical links

A logical link relates a physical link to the transactions and terminals that can use that physical link.

Each IMS in an MSC network has one or more defined logical links.

A maximum of 676 logical links are allowed for each IMS in the MSC network.

Two IMS systems that are defined to communicate with each other, each through a specific logical link, are called *partner systems*.

To establish connection between two IMS systems, each partner must have a logical-link definition. The two logical-link definitions must specify the same partner identifications and be assigned to the same physical link. The IMS system definition process assigns a number to each defined logical link. Logical link numbers are assigned sequentially, beginning with 1, in the order in which the links are defined. A logical link can be reassigned to a different physical link, but the two IMS systems must always communicate through a logical link partnership.

IMS also assigns a default name to each logical link, unless you specify a different logical link name in the label field on the MSLINK macro. Default logical link names are DFSLxxxx, where xxxx is the logical link number. You can display or modify the logical link name by using the type-2 commands QUERY MSLINK and UPDATE MSLINK.

Type-1 commands require you to specify the logical link number to identify the target logical link. Type-2 commands require you to specify the logical link name to identify the target logical link.

With VTAM, multiple sessions can use the link. During IMS system definition, you can specify how many sessions are to share the SDLC link. Each session that is activated becomes a logical link between partner systems.

The IMS system definition process does not require that a physical link be specified for each logical link. You can assign a physical link to the logical link online by using either the type-1 IMS command /MSASSIGN or the type-2 IMS command UPDATE MSLINK. No communication between partners can occur until the assignment is made.

Related reference:

-  UPDATE MSLINK command (Commands)
-  QUERY MSLINK command (Commands)
-  MSLINK macro (System Definition)

MSC logical link paths

Messages are routed in an MSC network by using *logical link paths*. Logical link paths identify the remote IMS system to which a message must be delivered and the local IMS system that is delivering the message.

You must define at least two corresponding logical link paths to submit a transaction at a terminal in an *input IMS system*, process the transaction in *destination IMS system*, and return a response to the terminal in the input IMS system: One logical link path in the input IMS system and another in the destination IMS system.

In the logical link path that is defined in the input IMS system, the input IMS system is defined as local and the destination IMS system is defined as remote.

In the logical link path that is defined in the destination IMS system, the destination IMS system is defined as local and the input IMS system is defined as remote.

In the definition of a logical link path, each IMS system is identified by one of its assigned system identifiers (SYSIDs), a numeric value from 1 to 2036.

IMS associates the SYSID of the input IMS system with the input transaction message and its corresponding response, so that the response can be returned to the input terminal.

Before it can be used, a logical link path must be assigned to a logical link. You can assign multiple logical link paths to a single logical link.

You define logical link paths by coding the MSNAME stage-1 system definition macro. Because of the macro name, the term *MSNAME* is often used as a synonym for logical link path.

Examples

Consider the MSNAME macro definitions of two logical link paths:

```
MSNAME  SYSID=(2,1)
MSNAME  SYSID=(3,1)
```

The first definition says that messages that use this logical link path are processed in the remote system whose local SYSID is 2. The second definition says that messages that use this logical link path are processed in the remote system whose local SYSID is 3. By using these definitions, the IMS system definition process assigns SYSID 1 to the IMS being defined and recognizes two remote systems with SYSIDs of 2 and 3. If a third path is defined with SYSID=(5,4), IMS would also assign SYSID 4 to the local system.

Transactions are assigned to logical link paths in the APPLCTN macro definition.

In the next example, consider the following application definitions, each with one transaction code defined:

```
APPLCTN  PSB=A
          TRANSACT  CODE=A
APPLCTN  PSB=B,SYSID=(2,1)
          TRANSACT  CODE=B
APPLCTN  PSB=C,SYSID=(3,1)
          TRANSACT  CODE=C
```

The SYSID keyword identifies the logical link path to be used for the transactions that are associated with the application. Transaction A is considered to be a local transaction, because the absence of the SYSID keyword indicates transaction A is only processed by the IMS being defined. Transactions B and C are remote transactions. Relating the application definitions to MSNAME definitions, IMS would return responses from transactions B and C to the IMS defined as SYSID 1, unless the application program specified an alternative destination for the response.

If messages that originate in the local system refer to any logical terminals in a remote system, the logical link definition must also include NAME macros to identify those remote logical terminals, unless directed routing is used.

The MSC network and routing

The following topics explain the concepts necessary for MSC administration.

Remote and local systems

In an MSC network, a *local system* refers to a specific IMS where a message is entered. All other IMS systems are considered *remote systems* in regard to the specific local system.

A *local transaction* is a transaction that is processed in the same IMS in which it is entered. A *remote transaction* is a transaction that is entered into a IMS from a terminal or a link that is not processed in that IMS.

The following figure shows local and remote transactions and systems. When Transaction A is entered in IMS A, which is the local system, transaction A is processed locally. When Transaction A is entered on IMS B, which this time is the local system, it is sent across the MSC link to remote system IMS A and is processed remotely.




Figure 110. Remote and local transactions and systems

In the previous figure, transaction A, when entered from IMS A and processed in IMS A, is considered a local transaction. When transaction A is entered from IMS B and sent across an MSC link to be processed in IMS A, it is considered a remote transaction.

Flow of data within multiple systems

The flow of a transaction in an MSC network requires additional steps as compared to one IMS.

The general steps are illustrated in the figure below, and explained as follows:

- In the local system, a remote transaction entered from an LTERM is placed on the message queue of the local system with the destination of the remote transaction name (1). (The message is queued to the MSNAME associated with the specified remote destination.)

- MSC removes the message from the message queue (2), sends it across the MSC link (3), and places it on the message queue of the remote system (4).
- The remote system sends the message from the message queue to the application program to be processed (5). After the application program processes the message, the program sends a reply (6).
- The remote system places the reply message, with a destination of the output LTERM, on its message queue (7).
- MSC removes the message from the message queue of the remote system (8) and sends it back across the MSC link (9).
- MSC places the message on the message queue of the local system (10) and sends it to the output LTERM (11).




Figure 111. Remote transaction flow

Message routing

The message-routing function of MSC supports several types of message routing.

The types of message routing supported by MSC include:

- Routing of transaction messages from a terminal in one IMS to an application program in another IMS. The transaction can be defined as recoverable, nonrecoverable, response mode, or conversational.

Restriction: Fast Path transactions across an MSC link are not supported. However, Fast Path transactions within the local system are supported.

- Routing of message switches from an LTERM in one IMS to an LTERM in another IMS and message switches between LTERMs in the same IMS. This support includes messages sent with a /BROADCAST command.
- Routing of response messages from an application program to the terminal that sends the transaction, or messages from an application program to an alternate terminal. Routing messages to alternate remote terminals (the transaction and the LTERM are in different IMS systems) requires that the alternate LTERM be defined as a remote LTERM. If directed routing or the TM and MSC Message Routing and Control User Exit routine is used, the alternate LTERM does not need to be defined as remote.

- Program-to-program switches between transactions in different IMS systems or within the same IMS.

Routing path

IMS passes messages from the local system to the remote system on a routing path. One or more IMS systems can be included in a routing path.

In the following figure, IMS B has a path to IMS A and back (path BA). Similarly, back-and-forth paths exist between IMS B and IMS D (path BD), IMS B and IMS C (path BC), and IMS A and IMS D (path AD). A path can go through an IMS in order to get to another IMS, such as path CAD between IMS C and IMS D. More than one path can exist between the same two IMS systems. IMS C and IMS D have two direct paths, CD1 and CD2, in addition to the indirect path CAD.

A total of 2036 paths are allowed in one MSC network.




Figure 112. Routing path

Logical destinations

Message routing for an MSC network uses logical destinations, as does a single-system environment.

A destination is either an LTERM or a transaction code. A local destination resides in the local system, and a remote destination resides in a remote system. Within each local system, all local and remote destinations must be defined with unique names. In the figure below, all of the local and remote destinations are uniquely defined within each local system. Destinations that are defined remotely are not also defined locally within the same IMS. Similarly, destinations that are defined locally are not also defined remotely.

The same destination names can be used for local destinations in different IMS systems that are connected by MSC. The destination names cannot conflict with the global intent of the destination within the MSC network. For example, in the figure below, TRANAB is a local transaction in IMS B and IMS A. It is a remote transaction in IMS C. IMS C is referencing the local TRANAB in IMS A only and not the one in IMS B. IMS C cannot remotely reference TRANAB in IMS B. The destination system that is referenced in a remote destination is determined by the system identification (SYSID) value.

With directed routing and the TM and MSC Message Routing and Control User exit routine, you can send messages from an application program to a remote destination that is not explicitly defined in the destination system. The validation of the destination name in the transaction processing system is delayed until the message arrives in the local system. Delaying this validation provides more flexibility in the resource naming requirements.




Figure 113. Logical destinations

Related concepts:

“System identifiers (SYSIDs)” on page 687

“Multiple Systems Coupling (MSC) directed routing” on page 693

Related reference:

🔗 TM and MSC Message Routing and Control User exit routine (DFSMSCE0) (Exit Routines)

Input, destination, and intermediate systems

Depending on the message, an IMS can be either an input system, a destination system, or an intermediate system.

For example, in the following figure, the transaction message originates from LTERMB in IMS B and is routed over path BD to TRAND in IMS D. IMS B is the input system and IMS D is the destination system. This message is called a primary message. When the message is processed by the application program in IMS D and a reply is returned to the input terminal in IMS B, this reply message is called a response. For this response message, the destination and input system is IMS B. The destination of the message is LTERMB in IMS B. The input that caused this response message is also from LTERMB in IMS B, which is the primary transaction.

If a message is sent through one IMS and routed directly to another IMS for processing, the routing system is called the intermediate system. For example, if TRAND is sent from IMS B through IMS A to IMS D (path BAD), IMS A is the intermediate system and IMS D is the destination system. Similarly, IMS A and IMS D are intermediate systems for TRANC when TRANC is sent from IMS B through IMS D through IMS A to IMS C (path BDAC).

Remote destination names must be defined as remote in the input system and as local in the destination system. Intermediate systems, however, do not need to define the input, destination name, or the path back (a local SYSID) for the routed message. Only the path to the destination system needs to be defined. The path back might not even be through the intermediate system. Input and destination name checking is not performed on a message when it is routed through an intermediate system.




Figure 114. Input, destination, and intermediate systems

System identifiers (SYSIDs)

MSC uses system identifiers (SYSIDs), two-byte numbers between 1 - 2036, to identify the IMS systems in an MSC network.

Each IMS system in an MSC network is assigned one or more SYSIDs. The SYSIDs are *local* to the IMS system that they are assigned to and *remote* to the other IMS systems in the MSC network.

The local SYSIDs must be unique within an MSC network. If an MSC network connects to an IMSplex with shared queues, the IMS systems in the shared queues group can share local SYSIDs.

A local SYSID and a remote SYSID are paired to define a *logical link path*. A logical link path is assigned to a logical link between the two IMS systems that the local and remote SYSIDs represent. The term *MSNAME* is often used as a synonym for the term logical link path.

As messages travel through an MSC network, the local SYSID identifies the point of origin of the message and the remote SYSIDs identify the destination of the message. Generally, the IMS system identified by the remote SYSID is where a transaction message is processed and the local SYSID is where the response to the transaction must be returned. However, a remote SYSID can also represent an intermediate IMS system that doesn't process the transaction, but instead passes it to another IMS system in the network.

Local SYSIDs are implicitly defined when MSNAME macros specify the paths to other IMS systems. That is, the local SYSID value of the SYSID parameter on the MSNAME macro determines which local SYSIDs are defined to an MSC system. Paths to other IMS systems are determined by the remote SYSID value on the MSNAME macro SYSID parameter.

You can change SYSIDs online by using the IMS type-2 command UPDATE MSNAME NAME(*msname*) SET(SIDL|SIDR).

In the following figure, IMS B has local SYSIDs 1, 2 and 3. IMS A has local SYSID 4. IMS C has local SYSID 5. IMS D has local SYSIDs 6 and 7. IMS B has local SYSIDs 1, 2, and 3 because of the three MSNAME macros with local SYSIDs of 1, 2, and 3.

In the figure, IMS B has three paths to remote SYSIDs 6 (IMS D), 7 (IMS D), and 5 (IMS C). IMS D and IMS B cannot route messages to IMS A, because they do not have paths to SYSID 4 (IMS A). IMS A does, however, have a path to SYSID 7 (IMS D) and SYSID 2 (IMS B). IMS A cannot send messages that originate in IMS A to SYSID 7 or SYSID 2, because the source SYSID (SYSID 4) is not recognized by IMS B or IMS D. In this configuration, IMS A can only function as an intermediate system for IMS B and IMS D. Another path (MSNAME) must be defined for IMS B to communicate with IMS A and IMS D.




Figure 115. System identifiers (SYSIDs)

Related concepts:

"MSC logical link paths" on page 679

Related reference:

➡ MSNAME macro (System Definition)

Routing messages with the destination name and SYSIDs

Messages in an MSC network contain information that makes it possible to route the message between IMS systems.

When a remote transaction is entered from an LTERM, a transaction message is built and queued on the message queue. This message contains information that is needed to route the message to its remote destination:

- Remote destination name (transaction code)
- Local or source LTERM name of the LTERM that entered the transaction
- Remote SYSID value
- Local or source SYSID value

If the transaction is processed by the application program and secondary messages are sent to other transactions (program-to-program switches), these messages have the destination name and SYSID of the switched-to transaction. The source (origin) name and SYSID remain the same. That is, the source SYSID and name never change. This facilitates sending the response message back to the input LTERM, regardless of where the transaction is processed.

Any IMS that locally processes a message that is received from a remote system must have a local SYSID defined to it that is the same as the remote SYSID of the message. It must have a path back to the source (origin) of the message. In the following figure, if TRANA is entered from LTERMB in IMS B, it is sent across the path MSNAME=(4,2) to IMS A. The destination name and SYSID of the message are TRANA and 4. The source name and SYSID are LTERMB and 2. IMS A accepts the message and processes it. It has SYSID=4 defined as local and has a path back to IMS B with the destination SYSID=2. If TRANA in IMS A issues a program-to-program switch to TRAND in IMS D, the destination name and SYSID are TRAND and 5. The source name and SYSID remain LTERMB and 2. IMS D accepts and processes the transaction. The transaction has SYSID=5 defined locally and has a path back to IMS B with destination SYSID=2.

If TRAND in IMS D sends a response back to input LTERMB, the response message has a destination name and SYSID of LTERMB and 2, and the source name and SYSID are also LTERMB and 2.

The input LTERM, LTERMB, is never defined as remote in IMS A or IMS D, yet the response message is returned to LTERMB in IMS B. IMS carries the originating LTERM name and SYSID in the primary message and all secondary messages. IMS knows where to route the response message when the application program responds to the input LTERM.




Figure 116. Message routing

Remote LTERMs

A *remote LTERM* is a logical terminal that does not reside on the local system.

Recommendation: Define the input LTERM as a remote LTERM only to send message switches, remote broadcasts, or messages from an application program to an alternate remote terminal (alternate PCB).

You can define remote LTERMs by using ETO MSC descriptors. The ETO MSC descriptor relates remote LTERMs to statically defined MSC links.

To help return response messages to the input (source or origin) LTERM, IMS carries the source LTERM name and SYSID in the remote message. To send a response message to an LTERM other than the source LTERM, you must define a remote LTERM. In the remote system, define NAME macros that have the name of the local terminal in the local system. Associate the NAME macro with the MSNAME that defines the destination SYSID of the local system.

For example, in the following figure, IMS A has local LTERMA, IMS B has local LTERMB, and IMS D has local LTERMD. IMS B can send message switches and remote broadcasts from LTERMB to LTERMD because LTERMD is defined remotely in IMS B. LTERMD is not actually defined with SYSID=(5,2), but it assumes those SYSIDs when the message is issued. LTERMD is associated with MSNAME BAD, which is defined with SYSID=(5,2) in IMS B. This association is established by placing the NAME macro for LTERMD after the following MSNAME macro:

```
BAD MSNAME SYSID=(5,2)
      NAME LTERMB
```

This naming differs from remote transaction definitions, which are explicitly defined with remote and local SYSIDs.

Also, LTERMA in IMS A can send message switches or remote broadcasts to LTERMD in IMS D. Similarly, LTERMD in IMS D can send message switches or remote broadcasts to LTERMA in IMS A. Both IMS A and IMS D have remote LTERM specifications for the other system's local LTERM.

TRAND in IMS D can send alternate messages to LTERMB in IMS B or to LTERMA in IMS A. MSNAMEs are defined with names for the purpose of:

- Displaying the queue counts for the named logical link path
- Stopping the sending of all messages from a terminal except those continuing a conversation
- Starting the logical link path
- Purging the logical link path in an MSC network for which input is stopped
- Allowing application programs to use directed routing
- Reassigning them to different links by using either the type-1 command `/MSASSIGN MSNAME msname LINK link#` or the type-2 command `UPDATE MSNAME NAME(msname) SET(MSLINK(linkname))`

In the following figure, the name of each MSNAME is the two-character name preceding MSNAME.




Figure 117. Remote LTERMs

Related reference:

➡ MSC descriptor format and parameters (System Definition)

Multiple Systems Coupling (MSC) directed routing

MSC directed routing is a function of MSC that allows an application program to specify the IMS name (MSNAME) and destination within that IMS for a message to an LTERM or an application program.

The receiving application program can determine the MSNAME of the IMS that originally scheduled it. With directed routing, the specified remote destination (a transaction or an LTERM) in another IMS does not need to be declared explicitly in

the IMS system definition for the sending IMS. These logical (local) names for terminals enable different IMS systems in the MSC network to use the same logical names for terminals and transaction codes. Names must still be unique within a given IMS. The Multiple Systems Verification utility (DFSUMSV0) cannot detect errors associated with MSC directed routing.

Restrictions:

- MSC directed routing does not support a program-to-program switch between conversational transactions.
- MSC directed routing does not support a program-to-program switch from a nonconversational transaction to a conversational transaction. For example, a conversational transaction in System A cannot use directed routing to perform a program-to-program switch to invoke a conversational transaction in System B.
- Response mode cannot be propagated on a DL/I ISRT call in a directed routing transaction.

Related concepts:

 Communicating with other IMS TM systems using Multiple Systems Coupling (Application Programming)

Remote destination verification

To maintain system integrity and prevent errors, an IMS in an MSC network verifies all specified destinations, unless MSC directed routing is used.

When MSC directed routing is used, IMS only ensures that a program-to-program switch is not being performed from a nonconversational transaction to a conversational transaction. Remote destination verification occurs when a message is received from a terminal or on receipt of an application program reply if a remote destination is specified for the message. Destination verification occurs as follows:

Destination

Verified For:

LTERM

Destination type: The original destination must have been a logical terminal.

Transaction

Destination type: The original destination must have been a transaction.

Transaction attributes: The following attributes must be consistent in the transaction definitions in the input and destination systems:

- Single segment or multiple segment
- Recoverable or irrecoverable
- Conversational or nonconversational
- Inquiry or update

The SPA size for a conversation from a remote system cannot be changed except by the remote system on an insert.

When an invalid destination is recognized, IMS cancels the message, sends an error message to the input terminal and to the master terminal of the local system, and logs an invalid request. If the message is conversational, the Conversation Abnormal Termination exit routine (DFSCONE0) is called in the input system, and the conversation is terminated.

Chapter 40. Administering Multiple Systems Coupling

The following topics describe the system administration activities required when you connect two or more IMS online systems in a network using MSC.

LU 6.2 application programs can process transactions on remote IMS systems. ETO also supports remote LTERMs.

Related concepts:

“Remote LTERMs” on page 691

Design considerations for multiple systems

The major design goals for MSC are to minimize resource consumption by defining suitable connections between the systems, balance resource demand by distributing functions among systems to obtain acceptable performance, and to program design considerations for multisystem conversational transactions.

The design and tuning recommendations that apply to a single IMS system are applicable to each IMS system in an MSC environment. Resource demand and consumption are taken into account when defining systems that are part of an MSC configuration.

An IMS transaction that is processed in a local system uses the same hardware and software resources that it would in a non-MSC environment. Transactions that are processed in a remote system require additional resources. In addition to resources used to transmit the transaction over physical links to the remote processor and to receive the response from the remote processor, resources are needed for message queuing and logging. Performance considerations are directly related to minimizing the resources consumed by remote processing and balancing the resource demand between several processors in an MSC configuration.

Minimizing resource consumption

You can minimize resource consumption in a number of ways when you design an MSC system.

To minimize resource consumption, do each of the following:

- Design the environment so that as many transactions as possible are processed locally.
- Provide physical links that go directly from local to remote systems; no intermediate systems should be involved in the transaction routing process. Transactions that must be routed through intermediate systems require additional processor activity, message queue activity, and logging activity.
- Design the message queue buffer pool in each processor to eliminate unnecessary message queue I/O activity.
- Design the queue buffer size to be large enough to hold the transaction input message and output response message in a single queue buffer.
- Define the physical link buffer sizes large enough to hold the message prefix plus all the segments of most messages. The physical link buffer size defines the default buffer size for all logical links assigned to a physical link. You can specify different buffer sizes for each logical link by using the UPDATE MSLINK

NAME(*linkname*) SET(BUFSIZE(*new_bufsize*)); however, the specifications for the buffer sizes for any given logical link must be the same in both of the IMS systems that the logical link connects.

Controlling the bandwidth of MSC links

You can control the bandwidth of CTC, MTM, and VTAM MSC link types by increasing MSC link buffer sizes and turning MSC bandwidth mode on and off.

When bandwidth mode is off, MSC sends a maximum of one message or response per I/O send or write operation. When bandwidth mode is on, IMS consolidates the following messages and responses in the same buffer:

- All messages that are queued and ready to be sent
- Any responses that are owed for messages received

Increasing the link buffer size allows more messages and responses to be sent simultaneously. Valid MSC link buffer sizes are 1024 bytes to 65536 bytes. The physical link defines the default buffer size for logical links associated with it. You can set the default value during system definition by specifying the BUFSIZE keyword on the MSPLINK macro. You can then modify the buffer sizes for each logical link in an online IMS system by using the UPDATE MSLINK command; however for any one logical link you must specify the same buffer size in both IMS systems at each end of the logical link.

To change bandwidth mode and link buffer sizes, use either the type-2 command UPDATE MSLINK or the type-1 command /UPDATE MSLINK. To display the status of bandwidth mode and the size of the buffers, use the type-2 command QUERY MSLINK or the type-1 command /DISPLAY LINK OPTION BUFSIZE.

You cannot set bandwidth mode during system definition. It must be set with a command. After you establish a bandwidth mode and buffer size, the mode and size selected remain in effect across link restarts and IMS warm starts. Bandwidth mode is set off and the buffer size is restored to the SYSDEF value across an IMS cold start.

Related reference:

-  MSPLINK macro (System Definition)
-  UPDATE MSLINK command (Commands)
-  QUERY MSLINK command (Commands)

Balancing resource demand

In an MSC environment with two or more processors, try to distribute the workload in a way that avoids excessive use of any one processor.

You can distribute the workload by distributing IMS applications and their associated transactions and terminals between the available processors. Depending on the complexity of the application and the capability of the processor, you can avoid overloading any single processor.

If the current design of the databases is such that the databases and their associated applications cannot be distributed across the available processors, you can:

- Duplicate inquiry-only databases; this allows more than one system to reference the whole of the database. (This is called vertical partitioning.)

- Split the databases into several component databases. (This is called horizontal partitioning.) The component databases must be completely independent for distribution among the available processors. For example, it might be possible to divide a database by key range intervals. The new databases and their associated applications can then be distributed among the existing IMS systems, and you can use the Terminal Routing exit routine to route incoming transactions to the correct IMS system. Another possibility is to divide the database by geographic area. Each IMS system could process the transactions that refer to the databases for its own geographic area and route transactions that refer to a remote geographic area.

In addition to balancing the workload across processors, you might also need to balance the workload on physical links. This occurs when a physical link between two systems is of the SDLC type and multiple physical links have been installed. You can balance the workload on physical links by:

- Specifying, during IMS system definition, proper logical link paths and logical links for each remote application.
- Using a user-written TM and MSC Message Routing and Control user exit routine (DFSMSCE0) to distribute the transaction load on each of the alternative physical links.

Planning for conversational processing

Conversational processing is available to terminals that are attached to any IMS in an MSC network to the same extent as if they were in a single-system environment.

The following differences apply to conversational processing in an MSC network:

- All transactions used in a conversation must be defined as conversational in each IMS of the MSC network.
- The input system controls the conversational resources for the duration of the conversation. When the input system receives a conversational transaction, it inserts the scratchpad area (SPA) as the first message segment, and routes the message to the destination application program.
- Any system in the MSC network can process any step in the conversation.
- Program-to-program switches can be routed from system to system.
- SPAs that are used in multisystem conversations must follow these rules:
 - For the SPA ISRT, conversational program-to-program switches can occur to a transaction with a SPA that is equal, larger, or smaller in size.
 - The minimum size of a SPA is 16 bytes (X'10'), and the maximum size is 32767 bytes (X'FFFF').

Generally, terminal operators and application programs are unaware of whether a conversation is multisystem.

Exceptions:

- Suppose a conversational program inserts a message to a response alternate PCB in a remote system. By implication, this destination is in the input system and is verified by the input system. Destination verification in this case involves confirming that the specified logical terminal is still assigned to the input terminal. If the logical terminal has been reassigned, the input system activates the Conversation Abnormal Termination exit routine, and terminates the conversation. The status code that is returned to the application program is

blank, indicating success. This status code is blank even when the actual result is unsuccessful due to the condition described.

- Suppose an application program executing in a system other than the input system uses the SPA to specify a transaction code, thereby passing conversation control to another program. If the specified transaction code is invalid, the input system activates the Conversation Abnormal Termination exit routine, and terminates the conversation. No status code is returned to the application program.

MSC support for APPC/IMS and OTMA includes the following IMS transaction types:

- Conversational
- Nonconversational
- Response mode
- Non-response mode

Restriction: MSC does not support Fast Path.

Routing exit routines with conversations

You can use the program routing and link routing entry points of the TM and MSC Message Routing and Control User exit routine (DFSMSCE0) to input messages that start a conversation.

It is not applicable at any other conversational step, because the application program, not the input terminal, provides the destination for continuation of the conversation.

Remote destination verification for conversations

Destinations for program-to-program switches are verified in the system in which the program requesting the switch executes, except where MSC directed routing is used.

If MSC directed routing is not used and the destination is valid, the system sends the SPA and the message to the destination transaction. If the destination is invalid, the system returns a status code to the application program.

If MSC directed routing is used, IMS ensures only that a program-to-program switch is not being performed from a nonconversational transaction to a conversational transaction. If the destination is valid, the system sends the SPA and the message to the destination transaction. If the destination is invalid, the system does not route the SPA and the conversation is terminated.

Destination verification for a message to the input terminal is performed by the input system. The specified logical terminal must still be assigned to the input terminal. The input system also verifies, except when MSC directed routing is used, the next transaction that is specified in the SPA. If the destination is invalid, the input system invokes the Conversation Abnormal Termination exit routine and terminates the conversation. No status code is returned to the application program.

Saving truncated data in the SPA

The SPA=STRUNC option on the TRANSACT macro applies to conversations that use program switches to other transactions that have different sized SPAs.

- If you use the SPA=STRUNC option, IMS preserves all of the data in the SPA, even when a program switch is made to a transaction that is defined with a smaller SPA. The transaction with the smaller SPA does not see the truncated data. However, when this transaction switches to a transaction with a larger SPA, the truncated data is used. IMS tracks the longest data that is inserted to the SPA to determine the truncated data length.

Example: If you have the three transactions:

```
TRANA SPA=100
TRANB SPA=50
TRANC SPA=150
```

If the application program for TRANA switches to TRANB, the last 50 bytes of the SPA for TRANA are not sent to TRANB. If the application program for TRANB subsequently switches to TRANC, the SPA that is received by TRANC contains the following 150 bytes:

- The first 50 bytes from the SPA that was inserted by TRANB
 - The second 50 bytes from the second 50 bytes that was inserted by TRANA
 - The third 50 bytes are binary zeros.
- If you do not use the SPA=STRUNC option, the truncated data is lost. In the previous example, the second 50 bytes that is received by TRANC would be binary zeros.

Restriction: Truncated SPA data from a previous transaction is lost if it is sent using MSC to an IMS 5.1 (or earlier) system.

Conversation termination

A conversation can be terminated by either the application program or the terminal operator.

An application program terminates a conversation by inserting a message to the input terminal with a SPA that contains either blanks as the transaction code or the transaction code of a nonconversational transaction.

A conversational transaction can also be terminated by the /EXIT command.

An IMS shutdown does not terminate conversations. The conversation is continued after IMS restarts, unless an IMS cold start occurs or the conversation is terminated by an operator command such as /EXIT. If the input system is shut down and subsequently cold starts, all the conversations that it controls are lost. If using Resource Manager (RM) and the status recovery mode is GLOBAL, an IMS conversation can survive an IMS cold start because the status is stored in RM. The input system cancels any conversational messages it receives for input terminals that were previously involved in active or held conversations. IMS retains conversations over a restart. An IMS cold start is an unusual procedure, and can cause many problems beyond losing conversation state.

If a remote system is shut down when a conversational step is processing or in its queue, and it is subsequently cold started, all references to the conversation are lost. A conversation that is lost in this way must be specifically canceled in the input system by the /EXIT command.

Abnormal conversation termination

Several different events can lead to the abnormal termination of a conversation.

A conversation is abnormally terminated if any one of the following events occur:

- The conversational application program abnormally terminates.
- An invalid destination is recognized in the input system or in the remote system (for a conversation response, a program-to-program switch, or in the SPA).
- A conversational message is inserted into a terminal whose conversation was terminated.
- Destination verification fails for a conversational message.
- No output was generated in the application program.

The conversation's SPA (along with an indication of the cause of termination) is passed to the Conversation Abnormal Termination exit routine in the input system.

Defining Multiple Systems Coupling resources

To generate an IMS online system that includes MSC facilities, your IMS system definition must include three macros: MSPLINK, MSLINK, and MSNAME.

You also need to define transaction codes for each IMS system that has any part in transaction entry or processing. You can define transaction codes dynamically or by using system definition macros. An individual system can play several roles: it can be the input system, it can be an intermediate system responsible for routing transactions, or it can be a destination system in which the transaction is processed.

Using MFS is the same in an MSC network as in a single-system environment. If a message is created in one IMS for a terminal that is attached to another IMS, the required message and format descriptions must be available to the IMS to which the terminal is attached, and definitions with the same name must be defined identically in each IMS.

Local system definitions

When you define an MSC network, you need to define transactions, logical terminals, and physical and logical connections in each local system.

More specifically, for each local system you need to define:

- All transactions entered or processed by that system
- All logical terminals that are attached to that system and all logical terminals in remote systems that are referenced by transactions processed in that system, or by terminal operators, unless either a Program Routing exit routine is used or MSC directed routing is used
- The physical and logical connections between that system and the remote systems that share in the processing of the specified transactions

The system definition macros that you need to prepare in each system are summarized in the following table. The first macro you use is IMSCTRL. Using the MSVID keyword, you assign a number in the range from 1 to 255 as a unique identifier of that system. This causes a control block, DFSMSxxx, to be built for use with the Multiple Systems Verification utility. (The 3-digit suffix, xxx, matches the MSVID parameter; one unique control block exists for each system in the MSC

network.) This offline utility helps verify that system definitions for all partner systems are consistent. The IMSMSV procedure is generated in IMS.PROCLIB in order to execute this utility.

Resource	Identification	Macro	Number coded
System	System ID	IMSCTRL	1
Programs	PSB name	APPLCTN	1 per PSB
Local transactions	Transaction code	TRANSACT	1 or more per APPLCTN
Remote transactions	Transaction code	TRANSACT	1 or more per APPLCTN in a remote system
Physical connections	Physical link name	MSPLINK	1 per connection
Logical connections	Logical link name	MSLINK	1 per MSPLINK ¹
Routing names	Logical path name	MSNAME	1 or more per MSLINK
Local terminals	LTERM name	NAME	1 per local terminal
Remote terminals	LTERM name	NAME ²	1 per remote terminal

Notes:

1. Multiple-session SDLC links that are VTAM have additional MSLINK macros.
2. Define with an ETO MSC descriptor if the corresponding local terminal is an ETO terminal.

Related reference:

 [Macros used in IMS environments \(System Definition\)](#)

 [Multiple Systems Verification utility \(DFSUMSV0\) \(System Utilities\)](#)

Defining partner systems

A relationship exists between the three system definition macros (MSPLINK, MSLINK, and MSNAME) that define the names of connections between systems and the logical names that are used in commands.

Defining the physical link

You can define different types of physical links to connect IMS systems. Your choice depends on the hardware that is required for each of the links.

The choices are:

- Channel-to-channel (CTC)
- Memory-to-memory (MTM)
- VTAM

You use the MSPLINK macro to declare which type of physical connection you are going to use. You must declare all physical connections that can potentially be used, even if several might be for backup purposes or might not be intended for continual use. You must declare the physical link for partner systems in both system definitions. Using the TYPE keyword, declare the kind of physical link to use.

You must declare the physical link for partner systems in both system definitions. Using the TYPE keyword, declare the kind of physical link to use.

Assign a name to the physical link. The name of the physical link is used for logical link definition to match a connection between systems with the physical device or transmission technique that is to be used. The name is also used to identify the link when issuing commands.

If part of a conversation is to be passed from one system to another, use the maximum scratchpad area (SPA) size, if that SPA size exceeds the message segment size.

You can change the attributes of existing physical links by using the type-2 UPDATE MSPLINK command. The changes remain in effect only until the next cold start of the IMS system.

Related reference:

-  MSPLINK macro (System Definition)
-  Macros used in IMS environments (System Definition)
-  QUERY MSPLINK command (Commands)
-  UPDATE MSLINK command (Commands)
-  UPDATE MSPLINK command (Commands)

Buffer sizes for physical links

The physical link buffer size defines the default buffer size for all logical links assigned to the physical link.

Use at least the size of the largest message segment that is to be transmitted across this physical link. The buffer sizes at each end of a physical link must be equal.

Specify a buffer size with the BUFSIZE keyword on the MSPLINK macro.

You can specify different buffer sizes for individual logical links by using either the IMS type-2 command UPDATE MSLINK NAME(*linkname*) SET(BUFSIZE(*new_bufsize*)) or the IMS type-1 command /UPDATE MSLINK NAME(*linkname*) SET(BUFSIZE(*new_bufsize*)).

See the description of the BUFSIZE keyword for the minimum and maximum sizes that you can specify for link buffers.

You can display MSC statistics to help determine the optimum buffer size. MSC statistics are displayed by the IMS type-2 command QUERY MSLINK NAME(*linkname*) SHOW(STATISTICS). Use the statistics Hi_Msg_Send_SZ and Hi_Msg_Rec_SZ to see the maximum message sizes being sent and received. You can also compare the number of send and receive I/O requests to the number of messages that are sent and received. If IMS requires multiple send and receive I/O requests to send a single message, the buffers are too small. The send and receive I/O requests are captured by the statistics Tot_Send_CT and Tot_Rec_CT. The number of messages sent and received are captured by the statistics Tot_Msg_Send_CT and Tot_Msg_Rec_CT.

Depending on the physical link type, storage for the link buffer sizes are allocated from different pools. VTAM link buffers are allocated from the high input/output

pool (HIOP). CTC link buffers are allocated from the communications input/output pool (CIOP). MTM link buffers are allocated from subpool 23 in the common storage area (CSA).

Related reference:

- [➞ MSPLINK macro \(System Definition\)](#)
- [➞ Macros used in IMS environments \(System Definition\)](#)
- [➞ QUERY MSLINK command \(Commands\)](#)
- [➞ UPDATE MSLINK command \(Commands\)](#)
- [➞ UPDATE MSPLINK command \(Commands\)](#)

Defining a CTC physical link

The CTC adapter is a channel-to-channel hardware connection.

CTC links are usually used only when the IMS systems are in the same data center. You can assign only one logical link to a physical link that uses the CTC connection type.

For CTC links, you must add a DD name and an address parameter value.

Also, during system definition, you must link module DFSCMC10 to the LPALIB or the MLPA library as IGG019xx, where xx is the CTC appendage suffix that is defined by the APNDG parameter in the IMSCTF macro for z/OS systems.

Related concepts:

- [➞ Installing required IMS links to z/OS \(System Administration\)](#)

Related reference:

- [➞ MSPLINK macro \(System Definition\)](#)
- [➞ Macros used in IMS environments \(System Definition\)](#)
- [➞ IMSCTF macro \(System Definition\)](#)
- [➞ QUERY MSPLINK command \(Commands\)](#)
- [➞ UPDATE MSLINK command \(Commands\)](#)
- [➞ UPDATE MSPLINK command \(Commands\)](#)

Defining a MTM physical link

The MTM link is a software link between IMS subsystems that are running in the same logical partition.

You can assign only one logical link to a physical link that uses the MTM connection type.

Related reference:

- [➞ MSPLINK macro \(System Definition\)](#)
- [➞ Macros used in IMS environments \(System Definition\)](#)
- [➞ QUERY MSPLINK command \(Commands\)](#)
- [➞ UPDATE MSLINK command \(Commands\)](#)

Defining a VTAM physical link

A VTAM physical link is supported by SNA VTAM.

For physical links that use VTAM, you must supply a NAME keyword that matches the label on the VTAM APPL statement for the remote system. The NAME keyword and the VTAM APPL statement both specify the VTAM node name.

The label on the VTAM APPL statement also serves as a default value for the ACBNAME parameter on the same statement. Regardless of whether you use the default value for the ACBNAME parameter, this value must match the APPLID parameter on the IMS COMM macro.

The SESSION parameter indicates the number of parallel sessions that can be active for the physical link.

Related reference:

-  [MSPLINK macro \(System Definition\)](#)
-  [Macros used in IMS environments \(System Definition\)](#)
-  [QUERY MSPLINK command \(Commands\)](#)
-  [UPDATE MSLINK command \(Commands\)](#)
-  [UPDATE MSPLINK command \(Commands\)](#)

Defining the logical link

The logical link definition is defined with the MSLINK macro, which, among other options, allows you to name the link, to associate the logical link with a logical link defined in a partner system, and to define the type of physical link the logical link can be used with.

You can assign a name to the logical link by using the macro label field. Logical link names are used by the type-2 commands QUERY MSLINK or UPDATE MSLINK to identify the target logical link. If you do not define a logical link name, IMS assigns the default logical link name DFSLxxxx, where xxxx is the logical link number.

You associate the logical link with the logical link defined in a partner system by specifying a two-character alphanumeric partner ID on the PARTNER keyword. The partner ID that you specify must match a corresponding partner ID specified in an MSLINK macro in the partner system. The matching values of the PARTNER keyword represent a logical connection between the two partner systems. For example, if B1 is defined as the partner ID in one logical link, this same B1 partner ID must be specified in the corresponding MSLINK macro of the partner system.

You match the type of physical connection that can potentially be used with this logical link by specifying the MSPLINK keyword. Depending on the type of physical link you specify, you can assign one or more MSLINK macros to the physical link:

- When the physical link type is CTC or MTM, one MSLINK macro exists for each kind of physical connection.
- When the physical link type is SDLC that was VTAM parallel sessions, multiple logical links can be assigned to one physical link. For example, System_B and System_C can have two logical links, one using MTM and the other CTC; System_C and System_D can have two logical links, both of which use one VTAM physical link.

You can define a maximum of 999 logical links for each IMS system, unless the IMS system is defined as a large system by specifying `SYSTEM=LGEN` in the `IMSCTL` macro. When `SYSTEM=LGEN` is specified, you can define only 676 logical links.

You can change logical link attributes that are defined by the `MSLINK` macro by using the `UPDATE MSLINK` command. The attributes you can change include the name, the physical link assignment, the partner ID, the buffer size, and others. The changes remain in effect until the next cold start of the IMS system.

Related reference:

➞ [UPDATE MSLINK command \(Commands\)](#)

➞ [MSLINK macro \(System Definition\)](#)

Defining a logical path

When the operator connects two systems, a choice of physical connection type might be possible, therefore the logical link name is used. The `MSNAME` macro lets you define that logical path name.

The `SYSID` keyword is used to declare the two systems that are joined in a pathway. You select the 1-digit identifiers from the range 1 to 2036, one for the remote system and the other for the local system. For example, (1,3) specifies that any message using this path is being sent to the remote system number 1, and that the local system number is 3.

You can change the remote and local system IDs of a logical path by using the type-2 command `UPDATE MSNAME`. The changes will remain in effect only until the next cold start of the IMS system.

The `MSNAME` macro can be followed by a set of `NAME` macros on which you can specify the `LTERM` names for terminals that are in remote systems. You do not need to declare every terminal in the remote system that is entering transactions, but only those that enter traffic destined for this local system. If the `LTERM` in the remote system is for an `ETO` terminal that enters transactions destined for this system, define the `LTERMs` using `ETO MSC` descriptors instead of `NAME` macros.

Setting link priorities for remote transactions

You can assign priorities to remote transactions using the `PRTY=` keyword on the `TRANSACT` macro during system definition in the destination system.

You might want to assign priorities because a remote input transaction might have a different scheduling priority in the destination system from that defined for it during system definition of the originating system. This means there could be a long wait for a response, even though the remote transaction had a high priority when it was defined with the `SYSID=` parameter of the `TRANSACT` macro.

For example, in the following figure, Terminal 1 is sending a remote transaction (`ASMB1`) to System B.




Figure 118. Link priorities for a remote transaction

In IMS system A, this transaction is enqueued to a control block (CB) and has a priority of 9 (defined with the SYSID= parameter). When ASMB1 is sent to system B, it becomes a local transaction and is processed. When the message router is called to send a response, it enqueues the response on the control block that represents the MSC link. If ASMB1 is defined in system B with a priority of 7, all transactions with priorities of 12-15 (in queue 1) are processed first, and 8-11 (in queue 2) are processed next. The ASMB1 (with a priority of 7 in queue 3) is processed next. However, as shown in the figure, the limit priority can be set at 8 and the limit count can be set at 3. As a result, when the number of message queues to be processed reaches 3, the priority of ASMB1 is changed to 8, and subsequent responses are placed on queue 2. On system B, all messages for BSMB1 are sent before responses are sent. Messages for BSMB2 are sent last in this example.

In a shared queues environment, however, priority only applies to messages received from the MSC link to be processed. Priority does not apply to response

messages (messages going back to the input system) because the coupling facility does not have the four queues. Instead, all response messages are sent FIFO.

Serial transaction processing in an MSC network

Serial transactions are processed in the order in which they are received relative to other transactions of the same type.

You can ensure serial processing of transactions in remote MSC IMS systems by taking the following steps:

- Define the transaction as serial in both the local and remote MSC IMS systems.
- Restrict the transactions to a single logical link path between the local and remote MSC IMS systems.
- Send all serial transactions of the same type to the same remote MSC IMS for processing.

Serialization is not preserved for transactions that:

- Are sent across different logical link paths
- Originate from different MSC IMS systems
- Are processed in different remote MSC IMS systems

The PRTY= keyword and output messages of serial transactions in an MSC network

The serial processing of the output messages of a serial transaction can become unpredictable in an MSC network if the *normal* and *limit* parameters of the PRTY= keyword are not equal in the TRANSACT macro.

If the *normal* and *limit* parameters are not equal and the number of output messages on a message queue becomes equal to or greater than the value of the *limit_count* parameter of the PRTY= keyword, output messages received by a remote MSC IMS system after the *limit_count* value has been reached might be processed before messages received by the remote MSC IMS system.

Related Reading: For more information on the PRTY= parameter of the TRANSACT macro, see *IMS Version 11 System Definition*.

Specifying exit routines

The following topics describe how to specify exit routines used with MSC.

Related concepts:

“TM and MSC Message Routing and Control user exit routine overview” on page 735

Routing messages with DFSMSCE0

The TM and MSC Message Routing and Control User exit routine (DFSMSCE0) is not included in the system definition process. Therefore, incorporate DFSMSCE0 in IMS.SDFSRESL.

The terminal routing entry points are invoked in the input system, possibly changing the destination. When transactions arrive at an IMS across the MSC link, the link-receive entry points can be invoked. These entry points can change the destination again. Because destinations can be changed at so many different points, clearly document how the destinations are changed, and explain the reasons for the new destinations, which are used for processing. You might need to correlate

the transaction code changes as information for master terminal operators so that they can interpret a display of queue status.

If the DFSMSCE0 exit routine is used in an IMSplex environment with shared queues, the DFSMSCE0 exit routine can be used in the IMS systems within the IMSplex for affinity routing, which establishes an affinity between a transaction message and a back-end IMS system. Affinity routing effectively suppresses the sharing of the transaction message on the shared queues so that only IMS systems that have affinity are notified when the transaction arrives on the shared queues.

DFSMSCE0 IMSplex affinity routing is also a way to route APPC synchronous transaction or an OTMA send-then-commit (CM1) transaction to back-end IMS systems in an IMSplex for processing when either RRS=N or AOS=N. Without assigning affinity, these APPC and OTMA transactions cannot be routed to back-end systems for processing. MSC must be enabled in the IMSplex to use DFSMSCE0 IMSplex affinity routing with APPC synchronous or OTMA send-then-commit (CM1) transactions.

Be aware of how transactions can be routed for different stages in the total processing. If input from IMS A is being processed in IMS C, the processing program can invoke the exit routine to determine whether to send its output to the original input terminal or to another location. Your documentation of the exit routine should show the patterns of alternative message destination. An end user could be expecting output and be unaware that it is being sent to another component.

Related concepts:

“TM and MSC Message Routing and Control user exit routine overview” on page 735

Related reference:

 TM and MSC Message Routing and Control User exit routine (DFSMSCE0) (Exit Routines)

Managing error messages with DFSCMUX0

The Message Control/Error exit routine (DFSCMUX0) allows you to manage and control messages that are in error on an MSC link or on a message queue. It is activated by an MSC link start, link termination, send error, or receive errors.

The Message Control/Error exit routine is called for these situations for APPC/IMS:

- If an LU 6.2 session fails while sending an output message to an LU 6.2 program
- If a send to an LU 6.2 application program is rejected with a Deallocate with Send_Error message
- When /DEQUEUE *luname tpname* is entered

You can customize the Message Control/Error exit routine to ask IMS to take any of these actions:

- Take the default action: discard the message and proceed with the /DEQUEUE command.
- Discard the message in error.
- Discard the message in error and notify the MTO or originating terminal of this error.
- Re-route the message in error to a different local or remote transaction, a local or remote LTERM, or an LU 6.2 application program.

The sample exit routine that is provided in IMS.ADFSSRC uses the default action.

Related reference:

 Message Control/Error exit routine (DFSCMUX0) (Exit Routines)

How network definition is affected by multiple systems

Each system definition must have all terminals that are connected to it defined using the `TERMINAL` and `NAME` macros. However, if terminals in other systems are capable of sending messages that are destined for the defined system, those terminals are logically part of the network for that system.

Although MSC physical and logical links continue to be predefined through the system definition process, you can dynamically create MSC remote LTERMs during IMS initialization with Extended Terminal Facility (ETO). Although you cannot define your MSC links using ETO, ETO does allow you to associate one or more message queues with an MSC logical link.

When a transaction is processed in a remote system, the input LTERM name in the local system is carried over as part of the message. If the processing program uses the alternate PCB to direct a message to another terminal besides the input terminal, those destinations need to be declared as remote, unless directed routing is used. Define the LTERM names for all inputs with `NAME` macros. Position the `NAME` macros in a group after the `MSNAME` macro. You now have a set of LTERMs that collectively can occur in several system definition decks. For example, `TERMA` can be present in the input system, in the intermediate system, and in the processing system.

When planning for the network, keep in mind that message queues for an input system or an intermediate system must allow for the remote transactions being queued. When allocating space for the message queues take into account the message lengths and their expected loads. In a similar way, you must allow for the presence of these messages in I/O buffers, even when they are not going to be processed in that system.

Related tasks:

Chapter 6, “Administering the Extended Terminal Option,” on page 87

Verifying transaction definitions across systems

After you have completed system definitions for several systems that are to be part of a multiple-system network, the transaction codes, LTERM names, and system identification numbers are all identifiable in control blocks.

Using the multiple systems verification utility

You can use the IMS Multiple Systems Verification utility offline in order to verify that the names you have used are consistent across system definitions.

Example: If you use `TRANX` as a code in two definitions but `TRANXX` in a third system, the Multiple Systems Verification utility sends you a message highlighting the single reference to `TRANXX`.

You can execute the Multiple Systems Verification utility by using the `IMSMSV` procedure in `IMS.PROCLIB`.

The Multiple Systems Verification utility also checks to ensure that transaction codes are not defined as local in more than one system.

The Multiple Systems Verification utility uses input control statements that specify the system identification numbers that you want to include in the checking. Its output is in the form of a multisystem path map.

Restriction: The Multiple Systems Verification utility cannot detect errors caused by improper use of MSC directed routing.

Related reference:

 Multiple Systems Verification utility (DFSUMSV0) (System Utilities)

Verifying the system definition status online

You can check the consistency of your system definitions or ETO MSC descriptors using the /MSVERIFY command when each of the component systems of your network is operational.

Recommendation: Use the /MSVERIFY command during the system test phase, rather than during production. The /MSVERIFY command generates considerable traffic.

The /MSVERIFY command validates that:

- Logical terminals exist for the remote LTERMs defined.
- Transactions are defined and have the same attributes in their remote system as in the local system.
- Logical path definitions are consistent and usable.

Recommendation: Define your remote LTERMs using ETO MSC descriptors. Otherwise, the Multiple Systems Verification utility recognizes remote LTERMs, but not the corresponding local LTERM in the target system.

The /MSVERIFY command verifies consistent definition between the local system and one remote system for each command. Verification begins with local SYSIDs and then proceeds to remote SYSIDs. As each segment of the processing begins, the MTO is notified of the local and remote SYSIDs to be verified. This message, DFS2234I or DFS2236I, also is time-stamped. As the specific local and remote SYSIDs are verified, the MTO receives verification-completed notifications with corresponding time stamps. If a definition or assignment error is discovered, error messages are returned.

During command processing, the local system sends all its own locally defined MSC elements (SYSIDs) to the remote IMS system specified in the /MSVERIFY command. The remote system responds by returning all destinations and their attributes that belong to the local system (as the remote system definition or assignments have them described). The command input system checks local and remote definitions for consistency.

The response list should be checked for completeness using the listing of SYSIDs that are to be verified given in the DFS2234I or DFS2236I IN PROGRESS messages and the correspondence of time stamps in responses.

The local system then sends all MSC elements to the remote system expecting them to be defined in that remote system with corresponding attributes (as the local system has them defined). The remote system performs consistency checks and routes error messages to the command input terminal.

MSC path consistency is only checked for current operational assignments of logical link paths to logical links, and logical links to physical links.

Recommendation: Ensure that the MSC definition and assignments are accurate, even when receiving favorable responses (that return no errors). Definition and assignment errors can prevent the return of a command response for some SYSIDs. Regard the absence of a response for a particular SYSID as significant.

Reviewing these error responses from the /MSVERIFY command might help you to avoid these definition and assignment errors:

```
DFS2235I SYSID __ is defined as local in both systems
DFS2241I __ is defined as remote transaction in both systems
DFS2242I __ is not defined as LTERM in both systems
DFS2243I __ is not defined as transaction in both systems
DFS2245I Multisegment transaction flag for __ not consistent
DFS2246I Non-inquiry only flag for __ not consistent
DFS2247I Conversational flag for __ not consistent
DFS2248I Irrecoverable flag for __ not consistent
DFS2249I Fixed length SPA flag for __ is not consistent
DFS2250I The SPA length for __ is not the same
```

The /MSVERIFY command cannot detect errors caused by improper use of MSC directed routing.

Security considerations for MSC

For those transactions that are processed in another system, perform as much security checking as is required for the primary message. RACF can be used to protect IMS resources in an MSC network.

Signon verification, combined with transaction authorization and password checking, allows you to control the processing at input time. The resource definition to RACF must declare the transaction name, even when the transaction is not processed in the system where the security tables are built.

Security controls in an MSC network are performed independently in each local and remote IMS. An intermediate IMS in an MSC environment, which is neither local nor remote for a given transaction, does not perform any security checking for that transaction.

RACF can provide transaction authorization checking when a destination system receives a message to process. The amount of checking that RACF provides depends on the MSCSEC= parameter in the DFSDCxxx IMS.PROCLIB member and on feedback from the DFSMSCE0 exit routine. The DFSMSCE0 user exit can optionally override or accept the system DFSDCxxx member security, on a message by message basis.

Transactions received in a remote IMS on an MSC link are passed to the transaction authorization module for authorization checking, but because the password is not passed across the link, transaction authorization checking fails if a password is required. Transactions that do not require a password can be accepted.

To allow a transaction to be scheduled in a remote destination IMS, you can authorize its processing with resource access security (RAS). To use RAS security, the transaction must be defined to RACF as authorized for use by the dependent region.

If the RACF security environment is not available in the destination system (as when a /SIGN ON command is entered with RACF), the security environment will be dynamically created to allow the transaction authorization to proceed.

Related concepts:

 IMS security (System Administration)

Operations for Multiple Systems Coupling

Each system in an MSC network is operationally an independent unit. Each system exclusively owns its own communication resources, which are controlled by its own master terminal.

In many cases, the MSC operations performed in one IMS system require corresponding operations to be performed in a partner IMS system.

Common operational tasks include starting and stopping communications on MSC links, modifying link assignments, and recovery.

Related concepts:

 MSC operations (Operations and Automation)

MSC link statistics

IMS maintains statistics for each MSC logical link. You can use link statistics for tuning MSC systems to process MSC messages more efficiently.

IMS records three types of MSC logical link statistics:

General statistics

Statistics start time, MSC ITASK dispatch counts, MSC ITASK processing times, and the rate and number of logger check writes

Send statistics

Messages sent, byte count sent, send message sizes, queue manager get counts and times, and send I/O times

Receive statistics

Messages received, byte count received, receive message sizes, QMGR insert counts and times, and receive I/O times

The I/O times for sends and receives are kept only when MSC bandwidth mode is enabled. If MSC bandwidth mode is disabled, the values for the I/O statistics are zero.

You can view the current statistics for a logical link by issuing the type-2 command `QUERY MSLINK NAME(linkname) SHOW(STATISTICS)`.

IMS maintains two copies of the statistics for each logical link: the running cumulative totals since the last start or restart of IMS and the running cumulative totals since the last checkpoint or manual reset of the link statistics. The command `QUERY MSLINK NAME(linkname) SHOW(STATISTICS)` shows only the statistics since the last checkpoint or manual reset, unless the link statistics have not been reset since the last restart of IMS. Otherwise, you can see the running totals since the last IMS restart only by looking at the X'4513 ' log records.

You can reset the statistics for a link manually or have IMS automatically reset the statistics for a link at each system checkpoint. Resetting link statistics can be useful

when running performance tests and comparing results. To manually reset the statistics for a link or to turn on or turn off automatic checkpoint reset for link statistics, issue the type-2 command `UPDATE MSLINK NAME(linkname) START(STATISTICS) OPTION(reset_option)` with the appropriate reset option. The logical link does not need to be stopped or idle to reset link statistics or to change from automatic to manual reset mode and back again.

Resetting MSC logical link statistics does not affect the logging of MSC logical link statistics. MSC logical link statistics are always logged and cannot be reset.

Related tasks:

 Diagnosing link problems by using MSC link statistics (Diagnosis)

Benchmark link activity

You can benchmark link processing by documenting what the normal and high levels of message activity are for your MSC logical links by querying the link statistics.

If you document link activity levels, you can then use them to help recognize problems when current activity levels differ from the documented levels, or you can compare previous levels to current levels to see if more logical links or network capacity is needed to ease capacity limitations or bottlenecks.

Determine your optimum MSC link type

You can determine the specific resource requirements of a message by resetting the logical link statistics, sending one typical message, and recording the statistics.

A single message represents an MSC unit of work, and the resulting statistics for the single message represent the resources an MSC unit of work require, such as time, calls, and so forth.

You can also determine the average resource requirements for MSC units of work by sending multiple messages and then dividing the resulting cumulative statistics by the number of messages sent.

After you know how much resources a unit of work requires on a specific link type, if your installation uses multiple physical link types (that is, CTC, MTM, and VTAM), you can compare the resource requirements of units of work on each type of link, and then use the link type that utilizes the least amount of resources.

Reset statistics regularly at system checkpoint

Unless you have a specific reason for resetting MSC logical link statistics manually, you should configure IMS to reset link statistics regularly at system checkpoint by specifying a reset option of `RESET,CHKPT` on the type-2 command `UPDATE MSLINK`.

Automatic reset at system checkpoint is the default reset option that is in effect when IMS restarts.

The longer the interval between resets, the less useful the statistics become for determining problems and other analyses. If you need to run a benchmark or gather statistics for an interval longer than the interval between IMS system checkpoints, you can turn off the automatic reset at checkpoints by issuing the `UPDATE MSLINK` command with the `NORESET,CHKPT` parameters.

When automatic resets at checkpoint are turned off, IMS does not reset link statistics, and you must manually reset the statistics to start a new recording interval. Statistics are manually reset by issuing the type-2 command `UPDATE MSLINK NAME(linkname) START(STATISTICS) OPTION(RESET)`.

Adjust link buffer sizes to the size of the messages

Use the MSC logical link statistics to determine what the messages sizes are for an MSC link and adjust the buffer sizes for the link and message queues so that most messages can be received with a single I/O operation and two QGET (one GU and one DEQ) calls.

The high, low, and average message size statistics returned by the type-2 command `QUERY MSLINK` are derived from the total sizes of the messages in the send and receive message queues, including the prefix data and user data segments, as they are logged in the X'01' and X'03' log records.

Adjust the buffer sizes so that IMS uses no more than a single I/O operation and two QMGR calls to either send or receive a message on an MSC link: two QGET calls (one GU and one DEQ) and two QPUT calls (one ISRT and one ENQ).

The QPUT calls on the sending side of an MSC link operate differently from the QGET calls on the receiving side of an MSC link. Each QPUT call is for a message segment, and each QGET GU call is for a message queue buffer. A message with two user segments in one queue buffer requires three QPUT calls (two ISRT calls and one ENQ call) and two QGET calls (one GU call and one DEQ call). For single segment messages, IMS uses the same number of QPUT and QGET calls on each side of an MSC link.

Adjust logical link capacity for MSC bandwidth mode

In MSC bandwidth mode, IMS can process multiple messages in single send or receive action. When IMS processes multiple messages at a time, the total number of sends and receives is less than the total number of individual messages sent and received.

Although processing multiple messages in a single send or receive action is good for processing efficiency, if a link buffer that is sized for multiple messages is always full when the messages are sent, the link might be operating at or near its bandwidth capacity and message might be stacking up while waiting for the buffer to be cleared again.

If your logical links are operating at or near their bandwidth capacity, you might be able to ease the load on the link by distributing the message workload across additional logical links.

Determining optimum MSC link buffer sizes

The primary way to determine the optimum size of a link buffer is to use MSC link statistics to compare the number of times buffer content is sent with the number of actual messages sent.

If you are using bandwidth mode, the number of messages sent should be equal to or greater than the number of buffers sent. A buffer size that is optimized for bandwidth mode allows for multiple messages to be sent in a single send of the buffer contents, if there are multiple messages on the output queue.

If you are using non-bandwidth mode, the number of messages sent should be approximately equal to the number of buffers sent. A buffer size that is optimized for non-bandwidth mode allows for one complete message to be sent in each send of the buffer content.

To determine how well your buffer size is set, issue the command `QUERY MSLINK NAME(linkname) SHOW(STATISTICS)` and compare `Tot_Send_CT`, which is the total number of buffer sends, and `Tot_Msg_Send_CT`, which is the total number messages sent.

If you are using bandwidth mode, the `Tot_Msg_Send_CT` value should be equal to or greater than the `Tot_Send_CT` value. If the `Tot_Msg_Send_CT` is greater than the `Tot_Send_CT` value, multiple messages are being sent in each send of the buffer contents.

If you are using non-bandwidth mode, the `Tot_Msg_Send_CT` value should be approximately equal to the `Tot_Send_CT` value, which indicates that each send of the buffer contents contains a complete message. If the `Tot_Msg_Send_CT` value is less than the `Tot_Send_CT` value, the buffer size is not large enough to contain a complete message and multiple buffer sends are required to send a complete message.

Related reference:

 [QUERY MSLINK command \(Commands\)](#)

Determining MSC link buffer sizes for bandwidth mode

When MSC links run in bandwidth mode, you need to account for some additional bytes when determining the size of your link buffers.

In bandwidth mode each send of the buffer contents can include multiple responses and messages. However, if nothing else is on the output queue, a send of the buffer contents can include only one response or one message.

Buffer work fields are not included in the contents sent from the buffer.

Table 126. Byte format of MSC BUFMS buffers in bandwidth mode

Link type	Buffer work fields	Buffer header	Each response	Each message	VTAM work area
CTC	96	128	96	Size of message data, as shown in the X'01' or X'03' log record, plus 144-byte bandwidth prefix	N/A
MTM	96	128	96	Size of message data, as shown in the X'01' or X'03' log record, plus 144-byte bandwidth prefix	N/A
VTAM	96	128	96	Size of message data, as shown in the X'01' or X'03' log record, plus 144-byte bandwidth prefix	196

To calculate buffer sizes to use with MSC links in bandwidth mode:

1. Include 96-bytes for the buffer work field. The buffer work field is always the first 96 bytes of each buffer. The buffer work field is not sent.
2. Include 128 bytes for the buffer header.

- 3. For VTAM links only, include 196 bytes for the VTAM work area.
- 4. Include 96 bytes for responses. Each response is 96 bytes. The number of responses depends on the number of messages sent. Responses are inserted into the buffer first.
- 5. Determine the size of the messages being sent.

The size includes both the message prefix, data, and a 144-byte bandwidth prefix.

You can check the sizes of messages sent by issuing the QUERY MSLINK NAME(linkname) SHOW(STATISTICS) and look at the Hi_Msg_Send_SZ and Avg_Msg_Send_SZ values. The sizes returned by the QUERY MSLINK command include only the message prefix and data, but not the bandwidth prefix. The sizes returned are the same as the type X'01' and type X'03' queue manager message sizes.

The Hi_Msg_Send_SZ and Avg_Msg_Send_SZ values are derived from the actual sizes of the messages sent, as recorded in the X'01' and X'03' log records. The sizes in X'01' and X'03' log records include the prefix data and user data segments of each message. The X'01' and X'03' log records are mapped by the QLOGMSGP macro.

Related reference:

➡ QUERY MSLINK command (Commands)

➡ Log records (Diagnosis)

Determining MSC link buffer sizes for non-bandwidth mode

In non-bandwidth mode, you do not need to account for responses when determining the size of MSC link buffers.

You also do not need to account for the bandwidth header that is sent with each message when bandwidth mode is used.

In non-bandwidth mode, each send of the buffer contents includes only a single message. Responses are not sent as separate objects in the buffer contents.

Buffer work fields are not included in the contents sent from the buffer.

Table 127. Byte format of MSC BUFMS buffers in non-bandwidth mode

Link type	Buffer work fields	Buffer header	Response	Message	VTAM work area
CTC	96	6	Not applicable	Size of message data, as shown in the 01 or 03 log record, minus 124 bytes	Not applicable
MTM	96	6	Not applicable	Size of message data as shown in the 01 or 03 log record minus 124 bytes	Not applicable
VTAM	96	23	Not applicable	Size of message data as shown in the 01 or 03 log record minus 124 bytes	196

To calculate buffer sizes to use with MSC links in non-bandwidth mode:

- 1. Include 96-bytes for the buffer work field. The buffer work field is always the first 96 bytes of each buffer. The buffer work field is not sent.
- 2. Include the bytes for the buffer header as determined by your link type:

- For CTC links, include 6 bytes for the buffer header.
 - For MTM links, include 6 bytes for the buffer header.
 - For VTAM links, include 23 bytes for the buffer header.
3. For VTAM links only, include 196 bytes for the VTAM work area. The data in the VTAM work area is not sent with the messages.
 4. Calculate and include the size of the messages that are sent.
 In non-bandwidth mode, the size of the message that is sent is the message size shown in the X'01' and X'03' log records minus 124 bytes. In non-bandwidth mode, certain message prefixes are not sent with the message data and instead are rebuilt by the receiving IMS system.
 You can check the sizes of messages sent by issuing the QUERY MSLINK NAME(*linkname*) SHOW(STATISTICS) and look at the Hi_Msg_Send_SZ and Avg_Msg_Send_SZ values.
 The Hi_Msg_Send_SZ and Avg_Msg_Send_SZ values are derived from the actual sizes of the messages sent by subtracting 124 bytes from the recorded size in the X'01' and X'03' log records. The sizes in X'01' and X'03' log records include the prefix data and user data segments of each message. The X'01' and X'03' log records are mapped by the QLOGMSGP macro.

Related reference:

 QUERY MSLINK command (Commands)

 Log records (Diagnosis)

Use high-value link statistics to help diagnose MSC link problems

You can use high-value link statistics to help diagnose certain problems that lead to a backup of messages on MSC links. In the output of the type-2 command QUERY MSLINK command, the fields for the high-value statistics all begin with HI.

For example, if a message backup on a link occurs, you might check for an unusually high value recorded for the Hi_SendIO_Time. An unusually high I/O time to send a message might indicate that an outage has occurred in the network connecting the two IMS systems and you can then pursue the cause.

The high-value statistics that you can check include:

Hi_Proc_Time

The highest (longest) time the link was dispatched to process

Hi_MSG_Send_SZ

The largest message size sent (type X'01' or X'03' message record)

Hi_QGET_Time

The highest (longest) QMGR call (GU or DEQ) to process a send message.

Hi_SendIO_Time

The highest (longest) I/O time to send a message

Hi_MSG_Rec_SZ

The largest message size received (type X'01' or X'03' message record)

Hi_QPUT_Time

The highest (longest) QMGR call (ISRT or ENQ) to process a received message

Hi_RecIO_Time

The highest (longest) I/O time to receive a message

Monitoring and tuning multiple systems

Plan to obtain both statistical and performance data for IMS online systems that are part of a MSC network.

You can use the same monitoring tools that are used for generating performance data for single IMS systems:

- You can execute the IMS Monitor in several systems concurrently. You obtain IMS Monitor reports for each individual IMS system and coordinate your processing analysis.
- The Statistical Analysis utility produces summaries of transaction traffic for each individual system. Again, you combine the statistics for a composite picture.
- The IMS Transaction Analysis utility enables you to trace transactions across multiple systems and examine the traffic using the various active physical links.

Coordinating performance information

If possible, expand your MSC network by increasing the number of SYSIDs and the number of physical links and logical links.

You can specify up to 2036 SYSIDs. You can define up to 999 physical links and 999 logical links, unless the IMS system is defined as a large system by `SYSTEM=LGGEN` in the `IMSCCTL`. When `SYSTEM=LGGEN` is specified, you can define only 676 physical links and 676 logical links.

By expanding the MSC network, you can:

- Access an IMS subsystem from many other IMS subsystems
- Route transactions
- Distribute transaction processing
- Increase network throughput
- Grow beyond the capacity of one IMS system
- Respond to capacity constraints or response-time constraints

The IMS system log for each system participating in MSC contains only the record of events that take place in that system. However, logging is performed to record traffic on the links. Add the SYSIDs of all coupled systems to the system log documentation that records the checkpoint intervals. This helps to interpret reports, because you are aware of transactions that might be present in message queues but are not processed, and you can expect additional transaction loads from remote sources.

Your analysis procedures should include ways of isolating the processing that is triggered by transactions originating from another system. Considerations for tuning buffers for the asynchronous communication processing should include a criterion that no exceptional conditions resulting from intersystem traffic exist.

To satisfy the need for monitoring with typical activity that includes cross-system processing, coordinate your scheduling of the DC Monitor and other traces between master terminal operators. The span of the monitoring does not need to

be exactly the same, but if it is widely different, the averaging of report summaries might make it more difficult to interpret the effect of the processing that is triggered by cross-system messages.

Reports generated by the IMS Monitor for MSC

The IMS Monitor Report Print Program includes three reports that highlight message events caused by system coupling.

The reports are:

MSC Traffic report

Shows for a specified interval, the counts of messages using the various link paths. This report can be used to assess cross-system queueing

MSC Summaries report

Shows summaries of the traffic queues for each input transaction name, each destination name, each link number, and each destination system. The MSC Summaries report can be used to assess link loads.

MSC Queuing Summary report

Shows how long messages spend on queues and the numbers of messages on queues. The MSC Queuing Summary Report is generated when intersystem messages are queued on the local system before the messages are sent to the destination system. The local system must be an intermediate system.

All three reports can have entries in the Distribution Appendix, so that you can examine the frequency distributions of the traffic if you suspect unusual transmission patterns. The Distribution Appendix displays multiple occurrences of an event across a range of values relevant to the event, such as wait times. For example, in the case of wait times, the Distribution Appendix allows you to easily see how often a given event occurs with a specific wait time and what the distribution of the event is across a range of wait times.

Related concepts:

-  [IMS Monitor reports \(System Administration\)](#)
-  [MSC Traffic report \(System Administration\)](#)
-  [MSC Summaries report \(System Administration\)](#)
-  [MSC Queuing Summary report \(System Administration\)](#)
-  [Interpreting the Distribution Appendix \(System Administration\)](#)

Extracting multiple-system transaction statistics

You can use the Log Transaction Analysis utility to obtain counts of the message traffic both in local systems and between systems.

The transmissions over the different types of physical links can also be examined. (The activity is summarized for each step of the logical link paths.) You must provide IMS system log input that reflects all partner system activity, that is, sets of system logs for each MSC system. To coordinate the sets of individual system logs, use the Log Merge utility. As many as nine separate system logs can be merged, each log being the output of a uniquely identified IMS system with MSC installed.

Related reference:

➡ Log Transaction Analysis utility (DFSILTA0) (System Utilities)

Controlling the log merge

You need to perform a number of steps to control log output.

To control the log output, you need to:

- Choose the required systems that participate in the logical link paths you want to examine.
- Coordinate the series of input logs for each system so that they cover a similar time span.
- Specify a start and stop time for the Log Merge utility control statements if you want to sample the cross-system processing for a particular interval.

You can give both start date (Julian) and time of day, or just time of day. It is the first system log (specified by the LOG01 DD statement) to which these times apply. Other log activity is collected if it falls between the initial and final events that are present on the first log.

- Specify MSG to select log records that are suitable for the transaction analysis step. (ALL records is the default, but this means the DL/I activity for several systems is included in the utility input, and this can cause extended processing time.)

Related reference:

➡ Log Transaction Analysis utility (DFSILTA0) (System Utilities)

➡ Log Merge utility (DFSMTMG0) (System Utilities)

Interpreting the Transaction Analysis report

Using the Log Analysis report produced by the IMS Transaction Analysis utility, you can obtain statistics for individual transactions that are processed in any system.

The Log Analysis report includes the following statistics:

- The total response time
- The time on input and output queues
- The processing time

The Log Analysis report also includes definitions for the format of the detailed report records produced by the IMS Transaction Analysis utility and a list of processing type codes. The absence of times for a message GU call or MPP termination in the report lines indicates an input source or intermediate system report line.

The processing type field is an important one for the interpretation of the detailed report lines. The S code indicates that this line shows a send or receive event for the transaction. You can trace the progress of a cross-system conversation using the codes C, D, P, X, and Y.

The report headings include a column headed ID after the column for the GU to the message queue time. The number shown in a report line under the ID heading matches the sequence in which log input is fed to the Log Merge utility. The field corresponds to starting position 102, the 3-digit field named SYSTEM ID, in the detailed report records.

You can use the sort step to reorder the report records in any order you want, for example, by system ID within transaction code. The default order is the input sequence.

Related concepts:

➡ Statistical-analysis, log-transaction reports, and analyzing log records (System Administration)

Related reference:

➡ Log Transaction Analysis utility (DFSILTA0) (System Utilities)

MSC and IMSplexes with shared queues

The following topics discuss the coexistence of a Multiple Systems Coupling (MSC) network with an IMSplex with shared queues. MSC and IMSplex coexistence can be temporary, such as when migrating from an MSC network to an IMSplex configuration, or permanent, such as when an MSC link connects an IMSplex to an IMS system outside of the IMSplex.

A primary concern when an MSC network and an IMSplex coexist is the proper routing and processing of your transaction messages across both the MSC and IMSplex environments, because each environment uses a different routing method.

Generally, MSC networks route transactions to specific IMS systems using SYSIDs, while IMSplexes with shared queues route transactions by making them available on the shared queue to any IMS system that registers an interest in the transactions. When an MSC network and an IMSplex with shared queues coexist, both of these methods of routing can apply to transactions.

Message routing across MSC and IMSplex environments

MSC uses local and remote SYSIDs and destination names to route messages across links between local, intermediate, and remote IMS systems in an MSC network. IMSplexes with shared queues use destination name registration, IMSIDs, a shared queue, and the registered interest of IMS systems to route messages between front-end and back-end IMS systems in an IMSplex. When an IMSplex and an MSC network coexist, both methods of routing can be used.

In MSC networks and IMSplexes with shared queues, IMS stores the SYSID, destination name, and IMSID value in the message prefix when IMS builds the message and places it on a local or shared queue. The SYSID, destination name, or IMSID stored in the message prefix must match the SYSID, destination name, or IMSID of an IMS system in either the IMSplex or the MSC network. If they do not match, any messages on the shared queue or in flight in the MSC network will trigger a routing error, such as a pseudoabend U0830, when an IMS system attempts to process the message. This condition can occur, for example, if the IMS system that owns the SYSID, destination name, or IMSID is down or has changed its value.

How messages are routed in an MSC network

In an MSC network, IMS uses remote transactions, remote logical terminal (LTERM) names, and SYSIDs to route the messages through the MSC network.

Remote transactions are defined by specifying remote and local SYSIDs in the TRANSACT macro. Remote LTERMs are defined in NAME macros immediately following an MSNAME macro. The remote and local SYSIDs in the MSNAME macro are then associated with the remote LTERM. IMS stores the remote

transaction codes, LTERM names, and SYSIDs in the message prefix of each message routed through the MSC network. In this context, the transaction codes and LTERM names are referred to as destination names.

MSNAME labels can also be used as intermediate destination names, when messages must pass through an intermediate IMS system prior to reaching the IMS system that contains the true destination name. In this case, the MSNAME label is also stored in the message prefix.

Related concepts:

“Routing messages with the destination name and SYSIDs” on page 690

Related tasks:

“Defining Multiple Systems Coupling resources” on page 700

Related reference:

➞ MSNAME macro (System Definition)

➞ NAME macro (System Definition)

➞ TRANSACT macro (System Definition)

How messages are routed in an IMSplex with shared queues

In an IMSplex with shared queues, IMS uses origin names and destination names registered with a coupling facility to route messages between front end and back end IMS systems.

Origin and destination names can be logical terminal names, transaction codes, or APPC or OTMA client names. You define static origin and destination names during system definition. You define in descriptor libraries or user exits the origin and destination names that are dynamic LTERMs, ETO terminals, or dynamic transactions. Origin and destination names for APPC or OTMA clients are sent, in the form of an 8-byte token, to IMS when an APPC or OTMA client allocates a conversation.

IMS registers transactions when they are started and a dependent region or an MSC link is started and ready to process messages. IMS registers LTERM names with the coupling facility when they are started or when they sign on and are ready to process messages. IMS registers APPC and OTMA tokens when an APPC or OTMA conversation is allocated. IMS deregisters APPC and OTMA tokens when the APPC or OTMA conversation is deallocated.

Messages carry these origin names and destination names in the message prefix. When an empty shared queue on a coupling facility receives a message from a front-end IMS system for a registered destination name, the coupling facility notifies all interested IMS systems that there are messages to process.

Message routing when an IMSplex and MSC network coexist

When an IMSplex with shared queues and an MSC network coexist, IMS routes messages between front-end, back-end, and remote IMS systems using origin and destination names, MSC SYSIDs, and IMSIDs.

IMS uses the coupling facility and the shared queues to route messages within the IMSplex. IMS uses MSC links only to route messages to IMS systems outside of the IMSplex. Any MSC link definitions that exist between IMS systems in the IMSplex are not used.

Transactions defined for routing in an MSC network can also be routed in an IMSplex with shared queues but the method used for routing is that of the IMSplex and not MSC; however, the IMSplex does recognize the MSC SYSID attributes of the transaction. The recognition of the SYSIDs by the IMSplex allows the IMSplex to handle remote transaction in a manner similar to the MSC network. This also has implications for local processing affinity.

Transactions not specifically defined for routing in an MSC network cannot be routed outside of the IMSplex using an MSC link.

Related concepts:

“Processing affinities in an IMSplex”

Processing affinities in an IMSplex

If a message must process on a specific IMS system in the IMSplex, IMS assigns the message affinity to that IMS system by appending the IMSID of the IMS system to the destination name. Then, only that IMS can process it.

You can also use the TM and MSC Message Routing and Control user exit routine (DFSMSCE0) to establish affinities between transaction instances and IMS systems in an IMSplex.

Processing affinities in an IMSplex-MSC coexistent configuration

In an MSC network, remote transactions bypass any local affinity the transactions might have to the IMS system that receives them as input. In an IMSplex, remote transactions also cause local affinity to be bypassed. Otherwise, the rules that govern processing affinities in an IMSplex remain unchanged.

Related concepts:

“TM and MSC Message Routing and Control user exit routine overview” on page 735

Related tasks:

“Managing remote transactions for APPC and OTMA when MSC and IMSplexes coexist” on page 728

Migrating from an MSC network to an IMSplex network

When migrating or converting an existing MSC network to an IMSplex with shared queues, you need to retain the MSNAME definitions in your IMS systems for as long as you have transaction messages that use those definitions on the shared queues.

You also need to retain any MSC link definitions to IMS systems outside the IMSplex. Once you have finished testing the IMSplex and are sure you will not need to return any IMS systems in the IMSplex to an MSC-only network, you can remove any MSC link definitions between IMSplex member systems.

In an IMSplex, if one IMS system is MSC capable, then all the IMS systems in the IMSplex must be MSC capable. To enable an IMS system for MSC, define an MSC link for the IMS system during system definition. This MSC link definition does not need to be functional.

MSC link definitions in an IMSplex

During the migration process from an MSC network to an IMSplex, you can retain your MSC link definitions for as long as you have a need for them. This allows you to test the IMSplex with shared queues environment and then revert to the MSC network to make adjustments.

After you have established the IMSplex with shared queues, you should not add new MSC links between the IMS systems in the IMSplex. If you attempt to start an MSC link between two IMS systems within the same IMSplex with shared queues, IMS issues message DFS2149 and aborts the link startup.

Sharing MSNAME definitions and SYSIDs in an IMSplex

In an IMSplex with shared queues, when you start an IMS system that includes MSNAME statements, the IMSplex shares the MSNAME statements and the local SYSIDs of the joining IMS system with the other IMS systems already in the IMSplex.

For each MSNAME statement in the joining IMS system, the IMSplex creates a duplicate dynamic MSNAME statement in each of the other IMS systems in the IMSplex, unless one already exists.

For each local SYSID owned by the joining IMS system, the IMSplex adds a duplicate local SYSID to the SYSID table in each of the other IMS systems in the IMSplex. This effectively makes the SYSID of the joining IMS system local to the IMSplex as a whole. The IMSplex updates the SYSID table of each IMS system in the IMSplex whenever an IMS system that has an MSNAME statement joins or rejoins the IMSplex.

The generation of dynamic MSNAME statements and the sharing of local SYSIDs throughout the IMSplex allow the IMSplex to function as a single IMS system on the MSC network. It also allows transactions defined to run in an MSC network to take advantage of the distributed processing of the IMSplex with shared queues.

To illustrate dynamic MSNAMEs and shared SYSIDs, The following figure shows a simple MSC network prior to introducing an IMSplex with share queues. Figure 120 on page 725 then shows the same MSC network after an IMSplex with shared queues has been created between two of the IMS systems.




Figure 119. MSC network without an IMSplex

In the previous figure, IMS1, IMS2, and IMS3 are each members of the MSC network. Each of their local SYSIDs are unique and their MSNAME statements only define the links that are defined locally in each IMS system.




Figure 120. MSC network coexisting with an IMSplex with shared queues

In the preceding figure, IMS1 and IMS2 are now part of an IMSplex with shared queues. Because of this, SYSID 1 becomes local to IMS2 and SYSID 2 becomes local to IMS1. Similarly, dynamic MSNAME MSN31 appears in IMS2, and dynamic MSNAME MSN32 appears in IMS1. IMS1 and IMS2 also remain connected to the MSC network and IMS3, although IMS1 and IMS2 now function as a single IMS system, or node, within the MSC network.

The following table represents the SYSID tables of the IMS systems shown in Figure 119 on page 725 prior to introducing the IMSplex:

Table 128. SYSID ownership in an MSC network without an IMSplex

SYSID	IMS1	IMS2	IMS3
1	Local	RMT/MSN12	RMT/MSN13
2	RMT/MSN21	Local	RMT/MSN23
3	RMT/MSN31	RMT/MSN32	Local

The following table represents the resulting SYSID tables in each IMS after the IMSplex is introduced into the MSC network, as shown in Figure 120 on page 725:

Table 129. SYSID ownership in an MSC network that coexists with an IMSplex

SYSID	IMS1	IMS2	IMS3
1	Local	Local	RMT/MSN13
2	Local	Local	RMT/MSN23
3	RMT/MSN31 Dynamic MSN32	RMT/MSN32 Dynamic MSN31	Local

When an IMS system leaves the IMSplex, the IMSplex does not change the SYSID tables or delete any dynamic MSNAME statements in the remaining IMS systems. This ensures that if there are any messages on the shared queues that use the MSNAME definitions of the departed IMS, the IMSplex can still process them. When an IMS system rejoins the IMSplex, the IMSplex always exchanges MSNAME statements again and revalidates the rejoining SYSID on all SYSID tables in the IMSplex.

Whenever an IMS system successfully joins or leaves an IMSplex with shared queues, the IMSplex issues message DFS0778I to the master terminal of each IMS system in the IMSplex.

MSNAME duplication in an IMSplex with shared queues

For each IMS system that participates in both an MSC network and a shared queues group, MSNAMEs and their associated remote and local SYSIDs are exchanged with the other IMS systems in the shared queues group when the IMS is initialized.

Dynamic MSNAMEs are created in each IMS system for the MSNAMEs that were defined on the other IMS systems in the shared queues group.

A common MSC SYSID routing table is built in each IMS system from the remote and local SYSID values that are defined in each MSNAME. SYSIDs that are local in each IMS system are also made local to the other IMS systems in the shared queues group. The same local SYSIDs on all IMS systems in the shared queues

group allow any message received from an IMS system outside the shared queues group through an MSC link to be routed and processed by any IMS system in the shared queues group.

Duplicate MSNAMEs in IMS systems in a shared queues group must have the same remote SYSID because MSNAMEs and remote SYSIDs are synonymous. The same MSNAMEs must relate to the same remote SYSIDs. Different local SYSIDs for the same MSNAME are allowed.

If the same MSNAME is defined with different remote SYSIDs in two IMS systems, the MSNAME is ignored when MSNAMEs are exchanged within the shared queues group. Dynamic MSNAMEs are not created in the IMS systems with the same MSNAME and each IMS retains its own version of the MSNAME.

Each of these two examples show valid combinations of MSNAMEs that have the same name in a shared queues group:

- A valid MSNAME pair:
 - (On IMSA) MSN12345 MSNAME SYSID=(1,2)
 - (On IMSB) MSN12345 MSNAME SYSID=(1,2)
- Also a valid MSNAME pair:
 - (On IMSA) MSN12345 MSNAME SYSID=(1,2)
 - (On IMSB) MSN12345 MSNAME SYSID=(1,3)

This is an example of an invalid combination of MSNAMEs that have the same name in a shared queues group:

- (On IMSA) MSN12345 MSNAME SYSID=(1,2)
- (On IMSB) MSN12345 MSNAME SYSID=(2,1)

Related concepts:

 Resource type consistency (System Administration)

Deleting MSNAME definitions from the SYSID tables in an IMSplex

To remove a SYSID or an MSNAME statement from the IMSplex, you must first remove it from the IMS system that owns it and then cold start the IMS system back into the IMSplex.

After you have removed the SYSID or MSNAME statement from its original owning IMS system, warm or cold start each of the other IMS systems in the IMSplex to rebuild their SYSID tables.

Removing MSPLINK and MSLINK definitions when an MSC to IMSplex migration is complete

After the migration of an MSC network to an IMSplex with shared queues is complete and you are sure you will not return any IMS systems to their MSC-only state, you should remove the MSPLINK and MSLINK definitions for the MSC links that you will no longer use.

Note: You must retain at least one MSC link definition in each IMS system in the IMSplex to maintain MSC enablement throughout the IMSplex. This link can be a functioning MSC link to an IMS system outside of the IMSplex or a non-functioning link defined only to attach MSNAME statements to and to ensure MSC enablement.

Managing SYSIDs when MSC and IMSplexes coexist

Local SYSIDs must be unique in an MSC network. Local SYSIDs are not required to be unique within an IMSplex; however, local SYSIDs owned by IMSplex members linked to an MSC network outside of the IMSplex must be unique in that MSC network.

Also, if multiple members of an IMSplex link to the same IMS system outside of the IMSplex, each IMSplex member must use a local SYSID that is unique to that outside IMS system.

You cannot use the same remote SYSID to point to different IMS systems in an MSC network.

Cloning MSC SYSIDs in an IMSplex

Even though the IMSplex adds every local SYSID in the IMSplex to each SYSID table in the IMSplex, you should clone all local SYSIDs within each IMS system in the IMSplex when removing the intra-IMSplex MSC links. To clone a local SYSID, clone an MSNAME statement that contains it.

Cloning local SYSIDs ensures that the IMSplex can route and process messages on a shared queue even if the IMS system that owns the SYSID is not available and the SYSID is not included in the SYSID tables in the IMSplex. MSC SYSIDs are not stored in the coupling facility within the IMSplex. If an IMS system cannot recognize the SYSID of a message on the shared queue as local when processing application program messages, the IMSplex issues pseudoabend U0830.

IMSIDs when IMSplexes and MSC coexist

The IMSplex uses IMSIDs to identify IMS systems and to route any messages that have an affinity to a specific IMS system. IMSIDs must be unique within each IMS in an IMSplex.

In an IMSplex-MSC network, IMSIDs must be unique throughout the entire MSC network, even the remote MSC IMS systems outside of the IMSplex.

IMSIDs are not required to be unique in MSC networks that do not include an IMSplex. When migrating an MSC network to an IMSplex with shared queues, make sure the IMSIDs of each IMS system are unique.

Managing remote transactions for APPC and OTMA when MSC and IMSplexes coexist

The following topics provide information about migrating from an MSC network to an IMSplex when synchronous APPC and OTMA messages were processed remotely in the MSC network, and processing APPC and OTMA messages on a remote IMS system outside of the IMSplex when an IMSplex and an MSC network coexist.

Remote processing of APPC and OTMA messages in an MSC network

In an MSC network, you can bypass local affinity for synchronous APPC and OTMA messages by queuing them to remote MSC transactions.

Remote IMS systems process APPC and OTMA messages asynchronously and disassociated with the APPC or OTMA conversation. IMS does not propagate or cascade APPC or OTMA conversations to the remote IMS system. IMS saves an APPC or OTMA token in the message prefix. When a response or

program-to-program switch returns to the originating IMS system, IMS attempts to restore the synchronous APPC or OTMA conversational mode to the message.

To enable the remote processing of an APPC or OTMA message in an MSC network, queue the message to a remote MSC transaction.

Back-end processing of APPC or OTMA transaction messages in an IMSplex with shared queues

In an IMSplex with shared queues, you can use remote MSC transactions to route APPC or OTMA messages to back-end IMS systems for processing.

Using remote MSC transactions avoids the local affinity restrictions that are otherwise imposed on transaction messages received from APPC or OTMA clients. Any IMS system in the IMSplex that defines the transaction as a local MSC transaction that is assigned to a region can then process the transaction.

If you are migrating from an MSC network to an IMSplex, you can use existing remote MSC transactions to process APPC or OTMA messages on back-end IMS systems in the IMSplex and bypass the APPC and OTMA affinity restrictions. When the transactions are converted to local IMSplex transactions, the APPC and OTMA affinity restrictions still apply unless you specify RRS=Y and AOS=Y.

When an APPC or OTMA message is queued in a front-end IMS system to a transaction that is defined as a remote MSC transaction, IMS inserts the transaction message to the shared queue without affinity to any IMS system.

When a back-end IMS system retrieves APPC and OTMA transaction messages from the shared queue, it saves an APPC or OTMA conversation token in the message prefix and processes the transaction message independently from, and asynchronous to, the APPC or OTMA conversation. IMS does not propagate or cascade the APPC or OTMA conversation to the back-end IMS system or to any other IMS systems, including the front-end IMS system, that might process APPC or OTMA messages after a program-to-program switch.

Note: If an OTMA transaction that is processing independently from the OTMA conversation issues a CHNG call to a modifiable PCB and IMS calls the OTMA Destination Resolution exit routine (DFSYPX0), the DFSYPX0 exit routine recognizes the transaction message as an OTMA message only if the transaction is processing within the IMSplex where the OTMA client originally submitted the transaction.

When the originating front-end IMS system receives the first or only response, the response is returned to the client in APPC or OTMA synchronous mode, assuming that the client is still connected in synchronous mode. IMS returns any subsequent responses to the client for the same interaction asynchronously. If the client has terminated before the first or only response is returned, the response is not discarded, but instead queued to the client asynchronously.

For example, if you define an MSC transaction as remote in a front-end IMS system without assigning the transaction to a started MSC link and then you define the same MSC transaction as local in the back-end IMS and assign the transaction to a started region, when an APPC or OTMA client initiates the transaction on the front-end IMS system in synchronous mode, the front-end IMS system saves a token in the message prefix to identify the client and places the transaction on the shared queue without any affinity to the front-end IMS system.

When the back-end IMS system retrieves the transaction from the shared queue, the transaction is processed asynchronously and disassociated from the APPC/OTMA client.

The following series of figures show various possible scenarios in which remote MSC transactions can be used to process APPC or OTMA transactions on the back end in an IMSplex, as well as outside of an IMSplex by using an MSC link.

The following figure shows a simple configuration in which a remote MSC transaction is used to process an APPC or OTMA transaction on the back end in an IMSplex. The transaction TRAN1 is defined to IMS 1 as a remote MSC transaction, but because TRAN1 is not assigned to an MSC link in IMS 1, IMS 1 queues it to the shared queue.

After submitting the transaction to IMS 1, the APPC or OTMA client waits in a receive state for the synchronous response; however, the queuing and processing of TRAN1 by IMS 1 and IMS 2 is handled by IMS independently from the synchronous communications mode that is maintained with the client.




Figure 121. Using a remote MSC transaction to route an APPC or OTMA transactions to a back-end IMS system

The following figure shows another scenario in which a remote MSC transaction is used to route APPC or OTMA transaction. This scenario involves three IMS systems and TRAN1 is routed to a remote MSC IMS system by way of both a shared queue and an MSC link.

When IMS 1 receives the transaction from the client, IMS 1 queues it to the shared queue. In IMS 2 the transaction TRAN 1 is defined as a remote MSC transaction and assigned to the MSC link to IMS 3. IMS 3 generates the response, which is then returned across the MSC link to IMS 2, and then returned to IMS 1 by way of the shared queue. IMS 1 returns the response to the APPC or OTMA client in the synchronous mode expected by the client.




Figure 122. Routing an APPC or OTMA transaction across a shared queue and an MSC link

In the scenario shown in the following figure, the APPC or OTMA transaction TRAN1 is routed to IMS 2, a back-end IMS system in a shared queues group, where a program-to-program switch to TRAN2 sends processing back to IMS 1 by way of the shared queue. TRAN2 is not an MSC transaction. TRAN2 then sends processing back to IMS 2 by issuing another program-to-program switch to TRAN3, also a non-MSC transaction. TRAN3 is queued to the shared queue and retrieved by IMS 2, where TRAN3 processes and generates a response for the client. The response is returned to IMS 1 by way of the shared queue. IMS 1 returns it to the client in the synchronous mode expected by the client.




Figure 123. Routing an APPC or OTMA transaction by using a remote MSC transaction and shared queues with multiple program-to-program switches

The scenario shown in following figure involves three IMS systems. IMS 1 routes TRAN1 to IMS 2 as a remote MSC transaction by way of the shared queue. TRAN1 processes locally on IMS 2 and issues a program-to-program switch to TRAN2, which is defined as a remote MSC transaction on IMS 2. IMS 2 sends TRAN2 to IMS 3 across an MSC link. TRAN2 is processed by IMS 3 and generates the response for the client. IMS 3 returns the response to IMS 2 across the MSC link. IMS 2 queues the response to the shared queue. IMS 1 retrieves the response and returns it to the APPC or OTMA client in the synchronous mode the client expects.




Figure 124. APPC or OTMA transaction routed using MSC and shared queues with a program-to-program switch

Related concepts:

Managing APPC and OTMA messages in a sysplex environment (System Administration)

Enabling back-end processing of APPC and OTMA messages using a remote transaction

To use a remote transaction to process APPC and OTMA messages on back-end IMS systems within an IMSplex, perform the following steps.

1. Enable MSC in all IMS systems in the IMSplex by defining at least one MSC link in each IMS system in the IMSplex. This MSC link can be either a functional link to a remote IMS system outside of the IMSplex or, if none exists, it can be a non-functional MSC link.
2. In the IMS system connected to the APPC or OTMA client, define a remote transaction for the APPC or OTMA messages by specifying the SYSID keyword on the TRANSACT macro. In the SYSID keyword, the remote SYSID can be any remote SYSID. The local SYSID should match the SYSID of the IMS system connected to the APPC or OTMA client.
3. In all back-end IMS systems, define a local transaction to process the APPC or OTMA messages.
4. Register the remote transaction with the IMSplex by starting the transaction and the dependent regions that process it.

Sending a transaction to an MSC system outside of an IMSplex

To send an APPC or OTMA transaction from an IMSplex to a remote IMS system outside of the IMSplex by using an MSC link:

1. Enable MSC in all IMS systems in the IMSplex by defining at least one MSC link in each IMS system in the IMSplex. This MSC link can be either a functional link to a remote IMS system outside of the IMSplex or, if none exists, it can be a non-functional MSC link.
2. Define an MSC link between the IMSplex and the remote IMS system.
3. In any IMS system in the IMSplex, define a remote MSC transaction and assign it to an MSC link by specifying remote and local SYSIDs in the SYSID keyword of the TRANSACT macro. The remote SYSID must match the SYSID of the remote MSC IMS system. The local SYSID must match the SYSID of the IMS system that is connected to the APPC or OTMA client.
4. In the remote MSC IMS system, define a local transaction to process the APPC or OTMA transaction message.

Avoiding pseudoabend U0830

When an MSC network and an IMSplex with shared queues coexist, errors can occur if you remove IMS systems from, or add them to, the IMSplex or MSC network while messages remain on the shared queues or in flight within the MSC network.

Errors can also occur if the definitions used to route the messages, such as SYSIDs, IMSIDs, or destination names, are changed.

In an IMSplex with MSC, IMS validates SYSIDs, IMSIDs, and destination names when IMS systems retrieve transaction messages from the shared queue for processing. If validation fails, IMS leaves the transaction message on the shared queue and issues a pseudoabend U0830. By taking these precautions you can avoid most U0830 pseudoabends:

- Prior to migrating or moving IMS systems, dequeue all transaction messages destined for the IMS systems being migrated or moved from the shared queues and the MSC network.
- Prior to changing any SYSIDs, IMSIDs, or MSNAMEs, dequeue all transaction messages that use these identifiers from the shared queues and the MSC network.
- After migrating new IMS systems into an IMSplex, wait for all the IMS systems to join the IMSplex and all SYSID tables to be updated before you start dependent regions to process messages
- When restarting multiple IMS systems with MSC definitions in an IMSplex, if all local SYSIDs in the IMSplex have not been cloned in each IMS system, wait for all of the IMS systems to finish restarting before starting any dependent regions. Until restart is complete, some SYSID tables might not include all of SYSIDs that are not cloned.
- Verify that MSC is defined to each IMS system in the IMSplex. Messages queued by an IMS that does not have MSC defined will not have the MSC extension prefix and will have zero SYSID values. If this message is processed by another IMS system in the IMSplex that has the MSC feature, a pseudoabend U0830 results with return code 08.

In an IMSplex, if one IMS system is MSC capable, then all IMS systems in the IMSplex must be MSC capable. To enable an IMS system for MSC, define an MSC link for the IMS system during system definition. This MSC link definition does not need to be functional.

For detailed information about the return codes associated with pseudoabend U0830 and other diagnosis information, see *IMS Version 11 Messages and Codes, Volume 3: IMS Abend Codes*.

VTAM Generic Resources (VGR) and MSC

MSC IMS systems outside of an IMSplex can use a VTAM generic resource group name (GRSNAME) to establish MSC link sessions with the IMSplex instead of using the APPLID of a specific IMS system in the IMSplex.

When you use a VTAM GRSNAME, you can move MSC link sessions between IMS systems within the IMSplex without having to change the VTAM node name on the NAME= parameter in the MSPLINK stage 1 system definition macro for each IMS system outside the IMSplex.

Using a VTAM GRSNAME can also make cloning IMS systems within an IMSplex easier because the same MSC link definition can be used by each IMS system within the IMSplex to connect to a remote IMS system. Each parallel link in the IMSplex can use a common GRSNAME to establish a session with an IMS system outside of the IMSplex. Other link attributes, such as partner ID, SYSIDs, and so forth, can also be cloned across the IMS systems within an IMSplex.

The use of VTAM generic resources in an IMSplex is transparent to the external MSC system. The external MSC system need only specify the GRSNAME that is used by the IMSplex, instead of a specific IMS system APPLID name, on the NAME parameter of the MSPLINK macro. In the IMSplex, the MSC VGR support is enabled by specifying MSCVGR=Y (MSCVGR=N is the default) in the DFSDCxxx PROCLIB member of each IMS system in the IMSplex. The IMSplex must also have VTAM generic resources enabled.

When MSC VGR support is enabled, also specify GRMESTAE=Y (the default) in the DFSDCxxx PROCLIB member to ensure that IMS deletes VTAM affinity if an IMS abend should occur. The type-1 command /DISPLAY AFFIN LINK and the type-2 command QUERY MSLINK SHOW(AFFIN) display current affinities between MSC links and the IMS systems in an IMSplex.

Related concepts:

 Planning for VTAM generic resource groups (System Administration)

Related reference:

 DFSDCxxx member of the IMS PROCLIB data set (System Definition)

 QUERY MSLINK command (Commands)

 /DISPLAY AFFIN command (Commands)

TM and MSC Message Routing and Control user exit routine overview

Message routing is automatic, according to the defined scheme, unless you use the TM and MSC Message Routing and Control user exit routine (DFSMSCE0), which provides routing options and control of messages.

The DFSMSCE0 exit routine provides a single parameter list and the ability to append optional user prefixes to messages for customizing message routing and security.

The DFSMSCE0 exit routine also provides an IMSplex affinity routing option for IMSplexes with shared queues. Affinity routing establishes an affinity between a transaction instance and a destination IMS system specified by the exit routine. When the DFSMSCE0 exit routine assigns an affinity to a transaction and the transaction is inserted to the shared queue, only the destination IMS system is notified.

Note: Using this exit routine does not require Multiple Systems Coupling (MSC), though most of its options work only when MSC is enabled.

This routing exit routine is called before the message destination is final. The entry points of the exit routine are:

- **Terminal routing (TR):** Receives control when a message is received from a terminal.
 - VTAM messages (TRVTAM)
 - APPC messages (TRAPPC)
 - OTMA messages (TROTMA)
- **Link receive (LR):** Receives control when a message is received on a MSC link.
 - Local transaction messages (LRTRAN)
 - Local LTERM messages (LRLTERM)
 - Local direct routing messages (LRDIR)
 - Intermediate messages (LRINT)
- **Program routing (PR):** Receives control when the application program issues a CHNG or ISRT call to insert a message.
 - Application program CHNG call (PRCHNG)
 - Application program INST call (PRINST)

The TM and MSC Message Routing and Control User Exit routine is loaded during IMS initialization, provided it exists in the JOBLIB, STEPLIB, or LINKLIST library that is concatenated in front of IMS.SDFSRESL. No system definition or startup parameter modules are needed to invoke this exit routine.

You can append optional user-defined prefixes to messages. Message prefixes can be used, for example, to customize message security, user accounting, and statistics requirements, and to increase routing control by allowing communications among exit routines. This user message prefix segment can be added to the message and updated as the message is routed through the MSC/TM network as each exit routine entry point is called. These other exit routines can read or update the prefix segment. The user prefix segment can be used offline. The user prefix size is limited to 512 bytes, and the total message prefix size is limited to the large message queue LRECL.

The TM and MSC Message Routing and Control User Exit routine uses a common parameter list interface (DFSMSCEP) for all of the entry points. DSECTS are provided to reference all parameter fields. You can select the entry points at which the exit routine should take control by changing and reassembling pointers in the DFSMCCSV macro in the front of the user exit module. IMS calls the exit routine if the entry point exists.

Recommendation: Provide a DSECT DFSMSCEP macro for common reference to the parameter fields to make it easier in the future to pass to the exit routines additional data, such as performance or path availability data for use in making routing decisions. Data might include SIO rates, response times, or queue counts.

Terminal routing

The terminal routing entry points of the TM and MSC Message Routing and Control User Exit routine exist in the input system. The exit routine is called when a message is received from a terminal. The exit routine can inspect the specified destination (LTERM or transaction code) and, if specified, reject it or change it to any local or remote destination. If the exit routine does not change the destination, the originally specified destination is used for routing. The exit routine can also override a local LTERM or transaction to route a message to a remote IMS instead.

IMS does not call the exit routine for commands or from a terminal that is continuing a conversation.

In a configuration using horizontal partitioning, the exit routine can be used to evaluate all input messages and route them to the appropriate processing system based on information in the first segment of the input message. If transactions and links are appropriately defined, the exit routine can also be used to set a common destination for a set of input messages. On arriving at the destination system, the messages are processed according to their individual transaction codes.

Link receive

The Link Receive entry points of the TM and MSC Message Routing and Control User exit routine (DFSMSCE0) can change the transaction code or LTERM name of a message when IMS receives the message from an MSC link. DFSMSCE0 can request that the message be processed locally in the current IMS system or it can reroute the message to a different remote IMS system. DFSMSCE0 can inspect the transaction code or the LTERM name that is used as the destination and, if specified, reject it or change it to another destination with the same attributes. The exit routine can also examine the first segment of the message to determine what the new transaction code should be. If the exit routine does not change the transaction code, the destination remains the location that is specified by the original transaction code. DFSMSCE0 can also determine what security to use, if any, on a message by message basis.

Program routing

By using the program routing points of the TM and MSC Message Routing and Control User exit routine (DFSMSCE0), you can control the routing of a message when an application program issues a CHNG or ISRT call. With these entry points, you can change the destination of the message or request to reject a message. The DFSMSCE0 exit routine can also request that the message be processed locally in the current IMS system or it can reroute the message to a different remote IMS system.

You can use the exit routine to avoid defining unique names for remote LTERMs and transactions. With the exit routine, you can have the same name for terminals throughout your MSC network. You can use this exit routine to execute the DL/I CHNG or ISRT call for the application program, and to change the destination from local to remote.

Related reference:

➡ TM and MSC Message Routing and Control User exit routine (DFSMSCE0) (Exit Routines)

➡ MSGQUEUE macro (System Definition)

The IMSplex affinity routing option of the DFSMSCE0 exit routine

In IMSplexes with shared queues, the IMSplex affinity routing option of the TM and MSC Message Routing and Control user exit routine (DFSMSCE0) can establish an affinity between a transaction instance and an IMS system in which the transaction is started with the affinity option.

MSC must be enabled in the IMS systems that use the IMSplex affinity routing option of the DFSMSCE0 exit routine.

When the DFSMSCE0 exit routine assigns an affinity to a transaction and the transaction is inserted to the shared queues, only the IMS system specified by the DFSMSCE0 exit routine is notified.

The IMSplex affinity routing option of the DFSMSCE0 exit routine enables back-end IMS systems in an IMSplex to process both APPC synchronous transactions and OTMA send then commit (CM1) transactions. The IMSplex affinity routing option overrides restrictions that would otherwise prevent the APPC and OTMA CM1 transactions from processing on back end IMS systems in the IMSplex. When the affinity option is invoked, APPC and OTMA transactions run in a mode that disconnects the transaction message from the client, similar to the mode in which MSC routes messages to remote IMS systems in an MSC network.

The DFSMSCE0 exit routine can be called to establish affinity at the following exit points:

Terminal routing (TR)

The DFSMSCE0 exit routine receives control and can set affinity for a message that is received from a terminal by an IMS system in an IMSplex.

Link receive (LR)

The DFSMSCE0 exit routine receives control and can set affinity for a message that is received from an MSC link by an IMS system in an IMSplex.

Program routing (PR)

The DFSMSCE0 exit routine receives control and can set affinity for a message when an application program running in an IMS system in an IMSplex issues a CHNG or ISRT call to insert the message.

The DFSMSCE0 exit routine establishes affinity for a transaction instance by appending the IMS ID or, in an XRF complex, the RSENAME of the destination IMS system to the shared queue name (SQNAME) in the message prefix.

Before a transaction instance can be processed by the IMS system with which it has affinity, the transaction type must be registered as having affinity on the IMS system.

To start a transaction with the affinity option in an IMS system, specify either of the following commands on the IMS system that processes the transaction:

- The type-2 command `UPDATE TRAN NAME(trannname) START(SCHD) OPTION(AFFIN)`
- The type-1 command `/START TRAN trancode AFFINITY`

IMS maintains the affinity status of a transaction type in an IMS system across warm starts and emergency restarts, but not across either a cold start or an export and import of a transaction definition to a resource definition data set.

To display transactions on the shared queues that have affinity to an IMS system, you can issue either of the following commands with the appropriate keywords:

- The type-2 command `QUERY TRAN QCNT(GT,0) SHOW(AFFIN)`
- The type-1 command `/DIS TRAN ALL QCNT`

After an IMS system retrieves a transaction with affinity from the shared queue, the affinity status of the transaction is displayed under `LclStat` as `AFFIN`. To see which transactions have local affinity status on an IMS system, you can issue the following type-2 command:

- `QUERY TRAN NAME(tranname | * ) SHOW(STATUS)`

For example, if you issue `QUERY TRAN NAME(APOL11 APOL12) SHOW(STATUS)`, the command output might look like the following:

Trancode	MbrName	CC	LclStat
APOL11	IMS1	0	
APOL12	IMS1	0	AFFIN

Related reference:

 [TM and MSC Message Routing and Control User exit routine \(DFSMSCE0\)](#)
(Exit Routines)

Part 9. ODBA and DRA connections

The Open Database Access (ODBA) and database resource adapter (DRA) interfaces provide direct access to IMS databases from applications running on the same z/OS image as IMS.

Application programs that use DRA or ODBA and DRA include:

- CICS Transaction Server for z/OS
- ODBA z/OS applications, such as DB2 for z/OS stored procedures

Chapter 41. Accessing IMS databases with CICS

CICS application programs can access DL/I databases through DBCTL. This means that DBCTL satisfies CICS DL/I requests by means of the CICS-DBCTL interface.

CICS/ESA and CICS Transaction Server for z/OS provide support for IMS DBCTL.

Installing data sharing with CICS

CICS users can optionally use the IMS data sharing facility. The Internal Resource Lock Manager (IRLM) is mandatory for block-level data sharing, but not for database-level sharing.

Installing CICS for use with IMS intersystem communication

Information about defining CICS as an Intersystem Communication node and defining CICS tables to be compatible with IMS in an Intersystem Communication network is provided in Part 7, “Intersystem Communication (ISC),” on page 457.

Related concepts:

 Data sharing in IMS environments (System Administration)

Coding considerations for PSBs

This topic provides some guidelines for coding PSBs for the DRA when using CICS.

PSBs for Online Transactions: A PSB is needed for each online program that accesses DL/I databases.

- The name of the PSB specified in the PSBNAME= keyword of the PSBGEN macro must be exactly the same as one of the entries in the PDIR.
- The name of the PSB must also be the same as the name specified in the scheduling call issued by an online transaction. If the online transaction does not specify any name in the scheduling call, the name of the PSB must be exactly the same as the name of the program associated with the transaction in the CICS program control table (PCT).

Example: Assume transaction X is associated with program Y. If program Y links or transfers control to program Z, which issues the scheduling call without specifying any PSB name, the default name for the PSB is Y, because it is program Y that is associated with transaction X in the PCT.

- The CMPAT=YES option on the PSBGEN statement can be omitted.
- There is no specific maximum size for an individual PSB that can be used in CICS/ESA. However, a limitation is set by the PSB pool size in DBCTL.

Additional Processing Intent Options (PROCOPT)

Two additional types of processing intent can be specified with the PROCOPT= keyword of the PCB or SENSEG statement. The two additional options are O and E.

PROCOPT=O specifies “read without integrity”; no dynamic enqueue is done by resource lock management for calls against the database. With the G intent option,

you can specify GO, GOP, GON, or GOT. This option is only valid for the PCB statement. Read note number 1 under **Attention** for more information.

PROCOPT=E forces exclusive use of this database or segment by the online transaction. Other application programs scheduling a PSB referring to this database/segment wait during their scheduling process. No dynamic enqueue by resource lock management is done, but dynamic logging of database updates will be done. PROCOPT E can be specified with G, I, D, R, and A.

Attention:

1. The PROCOPT=O option affects integrity in reading and using uncommitted data. When you specify this option, IMS does not check the ownership of the segments returned. This means that the read-only user might have access to a segment that had been updated by another user. If the updating user then abends and is backed out, the read-only user would have seen a segment that no longer exists in the same form in the database. Consequently, if you specify this option, do not use the data that is read as a basis for updating records in any database.
2. An abend might occur with PROCOPT=GO if another program updates pointers when this program is following the pointers. Pointers are updated during insert, delete, and backout operations.
3. If PROCOPT=O in the PCB statement, the SENSEG statement must not specify a PROCOPT of I, R, D, H, or A.
4. If the O or E option is used, it must be coded immediately after the associated function code; for example GO, not OG.

Related reference:

 Program Specification Block (PSB) Generation utility (System Utilities)

Using sequential buffering

You can use sequential buffering with CICS.

To use sequential buffering with CICS, you need to do two things:

- Put an SBONLINE control statement in the //DFSVMxx file. SBONLINE allows sequential buffering to be used.
- Specify programs that are to use sequential buffering. You can do this by coding during PSBGEN an SB= keyword on the PCB macro.

Related reference:

 Full-function or Fast Path database PCB statement (System Utilities)

 Specifying sequential buffering for an online system (System Definition)

CICS connected to DL/I

CICS can provide DL/I database support by using IMS.

There are two ways in which you can use DL/I support with CICS:

- Through CICS remote DL/I support, also known as function shipping
- Through DBCTL

For CICS/ESA, remote DL/I support and DBCTL support are included in the pregenerated CICS.

Configuring CICS CTL connections to IMS DBCTL systems

To connect CICS to databases managed by IMS DB, you must define the CICS resources and initialize the IMS DBCTL system.

- Define CICS resources

If you use DBCTL exclusively, define the PSBs and DMBs in IMS by using either the APPLCTN and DATABASE macros or, in systems that have dynamic resource definition enabled, the commands CREATE DB and CREATE PGM.

If you want to function-ship requests to a remote CICS system, in which the database manager can be DBCTL or remote DL/I (function-shipping), you need to generate a PDIR only.

CICS routes DL/I requests to remote DL/I or DBCTL according to the PSB that is named. If the PSB does not appear in the CICS PDIR, and CICS is connected to DBCTL, CICS routes the request to DBCTL.

Related reading: Refer to *CICS Transaction Server for z/OS CICS System Definition Guide* and *CICS Transaction Server for z/OS CICS Resource Definition Guide* for information on defining PDIRs and DDIRs.

- Initialize IMS DBCTL

Use the procedure library member DBC to initialize the DBCTL subsystem. Also generated are procedures for DBRC and DL/I, which are used to initialize the DBRC and DL/I address spaces. The DBRC and DL/I procedures are started automatically by DBCTL during DBCTL address space initialization. All three procedures use positional parameters on the EXEC statement:

```
PARM='region type, parm1, parm2, parm3,...'
```

The region types specified are:

```
PARM='DBC' for DBCTL  
PARM='DRC' for DBRC  
PARM='DLS' for DLISAS
```

When all three address spaces have been initialized successfully, DBCTL issues:

```
DFS989I DBCTL READY
```

When the “READY” message is received, the IMS console operator enters a /START command. The commands options are:

- /NRESTART CHECKPOINT 0 for a cold start with no previous shutdown
- /NRESTART for a warm start
- /ERESTART for an emergency restart after a failure

Related reading: See *CICS Transaction Server for z/OS IMS Database Control Guide* for additional details.

CICS tasks

The database resource adapter (DRA) is the interface between DBCTL and the transaction management subsystem. You need to put the DRA startup parameter table (DFSPZPxx) and the DRA startup router program (DFSPRRC0) in the CICS STEPLIB data set.

This topic describes CICS tasks associated with using DBCTL and CICS. The high level tasks include:

1. Define DRA resources

Example:

```
//STEPLIB DD DSN=CICSTSxx.CICS.SDFHAUTH,DISP=SHR  
// DD DSN=IMS.SDFSRESL,DISP=SHR
```

Related reading: For more information on the DRA startup table see *IMS Version 11 Application Programming APIs*, or “CCTL exit routines” in *IMS Version 11 Exit Routines*.

2. Connect CICS to DBCTL

After CICS has been started, the CICS operator can issue a CONNECT command to DBCTL through the CDBC transaction. This transaction is also used for disconnecting from DBCTL. The CDBI transaction provides the status of the connection.

Related reading: For more information on these transactions, see *CICS Transaction Server for z/OS IMS Database Control Guide*.

Chapter 42. Accessing IMS databases through the ODBA interface

Open Database Access (ODBA) provides a callable interface that enables any z/OS recoverable, resource-managed z/OS address space to issue DL/I database calls to an IMS DB subsystem.

The interface provides z/OS application programs (hereafter called ODBA applications) access to full-function DL/I databases and data entry databases (DEDBs). The ODBA application and IMS must coexist on the same z/OS image.

Application programs that use the ODBA interface can use the APPLCTN macro or, if dynamic resource definition is enabled, the CREATE PGM command to define the PSB names required by the ODBA applications.

The ODBA interface allows IMS DB and ODBA application programs to be developed, installed, and maintained independently. This independence provides failure isolation and resource recovery by using z/OS Resource Recovery Services (RRS).

You can achieve further failure isolation by configuring ODBA to use the CSL Open Database Manager (ODBM). ODBM can prevent an IMS abend 0113 from occurring if an ODBA application terminates unexpectedly during DL/I processing. You can configure ODBA to use ODBM by specifying the ODBMNAME and IMSPLEX parameters with the other IMS database resource adapter (DRA) interface startup parameters in the DFSxxxx0 member. The ODBA interface uses the DRA interface to communicate with IMS DB.

The ODBA interface resides in the z/OS address space and is recognized by IMS as an application region (hereafter called the z/OS application region).

Related concepts:

“RRS and distributed syncpoint/protected conversations” on page 40

Creating the ODBA DRA start-up table

Create the ODBA DRA startup table using the DFSPRP macro.

The startup table uses the DSNNAME to dynamically allocate the data set that contains the rest of the ODBA interface routines. The DDNAME is generated to allow multiple connections to IMS from the same z/OS application region. If you specify the DDNAME on the DFSPRP macro, it is ignored.

The default DSNNAME is IMS.SDFSRESL. This is the default name established by the IMS definition process. Make sure this data set is APF authorized.

Note: IMS.SDFSRESL library does not contain an DRA startup table. You must generate your own table by using the DFSPRP macro. Name the load module based on the following naming conventions:

- Characters 1-3 = DFS
- Characters 4-7 = specified 4-byte ID

The 4-byte ID should be the IMSID of the IMS system to which you will connect. However, this is not a requirement.

- Character 8 = 0 (zero)

Ensure that the DRA startup table module name is not the same as the name of an existing IMS module. To prevent accidental overlay, put the module in a load library that is accessible by the z/OS application region and not by the IMS region.

If you use a different library for your own versions, make sure that the library is APF authorized. The DRA callable interface dynamically allocates the library by using the data set name specified in the DRA startup table. The DDNAME is generated to allow multiple connections to IMS from the same z/OS application region.

The DFSPRP macro is documented in *IMS Version 11 System Programming APIs*.

Loading and running the ODBA and DRA modules in the z/OS application region

Place the following ODBA and DRA modules in the STEPLIB or JOBLIB in the z/OS application region. These modules are shipped with IMS in IMS.SDFSRESL.

DFSCDLI0

This module is bound or loaded by an application program. DFSCDLI0 also contains the ALIAS name AERTDLI.

DFSAERG0

This module is loaded by DFSCDLI0.

DFSAERM0

This module is attached by DFSAERG0 in the z/OS application region.

DFSAERA0

This module is attached by DFSAERM0 for initialization to the specified IMS DB subsystem.

Binding application programs

Bind the ODBA application programs with DFSCDLI0 (AERTDLI). As an alternative, you can issue a load and branch command passing the AIB call list in Register 1.

Related reference:

 The AERTDLI interface (Application Programming)

Establishing and defining security

IMS provides several options for establishing and defining security for application programs that use the ODBA interface.

The options that you select depend on the type of security environment and authorization method that you plan to use. In general, the process that IMS uses to secure PSBs involves one of the following types of security checking:

- Resource access security (RAS)

A security check is performed by RACF to determine if the user is authorized to use the PSB. RACF determines authorization by looking at the RACF security class profile defined for the dependent region.

- APSB security

A security check is performed to determine if the user is authorized to use the PSB.

The following table identifies the values that you need to specify to control data access for specific security implementations. The table also indicates the type of security checking that is performed for each set of specifications.

Table 130. Options for controlling data access for applications that use ODBA.

Security implementation	Authorization method	ISIS= specification	ODBASE= specification	Connection security	PSB security
Resource access security	RACF	R	N		X
	Resource Access Security exit routine (DFSRAS00)	C	N		X
	RACF and Resource Access Security exit routine (DFSRAS00)	A	N		X
	None	0 N	N		
APSB security	RACF	Not applicable	Y		X

To control data access for application programs that use the ODBA interface, follow the techniques discussed in the following topics to establish connection security and PSB security checking.

Related concepts:

 [IMS security \(System Administration\)](#)

Related reference:

 [Resource Access Security exit routine \(DFSRAS00\) \(Exit Routines\)](#)

RAS security

Use the ISIS and ODBASE execution parameters to control the authorization for a z/OS application region to use a PSB.

The following table describes the actions that you need to perform to set up security when specific options are selected.

Table 131. Options for defining RAS security for applications that use ODBA

Specifications	Actions to perform
ISIS=0 N and ODBASE=N	No action required. No PSB security checking is performed.
ISIS=R and ODBASE=N	Define the PSBs that you want protected by RACF to the IIMS or Ixxxxxxx resource class, and then define the user IDs of the dependent region that you want authorized to access the PSBs. The ODBA support for IMS will use the security environment (ACEE) passed in the dependent region's task (TCBSENV), if present, or the dependent region's address space (ASXBSENV), if the ACEE is not present at the task level.

Table 131. Options for defining RAS security for applications that use ODBA (continued)

Specifications	Actions to perform
ISIS=C and ODBASE=N	Create a Resource Access Security exit routine that is named DFSRAS00. This routine must determine if the user is authorized to use the PSB.
ISIS=A and ODBASE=N	1. Define the PSBs that you want protected by RACF to the IIMS or Ixxxxxxx resource class, and then define the user IDs of the dependent region that you want authorized to access the PSBs. The ODBA support for IMS will use the security environment (ACEE) passed in the dependent region's task (TCBSENV), if present, or the dependent region's address space (ASXBSENV), if the ACEE is not present at the task level. 2. Create a Resource Access Security exit routine that is named DFSRAS00. This routine must determine if the user is authorized to use the PSB. RACF is called first, and then the exit routine is called.

Defining APSB security

Use the ODBASE execution parameter to control the authorization for a user to use a PSB.

To set up APSB security when ODBASE=Y:

1. Define the PSBs that you want protected by RACF to the AIMS or Axxxxxxx general resource class (where xxxxxxx is the value specified on the RCLASS= parameter of the IMS SECURITY macro).
2. Specify RCLASS=IMS | RACFCOM | RACFTerm | RASRACF | RAS on the IMS SECURITY macro at IMS system definition time.

Part 10. Open Transaction Manager Access (OTMA)

The following topics about Open Transaction Manager Access (OTMA) are for IMS system and Transaction Manager administrators responsible for installation, design, customization, operation, and recovery procedures for OTMA servers or clients.

Chapter 43. Introduction to OTMA

IMS Open Transaction Manager Access (OTMA) is a transaction-based, connectionless client/server protocol.

Though easily generalized, its implementation is specific to IMS in a z/OS sysplex environment. The domain of the protocol is restricted to the domain of the z/OS cross-system coupling facility (XCF).

OTMA addresses the problem of connecting a client to a server so that the client can support a large network, or a large number of sessions, while maintaining high performance.

Other solutions available today use network-based protocols, such as Systems Network Architecture (SNA). These protocols require a great amount of overhead because they are not transaction based.

What is OTMA?

OTMA has similarities to network protocols.

Several architectural models for networks exist. The following figure shows two. The simplified four-layer model shown on the right is often used in descriptions of UNIX networks. In the open systems interconnection (OSI) model, shown on the left, OTMA is the session layer. Both models have a Transport, Network, and Data Link layer. The OSI model also includes layers for Application, Presentation, and Session, and the simplified model includes a process layer. In the four-layer model, OTMA is the process layer.




Figure 125. Network architecture models

OTMA, however, does not exactly conform to the OSI model, because OTMA can process several sessions simultaneously using a single transport connection, if the following are true:

- The z/OS cross-system coupling facility (XCF) is the transport layer.
- A session is the connection between IMS and a client.
- A client or server only creates a single XCF connection.

OTMA performs some of the basic functions of the OSI transport layer (those not performed by XCF), so it is simplest to think of OTMA as a combined session and transport layer, with the transport layer comprised of both XCF and OTMA.

Although you can think of OTMA as a session and transport layer in a network architecture model, OTMA is designed to be a high-performance comprehensive protocol that allows z/OS programs to access IMS applications.

Definitions: A *z/OS program* in this case means any z/OS application that is a member of an XCF group that includes IMS. The XCF group members that IMS communicates with are called *OTMA clients*.

By using OTMA, each client (z/OS application) can submit transactions to IMS or issue IMS commands and receive output from IMS application programs and from IMS itself.

Definition: Because IMS can communicate with, or serve, many OTMA clients, IMS is called the *server*. However, OTMA only operates in the following IMS environments:

- IMS TM and DB (the IMS DB/DC environment)
- IMS TM with DB2 for z/OS (the IMS DCCTL environment)

Related concepts:

Chapter 45, “The OTMA client,” on page 813

Capabilities of OTMA

OTMA capabilities include support for IMS applications, IMS and OTMA commands, resource monitoring, and more.

An OTMA client can issue most IMS commands, including both type-1 and type-2 commands, and receive responses as a result of those commands.

You can specify a level of security for each OTMA client individually. Alternatively, you can indicate that no security checking is to be done for its messages, thereby minimizing security-processing overhead.

The OTMA message flow and synchronization point protocols can be modified by an OTMA client for each transaction. In other words, the transaction-processing protocol used is not dependent on the current session.

OTMA can monitor how well IMS is processing incoming transactions and alert the client when the processing ability of the IMS system appears degraded.

OTMA can monitor transactions and cancel them if a specified transaction expiration time is exceeded before they are processed.

You can define a limit to the number of input messages from any one OTMA client that can be waiting at the same time to be processed. This prevents OTMA

messages from causing an abend as a result of using too much storage. OTMA issues warning messages as the message count approaches the limit. When the message count reaches the limit, OTMA identifies a *message flood condition* and rejects any new messages from the OTMA clients until the message count drops or a new limit is set.

You can stop incoming transactions and commands from individual OTMA clients without stopping the connection or affecting any transactions or commands that IMS is already processing.

OTMA also supports callout requests from IMS application programs running in IMS dependent regions. Callout is invoked with the DL/I ICAL call. Callout requests allow IMS application programs to request data or services from servers that are outside of the IMS installation. Depending on whether the callout request is processed synchronously or asynchronously, OTMA support for callout requests differs.

The IMS /DISPLAY TRANSACTION command output is in the form of an OTMA message returned to the client in the application-data section of the message prefix.

OTMA-initiated transactions are identified to z/OS Workload Manager using the OTMA transaction-pipe name, which identifies the logical connection between IMS and OTMA.

Existing IMS application programs that have previously not been used with OTMA can run without modification and interact with OTMA clients. APPC/IMS application programs that use IMS SETO calls might need some modification.

Related concepts:

Chapter 46, "OTMA support for callout requests," on page 845

"OTMA support for transaction expiration" on page 797

"OTMA resource monitoring" on page 786

Related reference:

"OTMA restrictions" on page 785

Benefits of using OTMA

IMS applications that do not use OTMA do not need to be modified to use OTMA, unless they use SETO calls, in which case, you might need to modify them to use OTMA.

The SETO call applies to APPC/IMS and SPOOL/API processing.

For each OTMA-originated transaction, the SETO call returns a status code. You can bracket the SETO call with an INQY call if necessary.

Full-duplex processing provides an environment in which transactions and output messages are sent and processed in parallel.

You can implement IMS device support outside IMS. You can also implement device support for your IMS subsystem that is different from what IMS provides, or enable device support that IMS does not provide. The following figure illustrates how IMS can communicate with a device, shown here as a workstation, using device support implemented within an OTMA client. IMS device support using Virtual Telecommunications Access Method (VTAM) is shown for

comparison.




Figure 126. IMS communicates with a device using device support implemented within an OTMA client

Flow-control and transaction-processing attributes are dynamically bound to the transaction.

Clients have high-performance access to IMS:

- OTMA uses the z/OS cross-system coupling facility application programming interface (API).
- OTMA does not use VTAM and IMS device-dependent support.

Transactions based on different protocols (that is, that have different processing requirements such as being recoverable or irrecoverable) can be associated with a single transaction pipe.

You can connect up to 255 clients to the OTMA group.

Messages can be extended using the user-data section of the message prefix, allowing additional user information to be sent with the transaction.

User information and transaction pipe name are included within the messages themselves.

Different clients can specify the same transaction pipe names, instead of needing to use uniquely named resources.

You do not need to use networking architectures, such as SNA (Systems Network Architecture).

Related concepts:

“Using transaction pipes with OTMA” on page 761

Related reference:

“Using DL/I calls in an OTMA environment” on page 806

Advantages of the OTMA protocol

OTMA treats transactions as data objects that have attributes independent of application-, session-, or transport-layer considerations. OTMA is, in effect, a transaction layer, independent of other layers.

As a unique layer, OTMA offers flexibility, simplicity, and performance that other solutions do not offer. This section outlines the transaction-specific services that OTMA provides the client.

Grouping of transactions using transaction pipes.

Security options (for example, the client can verify security or let the server verify the user ID).

Dynamically-bound flow control and processing. The client can decide how transaction requests and responses are to be processed by the server.

The ability for the client to query the server for transactions that the server supports.

Treating transactions as objects. The client can include any pertinent user data with the transaction, and allow that data to stay with all messages generated by the transaction.

The ability for the client to specify a client token with each transaction to correlate input with output.

The ability for the client to control transaction processing performed by the server, in terms of:

- The client can set a different level of security than that of the OTMA group to which it belongs.
- To improve performance, the client can eliminate the security-checking that the server performs on the messages it delivers.
- Transaction grouping, using the transaction-pipe token.
- OTMA clients that use `synclevel=confirm` or `synclevel=synchpt` can specify the timeout value for send-then-commit messages.

Client routing. An OTMA destination descriptor can be easily coded in the DFSYDTx PROCLIB member to reroute an output message that is inserted to an alternate PCB to any OTMA client or to IMS. Alternatively, an IMS exit routine can be coded to reroute output messages that are inserted to an alternate PCB.

Architected command output. The client can use the IMS /DISPLAY TRANSACTION command to query the server's transaction attributes and receive the reply in a structured format. Therefore, the need for automated operator scripting to control processing is reduced.

Unlike APPC, when using message flow through transaction pipes, no concept exists of a session that contains the flow-control parameters for all transactions and associated output data for the session.

How IMS messages flow in an OTMA environment

The key to message flow for OTMA is the *transaction pipe*, the logical connection between the server and the OTMA client.

An OTMA client includes the transaction-pipe name in the message-control information section of the message prefix for the input message. IMS then associates application output for an OTMA client with a specific transaction pipe.

Related concepts:

“Using transaction pipes with OTMA” on page 761

Basic OTMA message flow

The basic message flow for OTMA clients is relatively simple.

The basic message flow is:

1. The client submits a transaction or command to IMS.
2. IMS accepts IMS transactions as input from any client.

The IMS transaction code is specified in the application-data section of the input message.

If the client is submitting an IMS command, the command is included in the application-data section of the input message.

3. The input message is processed.

An IMS transaction is enqueued to the appropriate application program using an IMS scheduler message block (SMB).

An IMS command is processed by IMS. The output is sent to the client synchronously or asynchronously, depending on the type of request.

4. Application output is sent to the client.

Generation of output and commit are coordinated based on the commit mode specified in the state-data section of the message prefix for the input message.

The application output is enqueued to a dynamically created IMS transaction-pipe structure (specific to that client) before being sent to the client.

For an OTMA-submitted transaction, IOPCB output is returned to the OTMA client. By default, all alternate PCB output is also sent to the OTMA client. You can change this by coding either an OTMA destination descriptor in the OTMA DFSYDTx PROCLIB member or by coding the OTMA Prerouting exit routine (DFSYPX0) or the client's OTMA Destination Resolution exit routine (DFSYDRU0). You can also use these exit routines to route alternate PCB output from non-OTMA-submitted transactions to OTMA clients.

IMS delivers segmented messages in order, even though the z/OS cross-system coupling facility (XCF) does not guarantee sequential delivery of messages.

The following figure shows an example of the message flow in an OTMA environment. Two clients are shown side by side in the example; they can be a TCP/IP client, a WebSphere MQ Queue Manager client, or a client of any other network type. Message flow starts with the client, goes through the XCF group, and to IMS. Within the IMS address space, a control region contains OTMA; the message flow ends at a transaction-pipe. The IMS application program issues a Get Unique (GU) call in the dependent region.




Figure 127. IMS message flow in an OTMA environment

In the preceding figure:

1. The message prefix is always attached to the input transaction, even in the case of segmented input. This prefix contains important information, such as the transaction-pipe name and the client token.

A client application program can send several transactions specifying the same transaction-pipe name. The client token must always be present in the prefix, so that the client application program knows how to process the IMS output it receives.

2. OTMA clients do not need to predefine transaction pipes. Two different clients can use the same transaction-pipe name. Although many clients can use the same transaction-pipe name, each transaction pipe is unique. client 1 and client 2 both use tpipe1, yet each is a unique transaction pipe.

A client can create and use as many transaction pipes as it needs.

3. The transaction-pipe structure is created dynamically when OTMA receives output and is used as an anchor for the application output.
4. The IMS application program has no knowledge of the OTMA message prefix when it issues the GU call.

IMS supports a full-duplex message flow for a client/server session. The client can instead request a half-duplex message flow, but this flow must be implemented and managed by the client itself:

- A correlator token in the state-data section of the message prefix can be used to uniquely identify a transaction. IMS maintains this field in the message prefix for a transaction.
- The client can set the response-requested flag in the message-control information section of the message prefix to receive a response for a message.
- Any unsolicited output from IMS is easily identified by a client, because the message prefix specifies only the transaction-pipe name. The client can ask IMS to discard the output.

Unsolicited output should not interfere with half-duplex processing. That is, the client must be prepared for full-duplex flows while still maintaining a half-duplex flow on a user-token level. Contention should not be an error condition.

Sample commit-then-send transaction processing flows

The following figure shows a non-OTMA environment: a secondary logical unit type 2 (SLU 2) device communicates with IMS using VTAM and IMS device support (DDMs). The transactions are enqueued to the IMS message queues. Transaction output is returned to the SLU 2 device.




Figure 128. Standard SLU 2 transaction flow

The following figure shows the same transaction flow in an OTMA environment. The transaction still comes from a SLU 2 device, but the device communicates with IMS using an OTMA client, through an z/OS cross-system coupling facility (XCF) group, rather than VTAM.

The figure only shows the input flow, which begins with the SLU 2 device, goes to the OTMA client, through the XCF group, and ends at the OTMA server. The transaction is placed on the message queue, and the application issues get unique, insert, and get unique calls. Output follows the same path, in reverse. Of course, if a client is to send output to the SLU 2 device, the SLU 2 device must be defined to the client, and the client must be able to drive that device.




Figure 129. SLU 2 transaction flow using OTMA

It might seem that the OTMA flow is more complex, and for a SLU 2 device, perhaps it is. But you can use OTMA to allow any type of device to communicate with IMS, not just VTAM-supported devices. An OTMA client can also act as a gateway for another network, such as a TCP/IP network.

Using transaction pipes with OTMA

An IMS transaction represents a request for IMS to do some work. Many transactions also require a response after IMS has completed the work. So, each transaction has a source (the requester) and often a destination (for the response).

IMS uses the concept of a logical terminal (LTERM) to ensure that responses are associated with the correct requesters. An LTERM uses a queue where the transaction output is kept before it is returned to the requester.

Definition: A *transaction pipe* (tpipe) is a logical connection between a client and the server. It is analogous to an IMS logical terminal (LTERM). For each LTERM, IMS maintains a connection between the queue and the physical node that receives the output. OTMA does not use an LTERM but still must maintain a connection between the client and IMS. This connection is the transaction pipe, or tpipe.

Transaction pipes enable a client to associate its transactions with a transaction-pipe name. IMS uses the transaction-pipe name to associate all input and output with a particular client. The association between the transaction output and its ultimate destination (for example, a user at a terminal or a printer) is not made within IMS (as is the case with LTERMs), but is the responsibility of the client.

By using a transaction pipe, IMS does not know anything about the actual user of the transaction, often a user of the client application. Because IMS does not know anything about the actual user, the client has complete control over the output of transactions.

OTMA's use of transaction pipes provides:

- **Flexibility**

Many transaction outputs can flow through the same transaction pipe.

- **Performance**

Transaction pipes give the client the ability to specify and distinguish transactions based on their message-flow control and synchronization.

- **Resynchronization between a client and IMS**

Transaction pipes can be either synchronized or non-synchronized. For a synchronized transaction pipe, all output messages are serialized through a single process, and sequence numbers can be assigned to messages. By logging these serialized messages, IMS and the client can resynchronize in the event of an outage.

No resynchronization is required for a non-synchronized transaction pipe.

- **Object orientation**

A transaction can be thought of as an object because OTMA keeps the transaction message information (such as user data and transaction-pipe name) within the message.

IMS removes transaction pipes after they have been idle for three consecutive system checkpoints, except in the following circumstances:

- Commit-then-send messages are queued on the tpipe or the tpipe hold queue
- The tpipe is stopped A trace is set on the tpipe
- The tpipe is a synchronized tpipe, such as a tpipe used by MQSeries for commit-then-send input transactions
- The tpipe is in a WAIT state for a resume tpipe request that specified either the AUTO or the SINGLE-WAIT options

The following figure illustrates how transaction pipes fit in an OTMA client/server environment. As shown in the figure, transaction-pipe structures reside in the OTMA layer only for the server. z/OS cross-system coupling facility, which resides in the transport layer, can be thought of as an interprocess communication layer, because it provides communication between the client process and the server process.




Figure 130. How transaction pipes fit in an OTMA client/server environment

Differences in transaction pipes

IMS LTERMs and UNIX pipes both provide a one-way flow for message traffic. An OTMA transaction pipe provides a two-way flow.

The concept of a transaction pipe is applicable to any protocol. In a general way, the transaction pipe replaces the IMS LTERM because:

- Processing is full duplex.
- Multiple flow-control mechanisms are possible.
- The logical output entity (in other words, the LTERM) is dissociated from the node of the actual user.
- The transaction pipe is implemented as a protocol rather than as an API, which facilitates a client/server architecture.
- The transaction pipe sets up a data-control mechanism independent of session characteristics, and is therefore transaction specific.

Message flow using transaction pipes

The flow control of transactions is handled by the client.

The client dynamically binds flow-control parameters to the transaction by querying the transaction attributes in the server. Transaction pipes are not usually associated with flow control (except for synchronized transaction pipes using half-duplex processing).

The following figure shows the basic message flow between a client and a server, using the z/OS cross-system coupling facility (XCF). The order of processing is:

1. The client sends a transaction as input to the server (IMS).
2. The server returns transaction output messages to the client.




Figure 131. Basic transaction-pipe message flow

Within the server, the input transaction and the output messages are organized and synchronized using IMS queues, as shown in the following figure. The figure illustrates a commit-then-send transaction flow for a non-Fast Path environment.

The order of processing is:

1. The client sends a transaction to the server, and the server enqueues the transaction on a message queue.
2. The transaction is submitted to an application program for processing.
3. The application program prepares any output for the transaction and commits the output during sync-point processing.
4. The output is returned to the client.




Figure 132. Use of queues in the transaction-pipe message flow

In a full-duplex environment, transactions and output messages are being sent and processed in parallel, as shown in the following figure. This parallelism can be maximized by creating a process for every transaction and output message. The order of processing is:

1. The client sends a transaction (Tran 1) to the server, and the server's transaction pipe enqueues the transaction.
2. The transaction (Tran 1) is submitted to an application program for processing.
3. The application program enqueues any output (Message 1) for the transaction (Tran 1).
4. The client sends a second transaction (Tran 2) to the server and the server's transaction pipe enqueues the transaction.
5. The second transaction (Tran 2) is submitted to an application program for processing.
6. The output (Message 1) for Tran 1 is returned to the client.
7. The application program enqueues the output (Message 2) for the second transaction (Tran 2).
8. The output (Message 2) for Tran 2 is returned to the client.




Figure 133. Transaction-pipe flow in full-duplex environment

Q: Does specifying a transaction pipe as synchronized make the communication flow half-duplex?

A: No. Transaction pipes are always full-duplex.

Whether the communication flow is actually half-duplex depends on the client. For a synchronized transaction pipe, IMS processes all output messages in the order received. No new messages are sent for the transaction until IMS has received an ACK message for the previous message. A NAK message causes IMS to halt all output processing for that transaction.

While this output processing is taking place, the client could be sending new input transaction messages to IMS on that synchronized transaction pipe. If the client coordinates the sending of transactions with the receipt of IMS output, the client can effect half-duplex processing.

Related concepts:

"OTMA commit processing" on page 817

Chapter 44. Enabling and using OTMA

The following topics describe IMS tasks related to the OTMA environment, as well as how IMS operates in an OTMA environment.

Enabling OTMA

To enable IMS to use OTMA, specify the z/OS cross-system coupling facility (XCF) group name and IMS OTMA member name during system definition.

OTMA is installed with IMS TM. The IMS INSTALL/IVP Dialog is not used to install OTMA.

To start OTMA, you can use the OTMA=Y startup parameter in the IMS procedure during IMS system definition or, after an IMS restart, issue the type-1 command /START OTMA.

You can define various OTMA attributes during IMS system definition, including z/OS cross-system coupling facility (XCF) names, startup behavior, security, and so on.

Related reference:

-  DFSPBxxx member of the IMS PROCLIB data set (System Definition)
-  Parameter descriptions for IMS procedures (System Definition)
-  IMS procedure (System Definition)
-  DCC procedure (System Definition)

Summary of the OTMA configuration parameters

All the following OTMA configuration parameters can be specified in the DFSPBxxx member of the IMS.PROCLIB data set and either the IMS or DCC startup procedures.

GRNAME=

Defines the name of the XCF group that IMS creates and joins for OTMA communications.

OTMA=

Specifies whether OTMA starts during IMS startup and whether the /START OTMA command is recoverable across warm and emergency restarts.

OTMAASY=

For send-then commit messages, controls the synchronous or asynchronous scheduling of transaction originating from a program-to-program switch.

OTMANM=

Specifies the XCF member name that identifies OTMA within the XCF group.

OTMAMD=

Determines whether the member override field in the parameter list of the OTMA Destination Resolution exit routine (DFSYPX0) can be used to specify a different XCF member name for transactions invoked from an OTMA client.

OTMASP=

For WebSphere MQ, specifies whether a synchronized tpipe is used for OTMA output.

OTMASE=

Specifies the type of OTMA RACF security that you want to use, if any.

Related tasks:

“OTMA security” on page 799

Related reference:

 DFSPBxxx member of the IMS PROCLIB data set (System Definition)

 Parameter descriptions for IMS procedures (System Definition)

 IMS procedure (System Definition)

 DCC procedure (System Definition)

Defining the XCF group name

OTMA communications require XCF. Both IMS and the OTMA clients must join the same the XCF group.

You specify the XCF group name for OTMA communications on the GRNAME parameter in either the DFSPBxxx PROCLIB member or the IMS control region JCL. If a valid XCF group name is not specified, OTMA cannot be started.

When the GRNAME parameter is specified, IMS joins the XCF group only when OTMA is started. If you have configured OTMA to start when IMS starts, IMS joins the XCF group during startup. If OTMA is not configured to start when IMS starts, IMS joins the XCF group only after OTMA is started by the /START OTMA command.

If you specify GRNAME= and OTMA is started, you can use the /DISPLAY OTMA command to display the XCF status. You are not required to define any XCF information.

If you use RACF for security, the IMSXCF.group.member (client member name) must be defined in the RACF FACILITY class.

Related reference:

 Parameter descriptions for IMS procedures (System Definition)

Defining the OTMA XCF member name

XCF requires an XCF member name for each member of an XCF group.

You can specify an XCF member name for OTMA by using the OTMANM parameter. If you do not specify the OTMANM parameter, OTMA uses the APPLID of the IMS system as the OTMA XCF member name.

The OTMANM parameter is not used in Extended Recovery Facility (XRF) or Remote Site Recovery (RSR) installations. Instead, IMS uses a USERVAR name as the XCFmember name. The USERVAR name is specified in either the IMS control region JCL, the DFSPBxxx PROCLIB member, or the DFSHSBxx PROCLIB member.

Recommendation: Do not change the group name or the IMS member name during an IMS emergency or normal restart.

Related reference:

 [Parameter descriptions for IMS procedures \(System Definition\)](#)

Defining when OTMA starts up

You can configure OTMA to start when IMS starts or you can start OTMA after IMS is running.

Prerequisite: Before you can start OTMA, you must first specify the GRNAME execution parameter. The GRNAME parameter identifies the XCF group that OTMA uses for communications.

- To start OTMA during IMS startup, specify OTMA=Y in either the DFSPBxxx PROCLIB member or the IMS control region JCL.
- To start OTMA after IMS is running, specify OTMA=N, OTMA=M, or omit the OTMA parameter completely. After IMS is running, issue the /START OTMA command.

Note: The /START OTMA command is recoverable. That is, if OTMA was started by the /START OTMA command, IMS disregards an OTMA=N specification and restarts OTMA during warm and emergency restarts, even if OTMA=N or the OTMA parameter is not specified.

- To prevent IMS from restarting OTMA during a warm or emergency restart after the /START OTMA command has been issued, specify OTMA=M. When OTMA=M is specified, IMS does not recover the /START OTMA command during warm or emergency restarts.

Related reference:

 [Parameter descriptions for IMS procedures \(System Definition\)](#)

Defining the level of OTMA security checking

If you use a security product such as RACF, you can specify different levels of security checking by using the OTMASE parameter.

The value specified in the OTMASE parameter can be overridden when IMS is running by issuing the /SECURE OTMA command.

The levels of security checking include:

CHECK (C)

IMS commands are checked against the CIMS class. IMS transactions are checked against the TIMS class.

FULL (F)

The same type of security as CHECK, but additional checking is performed against dependent regions. F is the default value for the OTMASE parameter.

NONE (N)

OTMA RACF security is NONE. No calls to RACF are made.

PROFILE (P)

Each OTMA message defines the level of security checking to be done.

Related reference:

 [Parameter descriptions for IMS procedures \(System Definition\)](#)

 [/SECURE command \(Commands\)](#)

Specifying synchronized tpipes for WebSphere MQ

An OTMA client, such as WebSphere MQ, can require that an OTMA tpipe is synchronized for ALTPCB output.

You can specify synchronized tpipes in two ways: by setting the output flag in the OTMA User Data Formatting exit routine (DFSYDRU0) or by specifying OTMASP=Y. If the only reason you code the DFSYDRU0 exit routine is to set the synchronized output flag, you can use the OTMASP parameter instead.

By default, OTMA creates non-synchronized tpipes for the OTMA output.

If your organization uses both the DFSYDRU0 exit routine and the OTMASP parameter to control the tpipe that gets created for OTMA output, the following table shows when synchronized tpipes are created.

Table 132. Tpipes created when both OTMASP parameter and DFYDRU0 exit used

DFSYDRU0 is set to...	If OTMASP=Y, result is...	If OTMASP=N, result is...
Create synchronized tpipe	Synchronized tpipe	Synchronized tpipe
Create non-synchronized tpipe	Synchronized tpipe	Non-synchronized tpipe

Related reference:

-  Parameter descriptions for IMS procedures (System Definition)
-  OTMA Destination Resolution exit routine (DFSYPX0) (Exit Routines)

Enabling DFSYPRX0 member name override for OTMA clients

The parameter list for the OTMA Destination Resolution exit routine (DFSYPRX0) contains a member override field that allows you to route transaction received from non-OTMA LTERMs to OTMA clients.

By default, the member override field cannot be used to route transactions received from OTMA clients.

To use the member override field in the DFSYPRX0 parameter list for transactions received from both OTMA clients and non-OTMA LTERMs, specify OTMAMD=Y.

Specifying asynchronous delivery of program-to-program switch output messages

If you are using send-then-commit (CM1) messages that initiate multiple program-to-program switches, to ensure that only the appropriate output is returned to the OTMA client in synchronous CM1 mode, specify OTMAASY=Y.

When OTMAASY=Y is specified, if an input transaction triggers multiple program-to-program switches, only the response transaction is scheduled by OTMA synchronously.

Non-response transactions that originate from a program-to-program switch are scheduled asynchronously, regardless of whether they are returned to the client before the expected synchronous output. IMS does not issue DFS2082 messages for transactions that are scheduled asynchronously.

When transactions are scheduled asynchronously, IMS does not issue DFS2082 messages.

When OTMAASY=N is specified, or the OTMAASY parameter is omitted, if any output from a program-to-program switch is returned before the expected synchronous output, errors can occur.

Related reference:

 Parameter descriptions for IMS procedures (System Definition)

OTMA descriptors

OTMA provides two types of descriptors: an OTMA client descriptor and an OTMA destination descriptor.

The OTMA client descriptor specifies certain characteristics of an OTMA client, such as its DFSYDRU0 exit routine, its send-then-commit message time out values, or its message flood threshold value.

The OTMA destination descriptor defines destination attributes for:

- OTMA ALTPCB output sent to either IMS Connect or non-OTMA destinations such as an SNA terminal or printer.
- Callout requests from IMS applications programs that are sent to IMS Connect and that are ultimately destined for any of the following:
 - An enterprise Java bean (EJB) or message-driven bean (MDB) running under a web application server, such as WebSphere Application Server.
 - A Web service.
 - A user-written IMS Connect client.

Both OTMA client descriptors and OTMA destination descriptors are specified as entries in the DFSYDTx PROCLIB member.

OTMA destination descriptors can also be created and updated online by using the following type-2 IMS commands:

- CREATE OTMADESC
- DELETE OTMADESC
- QUERY OTMADESC
- UPDATE OTMADESC

Related reference:

 DFSYDTx member of the IMS PROCLIB data set (System Definition)

OTMA client descriptors

Use OTMA client descriptors to provide information about a specific OTMA client to IMS. OTMA client descriptor entries are identified in the DFSYDTx PROCLIB member by an M in column one of the descriptor entry.

The information in the OTMA client descriptor helps IMS manage messages from the OTMA client. The information included in an OTMA client descriptor can also be specified in an OTMA client-bid request or, in some cases, in a /START OTMA command. Information provided by these alternate methods might override the information provided in the OTMA client descriptor.

Using an OTMA client descriptor, you can specify to OTMA and IMS the following attributes of an OTMA client:

- The name of the OTMA Destination Resolution exit routine that the OTMA client uses.
- The timeout value that OTMA should use when waiting for an ACK or NAK response from the OTMA client for both commit-then-send messages and send-then-commit messages that use synclevel=confirm or synclevel=syncpt.
- The maximum number of input messages from the OTMA client that can be processing in an IMS system before triggering a message flood condition.

OTMA client descriptors are optional; however, they can be particularly useful for routing output to an OTMA client before the OTMA client has actually connected to OTMA. When the OTMA client finally does connect, the specifications in the client-bid request override the specifications made in the OTMA client descriptor.

The following defaults apply if you do not specify any of the parameters of the OTMA client descriptor:

- IMS uses the default OTMA Destination Resolution exit routine, DFSYDRU0, unless the client-bid request specifies a different exit routine name.
- For the acknowledgment timeout value, IMS uses a default of two minutes unless a different timeout value is provided in either the client-bid request or the TIMEOUT parameter of the /START TMEMBER command.
- For message flood detection, IMS uses a default threshold of 5000 messages, unless a different value is provided in either the client-bid request or the INPUT parameter of the /START TMEMBER command.

OTMA client descriptors are built during IMS initialization. The descriptors are included in DFSYDTx members of IMS.PROCLIB. The “x” on DFSYDTx is the IMS nucleus suffix. If you have multiple clients at the same IMS system, list each client name on a separate line. Ensure that columns 3-18 are different for each client.

You can define a maximum of 254 client descriptors. If you do not use the DFSOTMA descriptor, you can define 255 client descriptors.

Related reference:

 DFSYDTx member of the IMS PROCLIB data set (System Definition)

OTMA destination descriptors

Use OTMA destination descriptors to define destinations for messages that are routed through OTMA.

OTMA destination descriptors provide a simpler method for describing destinations than using the OTMA Destination Resolution exit routine (DFSYPX0) and the OTMA User Data Formatting exit routine (DFSYDRU0).

Note: When OTMA destination descriptors are used to describe a destination, IMS does not call the DFSYPX0 and DFSYDRU0 exit routines for messages sent to that destination. If you introduce a new OTMA destination descriptor that prevents the DFSYPX0 and DFSYDRU0 exit routines from being called for a given destination, make sure that any necessary processing performed by the exit routines for that destination is not lost when the exit routines are no longer called.

When OTMA destination descriptors are used to describe a destination, IMS does not call the DFSYPX0 and DFSYDRU0 exit routines. You can override this

behavior for descriptors with TYPE=IMSCON, TYPE=NONOTMA, or TYPE=MQSERIES by specifying EXIT=YES. In that case, the exits are called to process the message and determine whether the descriptor routing information must be modified.

OTMA destination types include IMS Connect destinations (TYPE=IMSCON) and non-OTMA destinations (TYPE=NONOTMA).

IMS Connect descriptors

OTMA destination descriptors support the routing of callout requests from IMS application programs to external data or service providers through IMS Connect. IMS Connect routes callout requests through one of the following IMS Connect clients:

- IMS TM Resource Adapter
- IMS Enterprise Suite SOAP Gateway
- User-written IMS Connect clients

For any of the following types of messages, define your destination type as IMS Connect (TYPE=IMSCON):

- Output for IMS Connect clients
- Callout requests from IMS application programs

If you are routing your messages to a destination such as an SNA terminal or a printer, define your destination type as non-OTMA (TYPE=NONOTMA).

The destination attributes that you can specify in the descriptor differ depending on the type of destination.

For IMS Connect destinations, you can specify the following destination attributes:

- An OTMA tmember name
- An OTMA super member name
- A tpipe name
- For IMS Connect XML conversion support for SOAP Gateway, an XML adapter name and an XML converter name.
- A timeout interval for synchronous callout requests that are made by IMS application programs that issue the DL/I ICAL call. If a response is not received in the specified amount of time, the request times out and the IMS dependent region is freed. If the timeout value specified on the OTMA destination descriptor differs from the timeout value specified by the IMS application on DL/I ICAL call, OTMA uses the smaller of the two values.

For IMS Connect destinations, the OTMA destination descriptor provides a layer of abstraction between the application program and the tpipes. This layer of abstraction provides flexibility in the routing of IMS Connect output. For example, the destination name specified in the descriptor can be explicit or generalized using a mask character. Also, the tpipe name specified in the descriptor can be an individual tpipe name, a super member name, or left blank, in which case the specified destination name is used as the tpipe name.

Related reference:

 DFSYDTx member of the IMS PROCLIB data set (System Definition)

Defining OTMA destination descriptors

You can create, delete, and update OTMA destination descriptors either by using the online type-2 IMS commands or by coding them directly in the DFSYDTx member of the IMS.PROCLIB data set.

You can define a maximum of 510 destination descriptors.

The type-2 commands for administering OTMA destination descriptors include:

- CREATE OTMADESC
- DELETE OTMADESC
- QUERY OTMADESC
- UPDATE OTMADESC

Changes made to the OTMA destination descriptors by any of the above type-2 commands become effective immediately and are recorded in both the IMS system logs as x'221B' log records and in the IMS checkpoint logs as x'4035' log records. The changes made by the type-2 commands are not stored in the DFSYDTx PROCLIB member and are retained only across warm starts and emergency restarts of IMS. The changes made by the type-2 commands are not retained across cold starts of IMS.

If you use the type-2 commands to create, modify, or delete OTMA destination descriptors, the changes override any existing definitions for the same descriptors that are in the DFSYDTx until the next cold start.

Changes made in the DFSYDTx member require a cold start of IMS to take effect, because IMS reads the DFSYDTx member only during a cold start.

In the DFSYDTx member, the OTMA destination descriptors are distinguished from the OTMA client descriptors, which are also stored in the DFSYDTx member, by a D in the first column of the descriptor entry.

You can update and delete OTMA destination descriptors that are coded in the DFSYDTx member with the type-2 commands; however the changes are not reflected in the DFSYDTx member. The modified or deleted descriptors are recorded in the logs, and the log records override the descriptor statements that are stored in the DFSYDTx member.

The QUERY OTMADESC command returns information about the OTMA destination descriptors as it is currently recorded in the logs.

Related reference:

 [CREATE OTMADESC command \(Commands\)](#)

 [DELETE OTMADESC command \(Commands\)](#)

 [QUERY OTMADESC command \(Commands\)](#)

 [UPDATE OTMADESC command \(Commands\)](#)

 [OTMA destination descriptor syntax and parameters \(System Definition\)](#)

Masking names of OTMA destination descriptors

When specifying a destination name in an OTMA destination descriptor, you can generalize destination names by using an asterisk as a wildcard character or as a mask of the final characters in the destination name field.

For example, if you specify OTICON* as the destination name, the descriptor applies to any output that specifies a destination name that begins with OTICON, such as OTICON01 and OTICON02.

The OTMA destination descriptors can be listed in any order in the DFSYDTx PROCLIB member; however, when IMS searches for a descriptor to route output, IMS searches the descriptors in order from the most specific destination name to the most general. That is, IMS always reads specific destination names before destination names that include an asterisk.

If a descriptor specifies a masked destination name, such as OTICON*, but does not specify a tpipe name, the output must specify the actual tpipe name in full. By default, if no tpipe name is specified on the descriptor, the destination name is used as the tpipe name. If the destination name is masked, then IMS must read the destination name specified by the output to determine the specific tpipe to use.

If you specify a tpipe name on a descriptor that uses a masked destination name, all output sent to destinations that match the masked destination name is routed to the same tpipe.

DFSOTMA descriptor

Use the optional DFSOTMA descriptor to define global and default attributes, limits, and types of support that apply to all OTMA clients.

The DFSOTMA descriptor is an OTMA type-M descriptor that sets system defaults and global parameters for all OTMA clients that connect to this IMS system.

The DFSOTMA descriptor does not support all of the parameters that you can specify on an OTMA client descriptor parameters. Parameters that are not supported, such as DRU=, T/O=, and INPT=, are ignored if they are specified in the DFSOTMA client descriptor.

The global values that you can specify for all OTMA clients include:

- Message flood protection
- CM1/CM0 ACK timeout value
- Global tpipe limit for all the OTMA members or clients

If DFSOTMA is used, IMS builds a type-M descriptor block to store the global information specified. The block counts toward the 255 maximum number of type-M descriptors, so when the DFSOTMA descriptor is used, the maximum number of OTMA client descriptors is 254.

The DFSOTMA descriptor is always loaded from the DFSYDTx PROCLIB member regardless of the IMS restart type. There are no checkpoint log records for the DFSOTMA descriptor.

Related reference:

 DFSOTMA descriptor syntax and parameters (System Definition)

OTMA-supported exit routines

Several exit routines delivered with IMS are specifically for OTMA. OTMA also supports several other exit routines delivered with IMS that are not specific to OTMA.

The following exit routines support OTMA:

- OTMA User Data Formatting exit routine (DFSYDRU0)
- OTMA Input/Output Edit exit routine (DFSYIOE0)

Restriction: The DFSYIOE0 exit routine is not supported for synchronous callout request messages, including those received by OTMA from IMS application programs that issue the DL/I ICAL call.

- OTMA Destination Resolution exit routine (DFSYPX0)
- OTMA Resume Tpipe Security exit routine (DFSYRTUX). You can use this exit routine to secure messages on the asynchronous hold queue.

In addition, the following exit routines are supported by OTMA:

- Command Authorization (DFSCCMD0)
- TM and MSC Message Routing and Control user exit (DFSMSCE0)
- Queue Space Notification (DFSQSPC0)
- Transaction Authorization (DFSCTRN0)

Related tasks:

“Securing messages on the asynchronous hold queue” on page 804

Related reference:

-  OTMA User Data Formatting exit routine (DFSYDRU0) (Exit Routines)
-  OTMA Input/Output Edit exit routine (DFSYIOE0) (Exit Routines)
-  OTMA Destination Resolution exit routine (DFSYPX0) (Exit Routines)
-  OTMA Resume TPIPE Security exit routine (DFSYRTUX) (Exit Routines)

Using DFSYPX0 and DFSYDRU0 OTMA exit routines to determine destination

Transaction pipe names can be the same as IMS LTERM names or APPC/IMS TP names.

To clarify whether a destination is for OTMA, IMS provides OTMA exit routines that can specify where IMS should look to resolve the destination names:

- The OTMA Destination Resolution exit routine (DFSYPX0)
- The OTMA User Data Formatting exit routine (DFSYDRU0)

In an IMS subsystem, you can have many DFSYDRU0 exit routines, but only a single DFSYPX0 exit routine.

The exit routines cannot change the actual destination name.

Determining the destination for an OTMA message requires two phases. In each phase, an OTMA exit routine can be called:

Phase 1

The OTMAYPRX user exit is called to determine the initial destination for the output.

The exit routine can determine whether the message should be directed to an OTMA client or to IMS TM for processing. The exit routine cannot determine the final destination (because insufficient parameters are passed to it).

Phase 2

The DFSYDRU0 exit routine is called to determine the final destination for the output. Each client can specify a separate DFSYDRU0 exit routine.

The name of the DFSYDRU0 exit is determined by the user or an OTMA client. Each client can have its own dedicated DFSYDRU0 exit. To view the name of the DFSYDRU0 exit routine associated with an OTMA client, issue the /DISPLAY TMBER command.

Both of these exit routines receive control when an IMS application program issues an ISRT call to an alternate program communication block (PCB), or issues CHNG or PURG calls. But if the destination is the name of an IMS scheduler message block (SMB), the DFSYDRU0 exit routine does not receive control. The following figure illustrates the two phases of message destination determination.




Figure 134. How DFSYPRX0 and DFSYDRU0 determine message destination

Recommendations:

- Because of the potential conflict with the SMB name, OTMA clients should avoid using a transaction pipe name as either the transaction name or the routing key.
- The DFSYPRX0 and DFSYDRU0 exits should be the same for the front-end and back-end IMS systems within a shared queues group. If the exit routines are different for one or more back-end IMS systems, asynchronous output might be sent to different destinations, depending on which back-end IMS system processed the input.

To ensure prompt delivery of the output, enable OTMA on every back-end IMS system in the shared queues group. If a back-end IMS system does not have

OTMA enabled, any asynchronous OTMA output that is inserted into an alternate PCB is simply queued and not delivered until the operator issues a /STA OTMA command.

- Specifying OTMAMD=Y in the IMS PROCLIB member DFSPBxxx can direct your OTMA message from the DFSYPRX0 exit to a different DFSYDRU0 exit without rerouting.
- Specifying OTMASP=Y in the IMS PROCLIB member DFSPBxxx always creates a SYNC tpipe for the ALT-PCB output message.

Note:

1. The SCD and PST addresses are available in the input parameter for both OTMA user exits.
2. The address of the first segment of the output message is not passed to either user exit.

Related reference:

 OTMA User Data Formatting exit routine (DFSYDRU0) (Exit Routines)

 OTMA Destination Resolution exit routine (DFSYPRX0) (Exit Routines)

Administering IMS for OTMA

The following topics describe IMS administration considerations for OTMA.

IMS conversations and OTMA

OTMA-to-IMS conversations are send-then-commit and are nonrecoverable.

IMS creates a unique conversation ID for each transaction pipe that is saved in the Server token field in the OTMA state-data section of the message. This ID must be passed to IMS with all subsequent client input for the conversation using that transaction pipe.

The server-state flag is set to Conversational State in the state-data section of the message prefix when IMS acknowledges the transaction. This flag must be set for all subsequent client input during the conversation.

An IMS conversation can be terminated by issuing the /EXIT command. Alternatively, if a client wants to terminate the IMS conversation, the client can send a message with the commit-confirmation flag set in the message-control information section of the message prefix (indicating a deallocate flow). IMS then terminates the IMS conversation. A deallocate flow must specify Conversational State for the server-state flag in the state-data section of the message prefix; it must also set the server token field.

MSC and OTMA transactions

You can use IMS Multiple Systems Coupling (MSC) with OTMA transactions.

MSC processing in an OTMA environment is similar to MSC processing in an APPC/IMS environment:

- A client sends a transaction to IMS. If the transaction is defined as remote for that IMS system, it is sent to a remote IMS system for processing. If the transaction is defined as a local transaction and performs a message switch to another IMS system, the switched message is sent to that remote IMS system for processing.

- Output from the remote application program is returned to the originating IMS.
- IMS recognizes that the data is OTMA data and uses the transaction pipe to send the data to the client.

If the remote application inserts to an alternate PCB that is a remote destination, the data is not routed to an OTMA destination. The remote destination does not route the output message to the OTMA client, even though the message has a prefix. If the message is to be properly routed back to the original client, the remote IMS must insert to a remote transaction. That transaction (at the original IMS site) must then send the message to the OTMA client using an alternate PCB and a Prerouting exit routine.

You can use MSC with OTMA in a shared-queues environment, as long as the MSC link exists in the front-end IMS subsystem that is connected to the OTMA client.

Fast Path and OTMA transactions

Fast Path transactions must run as send-then-commit transactions.

Any parameters with the OTMA transaction that contradict this commit mode cause the transaction to be rejected. Existing Fast Path application programs can run with OTMA if the client-entered transaction is properly defined.

IMS restart processing and OTMA

If an IMS subsystem connected to an z/OS cross-system coupling facility group must be restarted, IMS reconnects to the group during restart, but all clients must send new client-bid requests to IMS.

XRF processing and OTMA

XRF supports OTMA.

Recommendation: The active and alternate IMS subsystems should be in the same z/OS cross-system coupling facility (XCF) group. If they are not in the same XCF group, after an XRF takeover, the client must connect to the new XCF group. If they are in the same XCF group, after an XRF takeover, the client is automatically reconnected to the XCF group.

The active IMS joins the XCF group specifying the IMS USERVAR name as its member name. The alternate IMS joins the group only after an XRF takeover, and connects using the same USERVAR name as the member name. If the active and XRF alternate subsystems are in the same XCF group, during XRF takeover, the new active IMS:

- Ensures that the old IMS member is disconnected
- Rejoins the group
- Begins delivering any queued output destined for an OTMA client, following client-bid processing for that client

IMS XRF does not track the /START OTMA and /STOP OTMA commands. The IMS procedures for both the active and alternate IMS should both specify OTMA=Y in order to ensure that IMS automatically rejoins the XCF group during XRF takeover.

In an XRF environment, OTMA clients are equivalent to Class 3 terminals and are not automatically reconnected to IMS. The clients detect that IMS has left the XCF group and wait for IMS (or the XRF alternate) to rejoin the group. Then the clients send new client-bid requests to IMS.

The alternate system in an XRF complex does not track when resynchronization has begun or the resynchronization of a particular tpipe. If IMS fails during resynchronization, the client should detect the IMS failure, and try to resynchronize with the new XRF active system.

IMS Queue Control Facility and OTMA

The IMS Queue Control Facility (QCF) supports OTMA messages. You can use QCF to switch between all supported IMS releases, or between Shared Queues and non-Shared Queues.

You can access QCF both online with an ISPF interface and with control statements in a BMP environment.

TMEMBER and TPIPE are the operands for the INCLUDE and EXCLUDE control statements.

tmember

A 1- to 16-character OTMA target member (client) name. You can generically specify groups of names that begin with the same characters by using an asterisk (*) after the characters that are the same. An asterisk as the first character will include or exclude **all** OTMA transactions.

tpipe A 1- to 8-character OTMA transaction pipe name. You can generically specify groups of names that begin with the same characters by using an asterisk (*) after the characters that are the same.

To selectively recover OTMA messages, use the INCLUDE and EXCLUDE control statements. The format of the INCLUDE and EXCLUDE statements with OTMA operands is:

```
INCLUDE operand(,)
```

```
EXCLUDE operand(,)
```

operand must start in column 10 and is one of the following:

- TMEMBER=tmember
- TPIPE=tpipe

Example: To select all OTMA messages using transaction pipe name S4A1BV6, specify:

```
INCLUDE  TMEMBER=*,
          TPIPE=S4A1BV6
```

All messages with the same tmember are grouped together, and the count is reported by the tpipe name.

Example:

```
**** MESSAGES INSERTED BY DESTINATION ****
      BY OTMA DESTINATION
      TMEMBERNAME
      TPIPE1                      count
      TPIPE2                      count
```

If a client-bid request changes the name of the current OTMA Destination Resolution exit routine, any transactions enqueued before IMS terminates that are then reprocessed by the Message Requeuer might not use the changed exit routine name. Inserts to alternate PCBs use the exit routine name in the client descriptor.

With QCF, you can identify a category of message as well as the message type. The following table describes category parameters, and the supported message types and keywords associated with the parameter.

Table 133. Selecting messages by category type

Category parameter	Description	Supported message types and keywords
DESTTYPE	Checks the destination of a message for a selected message type	APPC, LTERM, MSC, OTMA, LTRAN, RTRAN, TRANS, VSP
SRCETYPE	Checks the source of the message for a selected message type	APPC, MSC, OTMA, VSP
MSGTYPE	Checks the source or destination of the message for the selected message type	APPC, LTERM, MSC, OTMA, VSP

Related Reading: For more information about message selection by category in QCF, refer to *IMS Queue Control Facility for z/OS User's Guide*.

Using shared queues with OTMA

This topic describes general information about using IMS shared queues with OTMA.

To ensure delivery of alternate PCB processing, enable OTMA on all IMS systems; assign each IMS system in the shared queue group a unique z/OS cross-system coupling facility member name.

Use the /DISPLAY TRANS ALL QCNT to view all the OTMA transactions currently in the shared queue group waiting to be processed. Use /DISPLAY QCNT OTMA MSGAGE 0 to view all the OTMA outbound messages on the shared queues.

Note: If OTMA outbound messages remain on a shared queue when the IMS system that was processing them is shutdown normally and then cold started, the outbound messages become stranded on the shared queue and cannot be delivered. However, if a normal shutdown is followed by a warm start, any outbound messages on the shared queue at the time of the shutdown can be delivered.

For OTMA hold-queue capable clients, such as IMS Connect, using the OTMA super member function can ensure that outbound messages do not become stranded on a shared queue and can be delivered to the client.

As the result of a temporary shortage in the HIOP storage pool, you might receive message DFS1269E, which notifies you of an internal IMS failure to register a shared queue resource. To re-register the shared queue resource for OTMA, issue the IMS commands /STOP OTMA and /START OTMA.

In a shared queues environment, IMS removes idle, non-queued, non-synchronized transaction pipes after three system checkpoints.

Related concepts:

“Using DFSYPRX0 and DFSYDRU0 OTMA exit routines to determine destination” on page 776

Related reference:

➡ DFS1269E message information (Diagnosis)

Related information:

➡ DFS1269E (Messages and Codes)

OTMA commit-then-send messages

OTMA commit-then-send (commit mode 0 or CM0) messages can be processed on any IMS system in the shared-queues group. Program-to-program switches can also be run on any IMS.

OTMA unsolicited messages

OTMA clients must connect to every IMS system in the shared queue group in order to receive unsolicited messages. The OTMA client connections are necessary because transactions that might cause unsolicited messages can run on any IMS within the shared-queues group.

If the IMS that processes an unsolicited message (the back-end system) is a different IMS than the one that receives the message, the unsolicited message is delivered by the back-end system. Therefore, OTMA must also be enabled on the back-end IMS.

OTMA send-then-commit messages

OTMA send-then-commit messages can also be processed on any IMS system in the shared queue group.

Synchronous and asynchronous transactions created by a program-to-program switch from an input synchronous transaction always run on the same IMS system as the transaction that initiated the program-to-program switch.




Figure 135. OTMA messages being processed on multiple IMS systems in a shared-queues group

In addition, program-to-program switching is not allowed for protected conversations (sync level 2).

Synchronous transactions which use send-then-commit processing support the following commit levels:

- NONE
- CONFIRM
- SYNCPT

Asynchronous transactions which use commit-then-send processing support the following commit levels:

- RESYNC
- NO RESYNC

The commit levels for synchronous and asynchronous transactions are shown in the following figure.




Figure 136. Synchronous and asynchronous transactions and their respective commit levels

Use the `/DISPLAY ACTIVE` command to determine whether the shared queue function for OTMA send-then-commit is active.

Using other IMS commands

The IMS command `/DISPLAY TMEBER membername TPIPE tpipename QCNT` shows the tpipe status and the output message queue count in a shared queue for a particular IMS system.

Related concepts:

“OTMA program-to-program switch processing” on page 807

Retrieving ALTPCB output from a back-end IMS system

In a shared queues environment, unless you take additional steps, the alternate PCB (ALTPCB) output placed on the shared queue by a back-end IMS system has affinity to the back-end system. This affinity prevents the originating front-end IMS system from retrieving the output unless you take additional steps.

To enable a front-end IMS system to retrieve ALTPCB output that is placed on the shared queue by a back-end IMS system, you can use the following solutions:

1. Enable the super member function and code the OTMA client to issue a resume tpipe call at the front-end IMS system. All retrieval options of the resume tpipe call are supported in the shared queues environment.
2. Disable the back-end affinity. This option is useful for OTMA clients that do not use the super member function, such as WebSphere MQ.

You can disable the back-end affinity by either of the following methods:

- In the OTMA client descriptor in the DFSYDTx PROCLIB member of the front-end IMS system, specify ALTPCBE=YES.
- In third byte of the state data prefix of the input transaction message, set the TMAMALTB flag (X'01').

When either of these two options are specified, the ALTPCB output generated at the back-end IMS system is delivered to the front-end IMS via the shared queue.

Related tasks:

“Sharing asynchronous commit-then-send output: the OTMA super member function” on page 843

IMS termination and OTMA

When IMS terminates, OTMA attempts to notify OTMA clients and might take additional actions depending on the circumstances.

Related tasks:

 Shutting down IMS (Operations and Automation)

OTMA restrictions

A number of general restrictions and requirements apply to OTMA.

The maximum total length of all prefixes for an OTMA message is 4096 bytes. This length does not include any application data.

Existing IMS application programs that use SET0 calls might not run as expected. APPC/IMS application programs using SET0 calls might require modification to use implicit OTMA support.

IMS conversational and Fast Path transactions must be defined as send-then-commit. Existing Fast Path applications can run with OTMA.

A transaction from an IMS terminal (for example, a SLU 2 terminal) cannot route output directly to a client, but must use an OTMA Prerouting exit routine (DFSYPX0).

OTMA does not support the IMS Message Format Service (MFS). However, the MFS message output descriptor (MOD) name can be specified by the client in the prefix of an OTMA message.

OTMA does not support IMS Front-End Switch.

OTMA messages cannot be encrypted.

All user IDs must be verified by RACF , unless the client specifies no security checking in the security-data section of the message prefix.

IMS modules that contain z/OS cross-system coupling facility (XCF) macros must be reassembled for new releases of IMS.

OTMA has read only access to the main storage data base (MSDB). No update access is available to MSDB from OTMA.

OTMA does not operate in the IMS DBCTL environment.

OTMA does not allow IMS terminal control commands, such as /FORMAT, /HOLD, /RCL, and /SIGN commands.

OTMA protected conversation messages that are part of the same z/OS Resource Recovery Services (RRS) unit of recovery and that access common database resources must be scheduled and executed on the same IMS system.

IMS shares locks and database buffers between protected conversation transactions, but this support is restricted to within a single IMS system. It does not function across shared queues. You must ensure such messages schedule on the same IMS system.

Managing system resources and OTMA

OTMA uses IMS system resources. A variety of factors affect how well OTMA uses those resources and there are a number of ways to monitor that usage.

The following topics describe these system resource considerations.

OTMA resource monitoring

OTMA monitors IMS system resources that are used to process OTMA transactions and notifies OTMA clients about how well the IMS system is processing OTMA transactions.

If an OTMA client receives a notification that the IMS system is not processing OTMA message normally, the OTMA client can then take appropriate action, such as rerouting OTMA transaction messages to a different IMS system.

OTMA sends notifications to OTMA clients as a *server state protocol command*. OTMA issues the server state protocol command in the following circumstances:

- When an OTMA client establishes a tpqe connection
- When a significant change occurs in the ability of IMS to process OTMA messages
- As a *heartbeat message* at 60 second intervals

The server state protocol message categorizes the overall state of the IMS system by using the following states.

Normal state (X'03')

IMS is available and is processing OTMA messages normally.

Degraded state (X'02)

IMS is processing OTMA messages slowly. OTMA issues a degraded state protocol command when one or more conditions indicate that IMS is not processing OTMA messages as quickly as it should.

Unavailable state (X'01)

IMS can no longer accept OTMA transactions for processing. OTMA issues the unavailable state protocol command to alert the OTMA client that one or more severe conditions prevent IMS from processing OTMA messages.

In addition to notifying the client of the overall state of IMS processing, if the IMS processing is in either a degraded or unavailable state, the server state protocol command can include additional information about the condition of specific resources associated with the degraded or unavailable state.

The server state protocol command is identified by a value of X'3C' in the protocol command type field (TMAMCTYP) in the message control information section of the OTMA message header. The server state protocol command itself is delivered in the state data section of the OTMA header and mapped with the TMAMRSIM field prefix in the DFSYMSG macro.

You can view information about OTMA clients, OTMA tpipe connections, and the OTMA messages that are currently being processed by IMS by issuing the type-2 IMS Command QUERY OTMATI.

Related concepts:

“Displaying the current transaction workload”

Related reference:

“Server state protocol command” on page 876

“Message-control information section” on page 855

Displaying the current transaction workload

You can view information about the number and type of OTMA messages that are currently being processed by IMS by issuing the type-2 IMS Command QUERY OTMATI.

Issuing the QUERY OTMATI command without specifying any parameters shows the current number of OTMA messages on each IMS system in an IMSplex listed with the OTMA client and the tpipe used to submit the messages. For example, in an IMSplex with three IMS systems the following information could be returned by QUERY OTMATI:

TMember	TPipe	MbrName	CC	MsgCnt
MQ	APPLA	IMSA	0	102
MQ	APPLB	IMSB	0	201
WAS	APPLC	IMSB	0	301

By specifying the parameters of QUERY OTMATI, you can filter and refine the information returned by the command.

For example, if you need to know if messages are waiting too long to be processed by IMS, the MSGAGE(*nnn*) subparameter returns a count of all messages in an IMS system that are older than *nnn*, where *nn* is a clock time specified in seconds. In the following example, QUERY OTMATI MSGAGE(3) was specified and the message count for each IMS system includes only those messages that are older than three seconds:

TMember	TPipe	MbrName	CC	MsgCnt	MsgAge
MQ	APPLA	IMSA	0	50	3
MQ	APPLB	IMSB	0	80	5
WAS	APPLC	IMSB	0	100	4

In the example above, the number shown in the MsgAge column is the age of the most recently submitted message among all messages that are older than three seconds. For example, among the 80 messages submitted by MQ on tpipe APPLB to IMSB that are older than three seconds, the message that has been processing for the least amount of time has been processing for five seconds.

The QUERY OTMATI command includes over 25 parameters and subparameters that return information about the various components and resources that are associated with the processing of messages from OTMA clients. In addition to message counts and ages, you can also retrieve information about such things as:

- The transaction or application program scheduled to process the message
- Security information
- Context tokens for two-phase commit processing
- OTMA commit mode
- Synchronization level
- Correlator tokens for callout processing
- Timeout setting for acknowledgments from the client

Impact of OTMA message TIBs on storage

For each message received from an OTMA client, IMS creates a *transaction instance block* (TIB).

Each TIB requires approximately 8 KB of extended private area storage.

You can view information about the TIBs and the OTMA messages that are currently being processed by IMS by issuing the type-2 IMS Command QUERY OTMATI. The information you can view includes message counts, message ages, and much more.

Normally TIBs are freed and the storage released by IMS either after IMS enqueues the commit mode 0 (CM0) input message of TIBs created for CM0 messages or after IMS returns a CM1 output message of TIBs created for CM1 messages.

There are circumstances in which IMS cannot free a TIB. When this happens, the TIB persists in storage as an *orphan TIB*.

An excessive number of TIBs is usually caused by either a message flood condition or orphan TIBs.

Message flood conditions usually occur as a result of:

- A stopped program
- The unavailability of message processing regions
- The rate of incoming messages exceeding the rate at which IMS can process the messages

To help avoid a message flood condition, you can enable message flood detection by specifying a maximum number of OTMA messages that IMS can process at any one time. IMS determines the number of messages that are currently being

processed by counting TIBs. When message detection is enabled, IMS issues messages as the number of TIBs nears the maximum.

Orphan TIBs are created when IMS processes an OTMA transaction, but does not generate either a CM1 output message or a DFS2082 message in place of the CM1 output message. For example, if a CM1 message triggers a program-to-program switch to an asynchronous transaction that is defined as a non-response mode transaction, and if OTMAASY=Y, IMS cannot delete the TIB even though the switched-to transaction generates a response to the IOPCB.

IMS includes orphan TIBs in its count of total TIBs when message flood is enabled. If OTMA input has been stopped due to a message flood condition, you can issue the /DISPLAY OTMA or /DISPLAY TMEMBER command to see the number of TIBs. If there is still a high number of TIBs after the OTMA input has been stopped, it is possible that there are orphan TIBs.

You can remove orphan TIBs related to IMS conversational transactions. You can identify orphan TIBs that are related to conversational transactions by issuing the command /DISPLAY CONVERSATION. If an orphaned TIB exists for the conversation, you can then issue the command /EXIT CONVERSATION to remove the orphaned TIB.

Related concepts:

“Displaying the current transaction workload” on page 787

“Message flood detection”

Message flood detection

If IMS cannot process the messages received from an OTMA client quickly enough or at all, the build up of messages in the IMS system can result in a *message flood* condition. When a message flood condition occurs, the storage required by both the messages and the transaction instance blocks (TIBs) that IMS uses to process the messages, can exhaust the below-the-line storage and potentially cause a z/OS S40D system abend.

To avoid the problems associated with a message flood condition, OTMA can monitor the number of input messages that are waiting to be processed in the IMS system. When monitoring is active, OTMA monitors both the number of input messages from each OTMA client and the total number of input messages from all OTMA clients combined. When the number of input messages from an individual OTMA client reaches its defined maximum, OTMA suppresses new input messages from that client. When the total number of input messages reaches the global maximum, OTMA issues a warning message, but still accepts new input messages.

You can modify the maximum number of input messages for both individual OTMA clients and for all OTMA clients combined. You can specify from 200–9999 as the maximum number of allowable messages or you can disable message flood detection by specifying 0.

The default maximum number for OTMA clients is 5000. The default maximum number for all OTMA clients combined is 8,000.

When the number of messages in the IMS system reaches 80% of the maximum number defined for either an OTMA client or for all OTMA clients, OTMA issues a server state protocol command to the appropriate OTMA clients and IMS issues message DFS1988W to the master console. IMS also issues DFS1988W at every 5% increase thereafter until the maximum number is reached.

When the number of messages in the IMS system for an individual client exceeds the maximum, OTMA issues another server state protocol command and IMS issues message DFS1989E. OTMA then rejects all subsequent input messages from the OTMA client until the message flood condition is resolved.

When the number of messages in the IMS system falls to 50 percent of the maximum, OTMA issues a new server state protocol command to the OTMA clients notifying them that the flood condition has been relieved.

To enable message flood detection for an individual OTMA client, specify the maximum allowable number of input messages from the OTMA client in one or more of the following places:

- The INPT= parameter of the OTMA client descriptor for the OTMA client
- The INPUT parameter of the /START TMEMBER command
- The client-bid request from the OTMA client in the two-byte field at byte 62 of the client-bid command message

To enable message flood detection globally for OTMA clients combined, use the type-1 command /START TMEMBER ALL INPUT *nmb*r.

Values specified in the OTMA client descriptor can be overridden by the /START TMEMBER command. If the client-bid request specifies a smaller number, the client-bid request can also override the value specified in the OTMA client descriptor.

Related concepts:

“Impact of OTMA message TIBs on storage” on page 788

Administering OTMA tmembers

IMS provides commands to get information about and dynamically modify OTMA target members (tmembers).

Use the START TMEMBER and STOP TMEMBER commands to start and shut down specific OTMA tmembers and tpipes. The commands also provide parameters to change the size of the input message hold queue and timeout length for enqueued CM1 and CM0 messages.

The /STOP TMEMBER command also clears the wait state (WAIT_A, WAIT_H, or WAIT_S) of all CM0 messages on the tpipe that are waiting for either an ACK or a NAK message from the client.

You can retrieve detailed information about some or all tmembers with the DISPLAY TMEMBER command. You can specify DISPLAY TMEMBER ALL to retrieve a report for all of the configured OTMA tmembers and their associated tpipes, or qualify the command with specific tmember and tpipe names to narrow the search. The response output shows the processing mode for each tpipe, security level, status, workload statistics, and other information.

The enqueue and dequeue counts on a tpipe are updated only for CM0 output messages. The counts are not updated for CM1 messages regardless of the sync level because CM1 messages are not recoverable.

You can also retrieve a list of OTMA tmembers with stopped tpipes by using the DISPLAY STATUS TMEMBER command. This command is useful for determining

whether any tmembers require administrator attention. The command also retrieves the name of the super-member for each tmember, if one is configured.

You might need to manually dequeue stalled or stale messages from an OTMA tpipe. For example, you might want to purge stale messages from a tpipe after modifying a client application so that new messages can be retrieved instead. You can use the DEQUEUE command to purge messages from a specific tpipe. The command provides options to dequeue either the first enqueued message for the tpipe or all enqueued messages.

If you need additional information about a tmember or a specific tpipe, you can enable the OTMA client activity trace with the TRACE TMEMBER command. You can also turn on tracing for an entire super member. Use the DISPLAY TRACE TMEMBER command to determine which tmembers and tpipes are currently being traced for an OTMA client. The tracing information is available in the OTMA log records.

A temporary tpipe is created when you issue a /TRACE *tpipe* or /STOP *tpipe* command against a tpipe that does not exist. A temporary tpipe is converted to a permanent tpipe if an input message reaches IMS through the tpipe or if an output message is queued to the output queue of the tpipe.

When you issue a /DISPLAY *tpipe* command against a temporary tpipe, the status of TMP is displayed.

Related reference:

 /DISPLAY TMEMBER command (Commands)

 /DISPLAY STATUS command (Commands)

 /DEQUEUE command (Commands)

 /START TMEM command (Commands)

 /STOP TMEM command (Commands)

Tpipe resource impact

Tpipes use significant amounts of IMS resources and processing time, so try to limit the number of tpipes that are created for each tmember.

IMS removes transaction pipes after they have been idle for three consecutive system checkpoints, except in the following circumstances:

- Commit-then-send messages are queued on the tpipe or the tpipe hold queue
- The tpipe is stopped A trace is set on the tpipe
- The tpipe is a synchronized tpipe, such as a tpipe used by MQSeries for commit-then-send input transactions
- The tpipe is in a WAIT state for a resume tpipe request that specified either the AUTO or the SINGLE-WAIT options

One way you can control the number of tpipes that are created for a particular OTMA client, is to set a maximum allowable number of tpipes for each OTMA client and for the IMS system.

Tpipe limits for individual OTMA clients

A maximum number of tpipes is specified for an OTMA client by specifying the MAXTP parameter on the OTMA client descriptor in the DFSYDTx member of the IMS.PROCLIB data set.

When the MAXTP parameter is specified for an OTMA client, OTMA monitors the number of tpipes created for the client and issues warnings when the number reaches certain levels.

After a MAXTP value is set for an OTMA client, OTMA monitors the number of tpipes that are created for the client. If the total number of tpipes reaches 80% of the maximum allowable number, IMS issues warning message DFS4382W to the system console and the master terminal operator (MTO). If the number of tpipes reaches 100% of the maximum number, IMS issues error message DFS4383E to the system console and the MTO. Any input transactions that require a new tpipe are rejected with a NAK with OTMA sense code X'29'.

After the warning or error messages is issued and the total number of the tpipes for the client drops 50% of the maximum allowable number, IMS issues message DFS4384I to indicate that the number of tpipes returned to normal.

Tpipe limits in the IMS system

You can define a global TPIPE warning threshold for all of the OTMA clients by defining an OTMA client descriptor with the name DFSOTMA and specifying the MAXTP parameter. The DFSOTMA descriptor is the OTMA system client descriptor that defines values for the IMS system that apply to all OTMA clients.

When the number of tpipes in the IMS system reaches 80% of the DFSOTMA MAXTP value, IMS issues message DFS4515W to both the system console and MTO and prevents all OTMA clients from creating new tpipes. If an input transactions requires a new tpipe, the transaction is rejected with a NAK message that contains OTMA sense code X'29' and message DFS4516E is sent to the system console and the MTO.

After a DFS4515W or DFS4516E is issued and the total number of monitored tpipes in the system drops down to 50% or another user-specified level of the global tpipe warning threshold, IMS issues message DFS4517I to indicate that the number of tpipes returned to normal.

If the OTMA DFSOTMA system client descriptor does not define a global maximum for the number of tpipes in an IMS system, and one or more OTMA clients have a maximum allowable number of tpipes defined in their OTMA client descriptors, the highest MAXTP value among all of the client descriptors serves as a global warning threshold for the IMS system. When the total number of tpipes in use by all OTMA clients that are subject to tpipe monitoring reaches the global warning threshold, IMS issues warning message DFS4385W to the system console and the MTO.

After a global warning is issued and the total number of tpipes drops down to 80% of the global warning threshold, IMS issues message DFS4386I to indicate that the number of tpipes returned to normal.

X'3C' protocol command notifications

OTMA also sends out the X'3C' protocol command to the OTMA clients at the various warning, error, and relief thresholds. Upon receiving the X'3C' protocol command for a warning or error, the client applications can reroute any subsequent transactions to a different IMS system as appropriate. When the number of tpipes drops below the relief threshold, the X'3C' protocol command is issued again with the warning or error flag turned off.

Displaying information about the number of tpipes

The IMS commands /DISPLAY OTMA and /DISPLAY TMEMBER can show the current number of tpipes for the OTMA clients that have a MAXTP value set. If an OTMA client reaches the maximum allowable number of tpipes, the command output shows MAX TPIPE as the USER_STATUS for the OTMA client.

The global warning threshold set by the highest MAXTP value among multiple OTMA clients is displayed under the output field TPNCT for the IMS server.

Automatic removal of idle tpipes

IMS scans tpipes during system checkpoint processing to determine if any tpipes can be deleted. If a tpipe meets certain qualifications, IMS attempts to remove the tpipe if it has been idle across three consecutive checkpoints.

IMS attempts to remove idle tpipes that conform to the following characteristics:

- The tpipe has not been stopped by a command.
- The tpipe is not the subject of a tpipe trace.
- The tpipe does not have any incomplete send-then-commit (commit mode 1 or CM1) messages.
- The tpipe is not used as a synchronous tpipe by WebSphere MQ.
- The tpipe does not have queued commit-then-send (commit mode 0 or CM0) output messages.

There are some internal circumstances in which IMS cannot remove an idle tpipe. In the event of such circumstances, IMS attempts to remove the idle tpipe again at the following system checkpoint.

You can view information about OTMA clients, OTMA tpipe connections, and the OTMA messages that are currently being processed by IMS by issuing the type-2 IMS Command QUERY OTMATI.

IMS message queue data set size and OTMA

Messages entering IMS from OTMA contain both the OTMA message prefix and other existing IMS message prefixes.

The OTMA message prefix is variable in length. Excluding the user data section, the OTMA message prefix can become very large, sometimes over 200 bytes in length. The OTMA message prefix, including the user data section, is stored on IMS message queue data sets, which increases usage of the queue buffer pool.

Recommendation: Because of this increase in queue buffer pool usage, try to increase the size of the message queue data sets.

Buffer pool usage for OTMA

If an IMS-OTMA environment has heavy OTMA traffic, a significant increase in LUMP and HIOP pool usage can occur.

Because LUMP and HIOP pools are allocated from private storage, you might need to increase the size of the IMS control region. Also, certain OTMA control blocks are allocated from extended common service area (ECSA), another limited resource.

Recommendation: Increase the ECSA size according to your workload. For example, if a client is sending more than 20 messages over 100 tpipes within a few seconds, try increasing the IMS control region size to 200 MB or more, and increase the ECSA size to 50 MB or more. If you cannot increase the IMS control region size or the ECSA size, try balancing your workload to allow IMS to reuse its buffers more effectively.

Dependent region occupancy and OTMA

A send-then-commit transaction remains in a dependent region while the output is being sent (before a sync point occurs). You can get additional information about a specific dependent region that is in an OTMA-related wait state by using the DISPLAY ACT command.

The following states (from the STATUS response field) indicate that an OTMA-related activity is in progress:

- WAIT-CALLOUT
- WAIT-RRS PC
- WAIT-SYNCPOINT
- TERM-WAIT SYNCPT
- TERM-WAIT RRS

If the region is waiting for a synchronous program switch response, the amount of time remaining before the region times out the request is displayed in the END TIME field of the output from the DISPLAY ACT REGION command.

Recommendations:

- If many of your transactions are send-then-commit (CM1) transactions, increase the number of dependent regions to improve throughput performance.
- For the CM1 messages that use synclevel=confirm or synclevel=syncpt, specify a timeout interval to free dependent regions if an OTMA client does not provide the required ACK or NAK response within a reasonable amount of time.
- Use as many commit-then-send (CM0) OTMA transactions as possible.

Related tasks:

“Timeout interval for acknowledgments to OTMA messages”

Related reference:

 /DISPLAY ACT command (Commands)

Timeout interval for acknowledgments to OTMA messages

You can specify an ACK timeout interval that determines how long OTMA waits for an ACK or NAK acknowledgment for OTMA output messages.

ACK timeout intervals can be specified for the following types of OTMA output messages:

- Some send-then-commit (CM1) response messages
- Commit-then-send (CM0) response messages

For a send-then-commit transaction that uses either `synclevel=confirm` or `synclevel=syncpt`, if OTMA does not receive the expected ACK or NAK response from the OTMA client before the timeout interval expires, OTMA aborts the transaction and IMS backs out the transaction, issues an OTMA CM1 deallocation message to the OTMA client, and issues message DFS0809E to the z/OS system console.

For a commit-then-send transaction, the action OTMA takes if an expected ACK or NAK response is not received before the timeout interval expires depends on whether the client is a hold-queue capable client, such as IMS Connect, or a non-hold-queue capable client, such as WebSphere MQ. For hold-queue capable clients, the action OTMA takes also depends on whether the output is the I/O PCB queue or the hold queue of a tpipe.

For a hold-queue capable client, if the output for a commit-then-send transaction is on the I/O PCB queue and the timeout interval expires, OTMA takes the following action:

1. OTMA attempts to deliver the CM0 output to a reroute tpipe.
2. If no reroute tpipe has been specified, OTMA attempts to deliver the output to a timeout queue.
3. If no timeout queue has been specified, OTMA delivers the CM0 output to the default timeout queue DFS\$\$TOQ.
4. Issues DFS3494E to the z/OS system console.

For a hold-queue capable client, if the output for a commit-then-send transaction is on the tpipe hold queue and the timeout interval expires, OTMA takes the following action:

1. OTMA attempts to deliver the output to a timeout queue.
2. If no timeout queue has been specified, OTMA delivers the CM0 output to the default timeout queue DFS\$\$TOQ.
3. Issues DFS3494E to the z/OS system console.

For non-hold-queue capable clients, OTMA takes the same action as for hold-queue capable clients when the CM0 output is on the I/O PCB queue.

For a synchronous callout requests from IMS application programs that issue the DL/I ICAL call, OTMA processes the request as a CM0 output message. OTMA sets the CM0 flag in the state data prefix and waits for an acknowledgment of receipt of the synchronous callout message. If an acknowledgment is not returned within the ACK timeout interval, OTMA takes the following action:

1. Issues return code X'100' and reason code of X'104' to the IMS application program.
2. Discards the synchronous callout request message. The message is not rerouted.
3. Issues DFS3494E to the z/OS system console.

You can specify a timeout interval for an OTMA client in the OTMA client descriptor. You can override the timeout interval in an OTMA client descriptor by specifying a different interval in the TIMEOUT parameter of the /START

TMEMBER command. You can also override the timeout interval in the OTMA client descriptor by specifying a smaller timeout interval in the OTMA client's client bid request.

For individual transactions, you can only specify a timeout interval that is shorter than the timeout interval specified for the OTMA client.

If you do not specify a send-then-commit timeout value, OTMA uses a default value of 120 seconds.

To view the current timeout interval set for an OTMA client, issue the /DISPLAY TMEMBER command.

Specifying an acknowledgment timeout interval at the OTMA client level

At the OTMA client level, you can specify an ACK timeout interval in one or more places.

You can specify a timeout interval in the following places:

- The T/O= parameter of the OTMA client descriptor
- The TIMEOUT parameter of the /START TMEMBER command
- The 1-byte field at offset 65 (X'41') of the state data in the client-bid request from the OTMA client

For commit-then-send transactions, you can specify the name of a timeout queue by:

- Specifying X'04' at byte offset 45 (X'2D') in the client-bid request
- Specifying the 8-byte character name of the timeout queue at byte offset 66 (X'42').

If no timeout queue name is specified, CM0 output that times out is routed to the default timeout queue DFS\$\$TOQ.

Related concepts:

"OTMA client descriptors" on page 771

Related reference:

"Server-Available and Client-Bid commands" on page 867

 /START TMEM command (Commands)

Specifying an acknowledgement timeout on CM1 transaction messages

In the OTMA header of an input message for a send-then-commit (CM1) transaction that uses either synclevel=confirm or synclevel=syncpt, you can specify a timeout interval to limit how long IMS waits for an acknowledgement (ACK or NAK) to the output response.

To set the time out interval:

- Specify the interval in minutes using decimal integers in the 1-byte field at offset X'1E' of the message control information prefix of the transaction message.
- Set flag X'08' at byte 5 of the state data

OTMA support for transaction expiration

You can specify an expiration time for a transaction to reduce processing costs by preventing IMS from processing transactions that the client can no longer use.

When a transaction specifies an expiration time, OTMA monitors the transaction and, if the transaction is not processed or enqueued before the time expires, OTMA discards the transaction.

OTMA stops monitoring a transaction for expiration when an IMS application program in an IMS dependent region retrieves the transaction for either MSC, Fast Path, or conversational processing. After either of these events, the expiration time no longer applies and IMS processes the transaction and returns the output to OTMA and the OTMA client.

You can enable transaction expiration either when you define the transaction to IMS or in the OTMA message header when the transaction is submitted to OTMA. An expiration time specified in the definition of a transaction becomes the default expiration period for that transaction type and applies to all instances of the transaction. An expiration time specified in an OTMA message header applies only to the transaction instance submitted with the OTMA header and overrides any expiration time specified in the transaction definition.

An expiration time specified in an OTMA message header can be specified as a point in time or as a length of time. If specified as a point in time, the expiration time is specified in store clock (STCK) format and represents the time of day at which a transaction expires. If specified as a length of time, the expiration time is specified in seconds and OTMA calculates the expiration time from the time at which OTMA receives the transaction from the z/OS cross-system coupling facility (XCF).

OTMA checks if a transaction is expired at three points:

1. When OTMA first receives a transaction from XCF. If the expiration time has already passed, OTMA discards the transaction and returns a NAK message to the client.
2. When OTMA enqueues a transaction in IMS. If a transaction expires before OTMA enqueues the transaction, OTMA discards the transaction and returns a NAK message to the client.
3. For transactions that are not MSC, Fast Path, or conversational, when an IMS application program issues a GU call to retrieve a transaction from the input queue.

If a transaction expires on the input queue before an IMS application program can retrieve it with a GU call, OTMA discards the transaction and, by default, issues message DFS3688I to the OTMA client. However, if the TMAMINPT flag is set in the OTMA input message that expires, instead of the returning a DFS3688I message, OTMA sends back the original input transaction data to the OTMA client.

For each OTMA client, you can optionally configure OTMA to create a symptom dump and issue a DFS554A message for transactions that expire on the input queue by specifying TODUMP=YES in the OTMA client descriptor in the DFSYDTx member of the IMS.PROCLIB data set.

You can also request a symptom dump and DFS554A message for individual messages by setting the TMAMDUMP flag in the OTMA state data prefix of input messages. Setting the TMAMDUMP flag on an individual message overrides a specification of TODUMP=NO for that message only.

For the following transaction types, OTMA monitors the expiration time only until the transaction is enqueued in IMS:

- MSC remote transactions
- Fast Path transactions

OTMA does not monitor the expiration time for the following transaction types:

- IMS conversational transactions
- Transactions queued as a result of a program-to-program switch

When you define an expiration time period in the definition of a transaction, you use the EXPRTIME parameter of either the TRANSACT stage-1 system definition macro or the type-2 dynamic resource definition commands CREATE TRAN or UPDATE TRAN. You can also specify an expiration time period for transactions created by the Destination Creation exit routine (DFSINSX0). Transaction expiration times enabled by any of these methods are not specific to OTMA.

Specifying OTMA transaction expiration time in STCK format

You can specify a transaction expiration time in store clock (STCK) format in the OTMA message header.

When specified in STCK format, the expiration time is the time of day that the transaction expires. If the transaction is either not retrieved by an IMS application program or enqueued for MSC or conversational processing before the expiration time of day, the transaction expires and OTMA discards the transaction.

To specify an expiration time period in STCK format for an OTMA transaction:

- In the state data section of the OTMA header, specify X'01' in the TMAMHIST field.
- In the state data section of the OTMA header, specify the byte offset of the time specification in the TMAMOSXP field. This is the offset of the time specification in the user data section of the OTMA header.
- In the user data section of the OTMA header, code the OTMA client to specify in STCK format the time of day at which the transaction expires.

Related reference:

"Transaction and callout messages" on page 871

"User data section" on page 882

Specifying OTMA transaction expiration time in seconds

You can specify a transaction expiration time as a length of time in seconds in the OTMA message header.

When specified in seconds, OTMA calculates the expiration time by adding the length of time specified to the time at which OTMA receives the transaction from the z/OS cross-system coupling facility (XCF).

For example, if the length of time specified is 10 seconds, the transaction expires 10 seconds after OTMA receives it from XCF. If the transaction is neither retrieved by an IMS application program nor enqueued for MSC or conversational processing before 10 seconds pass, the transaction expires and OTMA discards the transaction.

To specify a length of time for transaction expiration, code the following fields in the state data section of the OTMA header:

- In the TMAMHIST field, specify X'02'.

- In the TMAMOSXP field, specify the length of time in seconds that OTMA uses to calculate the expiration time. OTMA calculates the expiration time by adding this value to the time of day that OTMA receives the transaction from XCF.

Related reference:

“Transaction and callout messages” on page 871

OTMA security

Security for OTMA is enforced by RACF or a similar security product. The following topics describe implementing security for OTMA and OTMA clients by using RACF.

RACF security levels for OTMA

OTMA uses four levels of RACF security: CHECK, FULL, JOIN, NONE, PROFILE. OTMA uses four levels of RACF security: CHECK, FULL, NONE, PROFILE.

Only one RACF security level can be in effect at one time for OTMA globally; however, using the /SECURE OTMA TMEMBER command, you can specify different security levels for individual OTMA clients.

NONE

A system-wide security level. RACF is not called for messages received through OTMA. Specifically:

RACF is not called when IMS receives the connection request (client-bid) from WebSphere MQ or IMS Connect.

RACF is not called to verify that the user ID in the incoming message is a valid user ID (one that has been defined to RACF).

RACF is not called to verify that the user ID in the incoming message is authorized to the IMS command or IMS transaction requested in the message.

The user ID caching scheme is not used.

PROFILE

A message-by-message security level. In other words, each incoming message entered through OTMA is checked to determine whether or not RACF will be called. IMS checks each incoming message independently to see if the security value is set to NONE, CHECK, or FULL. Specifically:

Messages entered from IMS Connect will contain a 1-byte security flag field. The value in this field determines whether or not RACF is called.

Messages entered from the WebSphere MQ-IMS Bridge application will contain a SecurityScope field in the MQIIH structure. The value in this field will determine whether or not RACF is called.

Tip: Consider using the PROFILE security level for situations when application developers set the RACF security level as N (NONE), C (CHECK), or F (FULL) in each incoming message. In this case, the security level set in each message determines whether IMS calls RACF for security checking related to that message. You might not want application programmers deciding on the security for IMS commands and IMS transactions. RACF is called when IMS receives the connection request (client-bid) from WebSphere MQ or IMS Connect.

CHECK

A system-wide security level, which means that RACF is called for messages received through OTMA. Specifically, RACF is called:

- For client-bid connection requests. A cache, or hash table, is built for each OTMA client if the client-bid is successful.

A user ID caching scheme is used in IMS/OTMA environments. The caching scheme also improves authorization checking performance.

A cache, or hash table, is used to store previously verified user IDs. Each OTMA client (IMS Connect, WebSphere MQ for z/OS, or others) has a hash table created in the IMS control region after a successful client bid. Use of the hash table minimizes the number of calls to RACF to VERIFY user IDs. This way, if the same user ID enters multiple messages destined for IMS/OTMA, IMS can check the hash table for a valid entry for the user ID and might be able to avoid the VERIFY call to RACF. The entry for the user ID in the hash table contains a pointer to the accessor environment element (ACEE) for the user ID. The ACEE that is pointed to can be used for resource (command and transaction) FASTAUTH calls to RACF.

- To VERIFY that the user ID in the incoming message is a valid user ID (one that has been defined to RACF).

If the OTMA client (IMS Connect or WebSphere MQ for z/OS) supplied a UTOKEN in the incoming message, IMS supplies the address of the UTOKEN on the VERIFY call to RACF. Use of the UTOKEN in VERIFY processing improves performance. RACF returns an ACEE security control block to IMS for verified user IDs.

- To verify that the user ID in the incoming message is authorized to the IMS command or IMS transaction requested in the message. The address of the ACEE, previously built by RACF during verify authorization processing, is supplied by IMS on the FASTAUTH call to RACF.
- To verify that the user ID in the incoming message is authorized to the IMS transaction code set as the destination on a DL/I CHNG or AUTH call. A cached ACEE is used for these calls, which eliminates performance concerns for application programs that issue many CHNG or AUTH calls.

FULL A system-wide security level, which means that RACF will be called for messages received through OTMA.

FULL has the same characteristics as CHECK, with the following exceptions:

- During the verify processing, RACF is called a second time to build an additional ACEE security control block in the dependent region.
- The dependent region runs under the requestor's user ID. Any resources, such as files, that are accessed from the dependent region must be authorized to that user ID.

Specifying OTMA security

You can make most security specifications for OTMA and OTMA clients by using the system execution procedures, the transaction message prefixes, and the /SECURE OTMA command.

Not every method of specifying OTMA security has the same scope of affect. Some methods allow you to specify security for all members of the OTMA z/OS

cross-system coupling facility (XCF) group, and other methods allow you to specify security for only an OTMA client or an OTMA transaction or command message.

Specifying OTMA security during system definition

The OTMASE= execution parameter in the IMS and DCC execution procedures sets the RACF security level globally for all the OTMA clients in an OTMA z/OS cross-system coupling facility group.

Specifications made using the OTMASE= execution parameter can be overridden by issuing the /SECURE OTMA command.

The levels of RACF security that you can specify by using the OTMASE= parameter include:

- CHECK
- FULL
- NONE
- PROFILE

Related tasks:

“Defining the level of OTMA security checking” on page 769

Related reference:

“RACF security levels for OTMA” on page 799

 Procedures used in IMS environments (System Definition)

Modifying OTMA security online

You can modify security for all OTMA clients globally and for OTMA clients individually by using online commands.

You can modify security for all OTMA clients globally and for OTMA clients individually by using the online type-1 command /SECURE OTMA. You can use the /SECURE OTMA command to override security specifications made by using the OTMASE= parameter in the IMS or DCC execution procedures.

The /SECURE OTMA command can specify the RACF security level globally for the OTMA z/OS cross-system coupling facility (XCF) group or disable RACF security for the OTMA XCF group by using the following parameters:

- /SECURE OTMA CHECK
- /SECURE OTMA FULL
- /SECURE OTMA PROFILE
- /SECURE OTMA NONE

You can specify a RACF security level for individual OTMA clients by issuing the following commands:

- /SECURE OTMA CHECK TMEMBER *tmember_name*
- /SECURE OTMA FULL TMEMBER *tmember_name*
- /SECURE OTMA PROFILE TMEMBER *tmember_name*
- /SECURE OTMA NONE TMEMBER *tmember_name*

Security specifications made for individual OTMA clients override the global security settings made for the rest of the OTMA XCF group.

If you specify `/SECURE OTMA NONE`, IMS does not use RACF for security verification, regardless of what security is specified by the class for a client-bid request or for transactions.

When RACF security checking is disabled for OTMA, you can issue only the following default IMS commands through OTMA:

- `/BROADCAST`
- `/LOCK`
- `/LOG`
- `/RDISPLAY`
- `/UNLOCK`

Complete information for how to use these command is provided in *IMS Version 11 Commands, Volume 1: IMS Commands A-M* and *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

Refreshing RACF ACEEs for OTMA

When OTMA security is enabled, for performance purposes OTMA caches in memory RACF user profiles as a RACF accessor environment element (ACEE).

The ACEE is then used in subsequent calls to RACF to determine the user ID's authorization to the IMS command or IMS transaction requested in the input message. Using the ACEE reduces I/O to RACF.

When changes are made to a user profile in the RACF database on DASD, the changes are not reflected in the cached ACEE until the ACEE is refreshed. To ensure that cached ACEEs reflect the latest RACF security definitions, OTMA clients can specify an aging value that defines the length of time between automatic refreshes of cached ACEEs.

The ACEE expiration value is specified during the client-bid time.

You can refresh ACEEs manually by issuing the `/SECURE OTMA REFRESH` command. You can refresh all OTMA ACEEs globally, all of the ACEEs for specific OTMA clients, or all of the ACEEs for an individual user ID.

To refresh all OTMA ACEEs globally, issue the command `/SECURE OTMA REFRESH`. To refresh the ACEE for an individual OTMA client, issue the command `/SECURE OTMA REFRESH TMEMBER tmember_name`. To refresh only those ACEEs for a specific user ID, issue the command `/SECURE OTMA REFRESH USER user_id`.

You can also rebuild the ACEE table by issuing the `/STOP` and `/START OTMA` commands.

For example, you might need to refresh the ACEEs for an individual user ID if your RACF Security Administrator modifies a user's access privileges for OTMA in RACF. The user's old access privileges might remain unchanged in online memory unless the RACF ACEEs cached by OTMA are refreshed, either automatically when the ACEE aging limit is reached or manually by issuing the `/SECURE OTMA REFRESH` command.

To ensure that the cached ACEEs are refreshed promptly, the RACF security administrator should ask the IMS operator to refresh all OTMA cached ACEEs. The RACF security administrator must also provide the RACF user profile name that must be entered in the `USER` parameter of the `/SECURE OTMA REFRESH USER`

user_id command. When the IMS operator issues the /SECURE OTMA REFRESH USER *user_id* command, IMS scans all OTMA cached ACEEs and refreshes only the ACEEs that contain the specified user ID.

Related reference:

 /SECURE command (Commands)

Security specifications in OTMA message prefixes

Security specifications in the OTMA message prefix are in the security-data section. The security-data section is mandatory for every transaction or command and is optional for OTMA protocol commands.

The specifications you can make for security in an OTMA message prefix include:

- The RACF security levels to be used with this message, including:
 - FULL
 - CHECK
 - NONE
- The user token
- The user ID
- The SAF profile

In the state data of a client-bid request, you can specify an accessor environment element (ACEE) aging interval. An ACEE aging interval represents the amount of time an ACEE can be used before being refreshed. If the aging interval has expired when an input message is received by OTMA, OTMA refreshes the ACEE before validating the message.

Related reference:

Chapter 47, “OTMA message prefix,” on page 855

RACF security classes used by OTMA

To have RACF secure transactions or commands submitted to IMS through OTMA, you must define the security to be enforced for them to RACF.

Transaction security definitions are stored in the RACF TIMS class. Command security definitions are stored in the CIMS class. If a transaction is not in the TIMS class or a command is not in the CIMS class, the transaction or command is allowed regardless of any options you might specify by using the /SECURE OTMA command.

When you enter an OTMA command, OTMA issues a RACHECK to validate the command. OTMA passes only the command verb to DFSCCMD0 for verification, not the entire CVB control block.

If you are using RACF to secure asynchronous hold queues from unauthorized users of the RESUME TPIPE call, you must define either the default RIMS resource or a Rxxxxxxx resource class, where “xxxxxxx” is the value of the RCLASS parameter in either the SECURITY macro statement or the DFSDCxxx PROCLIB member. The RACF resource class you define must include the names of the protected asynchronous hold queues and the user IDs that are authorized to access each queue.

Securing messages on the asynchronous hold queue

You can protect messages on asynchronous hold queues from unauthorized use of the RESUME TPIPE call by using either RACF, the OTMA Resume TPIPE Security exit routine (DFSYRTUX), or both.

When security is enabled for tpipe hold queues, the user ID issuing the RESUME TPIPE call must be authorized to access the TPIPE name that is contained in the RESUME TPIPE call message before any of messages are sent to an OTMA client.

The security checking performed by RACF and the security checking performed by the DFSYRTUX exit routine are optional. They can be used in combination or either one by itself. If both RACF and the DFSYRTUX are used, RACF is called first before giving control to the DFSYRTUX exit routine, in which case, the DFSYRTUX exit routine can override RACF, depending on your needs.

Securing asynchronous hold queues by using RACF

When RACF security checking is enabled for an asynchronous hold queue, the authorization logic verifies and validates the security header and authorizes the user ID under the TPIPE name.

When a Resume TPIPE call is received, RACF security checking is performed only if a Resume TPIPE resource class (RIMS or Rxxxxxxx) exists and the tpipe name specified on the call is defined in the resource class.

Regardless of whether RACF security is enabled, you can use the OTMA Resume TPIPE Security exit routine (DFSYRTUX). If both the DFSYRTUX exit routine and RACF are used, the RACF security is always called first. In such a case, the DFSYRTUX exit routine can override the results of the RACF procedure.

To enable Resume TPIPE security:

1. If one does not already exist, define a Resume TPIPE resource class (RIMS or Rxxxxxxx) by using the RCLASS keyword on the SECURITY system definition macro. During IMS startup, if no Resume TPIPE resource class is defined to IMS, IMS issues message DFS3187W. After IMS is running, no further warnings are issued to alert you to the absence of a Resume TPIPE resource class.
2. In the Resume TPIPE resource class, specify the tpipe names of the asynchronous hold queues to be protected and the user IDs that are authorized to access the queues.
3. Specify an appropriate level of RACF security for OTMA by using either the OTMASE startup parameter or the /SECURE OTMA command. The appropriate levels of OTMA security that you can specify are FULL, CHECK, or PROFILE. If a security level of PROFILE is specified, the resume tpipe request message must specify either FULL or CHECK.
4. Code the resume tpipe request messages to access the RACF-protected asynchronous hold queues. The resume tpipe request messages must include:
 - The tpipe name in the control data section of the OTMA prefix
 - A user ID in the security section of the OTMA prefix
 - If the OTMA security level is PROFILE, a security flag setting of either FULL or CHECK

Related information:

 DFS3187W (Messages and Codes)

Securing asynchronous hold queues with the OTMA Resume TPIPE Security exit routine (DFSYRTUX)

The OTMA Resume TPIPE Security exit routine (DFSYRTUX) provides one of two possible layers of security for RESUME TPIPE calls issued to retrieve message queued to the OTMA asynchronous hold queue.

When security for RESUME TPIPE calls is enabled, the DFSYRTUX user exit checks the caller's authority when the RESUME TPIPE call is initiated, but before retrieving the messages from the hold queue.

You can use DFSYRTUX exit routine either with or without RACF security checking for the RESUME TPIPE call. If both the DFSYRTUX exit routine and RACF are used, the RACF security is always called first. In such a case, the DFSYRTUX exit routine can override the results of the RACF procedure.

When security for the RESUME TPIPE call is enforced by both RACF and the DFSYRTUX exit routine, the DFSYRTUX exit routine is invoked regardless of the success or failure of the RACF security procedure. The DFSYRTUX exit routine can accept the results of the RACF security check, override the results, or enforce more restrictive security rules. An example of a more restrictive rule might be to authorize a user to access the output messages only within a specific period of time during the day.

When authorization is successful, output messages in the hold queue are returned to IMS Connect.

When authorization fails, a rejection message (NAK) of the RESUME TPIPE call is sent to the client.

To bypass the DFSYRTUX exit routine, delete the DFSYRTUX module from the library.

Related reference:

 OTMA Resume TPIPE Security exit routine (DFSYRTUX) (Exit Routines)

General OTMA security considerations

A number of general security considerations exist for OTMA that you should be aware of.

- If you use RACF (or an equivalent product) for security, define the `IMSXCF.group.client_member_name` in the FACILITY class.
If you define the `IMSXCF.group.client_member_name` in the FACILITY class, and if IMS security is not set to NONE, the user token for the client-bid request must be valid and the user must have READ access to the FACILITY class.
If the user token for a client-bid request fails RACF verification, the client receives a NAK message from the server.
- Authorize the z/OS cross-system coupling facility client for z/OS.
- If you define your OTMA applications with full security, the security environment is kept until the application ends.
- After IMS receives messages from OTMA, when OTMA security is activated, IMS calls RACF to verify that the user ID in the incoming message is a valid RACF user ID. IMS is not passed a password for the user ID, so the call to RACF is to verify the user ID only. If a password must be validated, it must be validated before sending the message to IMS.

- IMS uses the UTOKEN in the input message in the call to RACF not only to verify the user ID, but also to create a security control block in the IMS control region to represent each verified user ID. The security control blocks built in the IMS control region, representing verified RACF user IDs, are called accessor environment elements or ACEEs.

Related concepts:

 IMS security (System Administration)

Related reference:

 ICAL call (Application Programming APIs)

Using DL/I calls in an OTMA environment

Certain DL/I calls have special considerations when used with OTMA.

CHNG

If a CHNG call is issued from an OTMA submitted transaction, the destination is assumed to be the same OTMA client (the tpipe name is set by the CHNG call). This behavior can be altered by the OTMA Prerouting and Destination Resolution exit routines.

An IMS application program that issues a CHNG call to an alternate PCB (specifying an options list) does not cause IMS to call the OTMA Prerouting and Destination Resolution exit routines to determine the destination. However, an IMS application program that issues a CHNG call to an alternate PCB (specifying an APPC descriptor) does cause IMS to call the OTMA exit routines to determine the destination.

The application program can still issue ISRT calls to the I/O PCB to send data to an OTMA destination.

OTMA application programs can use CHNG and ISRT calls for APPC destinations.

INQY (null)

An INQY call issued for an OTMA destination returns the following information: the transaction pipe name, the client z/OS cross-system coupling facility member name, the user ID, the group name, and the synchronization levels.

ICAL IMS application programs issue the ICAL call to send synchronous callout requests to a data or service provider that is external to the IMS installation. OTMA routes callout requests initiated by the ICAL call to a hold-queue capable OTMA client and routes the reply back to the waiting IMS application program. OTMA provides the ability to specify a timeout value in the OTMA destination descriptor for synchronous callout requests, which can be overridden by the a timeout value specified in the ICAL call.

PURG An IMS application program that issues a PURG call causes IMS to call the OTMA Prerouting and the Destination Resolution exit routines to determine the destination.

SETO An IMS application program that issues a SETO call does not cause IMS to call the OTMA Prerouting and the Destination Resolution exit routines to determine the destination.

Existing IMS application programs that issue SETO calls might not run as expected because a return code is returned to the program if it is

processing an OTMA-originated transaction. APPC/IMS application programs that issue SET0 calls might need modification if they require implicit OTMA support.

One way to make these application programs work is to use an INQY call before issuing the SET0 call. The application program can use the output from the INQY call to determine if a transaction originated from an OTMA client, and not issue the SET0 call.

For those DL/I calls that cause IMS to call one of the OTMA exit routines, IMS only calls the exit routines if the destination has not yet been set (for example, by another DL/I call).

To initiate protected conversations (such as accessing multiple resource managers' resources under one unit of recovery in an z/OS Resource Recovery Services environment), the client-adaptor code (OTMA user) must acquire and own a private context and provide the context ID in the state-data section of the message prefix.

Definition: A *context* is a z/OS entity under which resource managers perform work; a private context is required in this environment.

During message traffic between IMS and the client, if the context-ID field in the message header is non-zero, protected conversation processing occurs.

Related reference:

 Transaction management (Application Programming APIs)

OTMA program-to-program switch processing

Two types of message switch occur in OTMA: commit-then-send, and send-then-commit.

For OTMA commit-then-send input messages (also called asynchronous or commit mode 0 messages), the program switch always results in another commit mode 0 (CM0) message. For OTMA send-then-commit input messages (also called synchronous or commit mode 1 messages), the program switch results vary, depending on whether:

- there is an ISRT call to the I/O PCB
- an express PCB is used for the switch
- there is a switch to multiple programs
- the IMS start-up parameter OTMAASY=Y is specified
- the transaction is protected

A P2P switch for a commit mode 1 (CM1) input message, therefore, could be another CM1 message, a DFS2082 message, or a CM0 message. In addition, some OTMA clients, for example WebSphere MQ, can accept a CM0 output message for a CM1 input message; others, however, may not.

The following topics provide usage scenarios for different send-then-commit message switches.

OTMA single-stream program switch

The single-stream program switch is shown in the following figure.

Program A switches to program B, and Program B switches to Program C, which then inserts back to the I/O PCB. This model of program flow delivers the send-then-commit output message successfully. Single-stream means that the program switches occur one after another. No express PCBs are used in the P2P message switches.




Figure 137. Single-stream program switch

OTMA program switch without ISRT to I/O PCB

When several switches occur in sequence and none inserts back to the I/O PCB, message DFS2082 is sent back to the OTMA client.

The last program switched to, Program C, does not insert back to the I/O PCB. IMS therefore generates message DFS2082 for the OTMA client.

Attention: If program C runs in a remote IMS through MSC and does not insert back to the I/O PCB, the remote IMS does not issue message DFS2082. However, in this case, the OTMA client program might hang and the front-end IMS control region will experience a build up of its control blocks. This kind of build up could result in a storage-related system outage. Restarting IMS releases the control blocks.

OTMA program switch with express PCB

A P2P message switch with an express PCB can lead to a commit-then-send output message.

Program A uses an express PCB rather than a non-express PCB to perform the P2P message switch. The output from Program B is commit-then-send because using the express PCB forces Program B to be processed asynchronously. When a program is processed asynchronously and inserts back to the I/O PCB, the output message is sent as a commit-then-send message. However, if Program A also

switches to Program C using a non-express PCB, Program C then inserts back to the I/O PCB. The output from C will be a send-then-commit message.

OTMA program switch to multiple programs

After a program inserts back to the I/O PCB, the rest of the program to program message switch, if any, is processed asynchronously.

For example, Program A switches to Program B, which inserts back to the I/O PCB. The output from Program B will be a send-then-commit message. Program B then switches to program C, which will be processed asynchronously.

A “race” condition can occur when a program switches to multiple programs. Program A switches to multiple programs using non-express PCBs. Only one switched-to program, the one scheduled first, is processed synchronously. The rest of the switched-to programs are processed asynchronously. If the program processed synchronously inserts back to the I/O PCB, the output message is a send-then-commit message.

In some cases, one of the multiple programs could be a remote program. This program flow is shown in the following figure. Program A switches to remote Program B through MSC. Program B first launches a new program, Program C, in the local IMS and then inserts a response to the OTMA client through the local IMS. Depending on what happens first (scheduling of Program C or the processing of the response for the OTMA client) in the local IMS, an unwanted DFS2082 message could be sent to the client. This is also a race condition. If Program C gets processed first in the local IMS, a DFS2082 message is sent. If the response is processed first, the expected output from Program B is delivered synchronously using send-then-commit.




Figure 138. Race condition resulting from program switch to multiple programs

To avoid creating the race condition in these circumstances, you can do any one of the following:

- Modify your programs to avoid multiple program-to-program switches within the same transaction for send-then-commit (CM1).
- Use commit-then-send (CM0) input when performing multiple program-to-program switches within the same transaction.
- Use the IMS start-up parameter OTMAASY to serialize P2P message switch processing.

When you use the OTMAASY parameter to avoid a race condition, you can create a program switch model similar to the single-stream model. For example, in the preceding figure, Program B, for which the response mode transaction is processed synchronously, can deliver the send-then-commit (synchronous) output message. Program C, which is running with a non-response mode transaction, processes the message asynchronously.

Related tasks:

“Specifying asynchronous delivery of program-to-program switch output messages” on page 770

Related reference:

 Parameter descriptions for IMS procedures (System Definition)

OTMA program switch for protected transactions

For a non-conversational program that performs a P2P message switch for a protected transaction, ABENDU0711, with reason code 1D, will be returned. For a conversational program, the program receives an X6 status code.

Other OTMA program switch considerations

The following considerations also apply to program-to-program switching.

- The P2P message switch is not supported for OTMA protected messages (send-then-commit input with synclevel = SyncPt).
- If a non-conversational program performs a program-to-program message switch to a program in a Shared Queues environment, the program in the SQ environment must be running on the same IMS where the first program gets scheduled.
- If an input conversational transaction occurs, only the message-switched-to continuation of the conversation is scheduled synchronously. All other transactions are scheduled asynchronously.
- In a shared queues environment that has both synchronous APPC/OTMA support (AOS=Y on the DFSDCxxx PROCLIB member) and z/OS Resource Recovery Services (RRS) support (RRS=Y on the startup procedure) enabled, an application program running on a back-end IMS system that initiates an outbound APPC protected conversation with another IMS system is restricted to a single program-to-program switch to the same destination transaction.

If an application program performs multiple program-to-program switches after allocating an APPC outbound protected conversation on another IMS system, the results are unpredictable and can include a WAIT-RRS/PC condition in the MPP dependent region.

Chapter 45. The OTMA client

The OTMA environment includes a server and one or more clients.

This chapter explains how a client interacts with the server to process IMS transactions.

What is an OTMA client?

An *OTMA client* is a z/OS application program that sends transactions to an IMS server and receives output. The application program must be a member of an z/OS cross-system coupling facility (XCF) group and use the OTMA protocol.

Heterogeneous (non-z/OS) networks can connect with z/OS in many ways. The following figure shows some of the possible applications that use XCF. These include:

- WebSphere MQ applications
- OEM applications
- IMS Connect applications
- DCE/RPC applications
- Other IBM applications

Any of these can connect to an OTMA client to communicate with IMS.




Figure 139. Applications that use XCF to connect to IMS on z/OS

An OTMA client is the gateway by which transactions from outside IMS can enter IMS.

OTMA processing involves:

1. A client sends a transaction or command to the server (IMS).
2. The server returns output to the client.

OTMA naming conventions

The names of OTMA clients and transaction pipes must follow a convention defined by IMS.

Names for OTMA clients and transaction pipes must adhere to the following conventions:

- Must be character type A (A-Z, 0-9, @, \$)
- Must begin with a non-blank character
- Must be padded with blanks if shorter than the maximum length (16 for a client name, 8 for a transaction-pipe name)
- Cannot contain embedded blanks
- Cannot be a reserved word (for example, "TO" or "SECURITY")
- Cannot begin with "DFS" or "DBCDM"
- Cannot be an IMS keyword (for example, "LINE" or "NODE")

In addition, transaction-pipe names cannot:

- Duplicate an IMS transaction name.
- Have the same name as the z/OS system console (for example, "WTOR"), IMS MTO, or secondary MTO.

IMS does not perform uppercase translation. If lowercase characters are used, the client receives a negative acknowledgment (NAK) response from the server.

Messages sent by OTMA clients

An OTMA client communicates with IMS by sending messages.

First, a user enters application data using a device or program that is connected to the client. Next, the client adds some information (the message prefix) and sends the message to IMS. Output from IMS is sent to the client as a message, and the client uses the message prefix to route the data to the correct device or program.

Parts of the OTMA message prefix

The OTMA message prefix has the following sections:

- Message-control information
This section includes the transaction-pipe name, message type, sequence numbers (if any), the time out value for send-then-commit messages, and various flags and indicators.
- State data
This section includes a destination override (if any), map name, synchronization level, commit mode, tokens, and server state.
- Security data
This section includes the user ID, user token, and security flags.
- User data
This section includes any special information needed by the client.

Following the message prefix is the application-data section of the message. This section contains either the data to be sent to IMS for processing or the IMS response.

OTMA message-prefix rules

Because a message can have a single segment or multiple segments, the following rules apply to OTMA message prefixes:

- Single-segment messages can have the full prefix (message-control information, state data, security data, and user data).
- Only the first segment of multi-segment messages has the full prefix. Subsequent segments are sent with only the message-control information and application-data sections.
- Acknowledgment (ACK or NAK) messages sent by IMS only return the first input buffer. This message carries the full prefix, and the application-data section (if it is included in the client request).

Sequence numbers used by OTMA

OTMA uses two types of sequence numbers for messages: send-sequence numbers and recoverable sequence numbers. Send-sequence numbers and recoverable sequence numbers are used differently in OTMA.

Using send-sequence numbers

Send-sequence numbers are used for input and output messages. Send-sequence numbers should be incremented by a client for every input message. When IMS sends output to a client, the send-sequence numbers in the output message are also incremented. The send-sequence numbers are used for all the OTMA input/output messages. The send-sequence numbers in the input messages are also used to identify multi-segments.

Example: There is a two-segment OTMA input message. The first segment message will have send-sequence number=XXX and segment number=1. The second segment message should have the same send-sequence number=XXX and segment number=2. OTMA chains the two-segment message together because the send sequence numbers are the same.

OTMA uses send-sequence numbers in the following ways:

- All ACK and NAK messages from IMS use the send-sequence numbers submitted by the client on input.
- All OTMA commands that IMS sends to the client have send-sequence number 0 (zero). And, except for the resynchronization flows, these OTMA commands are all single segment.
- Send-sequence numbers for IMS error messages and IMS transaction output are set for each transaction pipe. The send-sequence number for a given transaction pipe is incremented by one for each message, and it is never 0 (zero). When the sequence number exceeds 4 294 967 295 (the 32-bit maximum), it is reset to 1.

Using recoverable sequence numbers

Recoverable sequence numbers are used only to control resynchronization. If a client does not support resynchronization, recoverable sequence number=0 (zero). Resynchronization is only valid for synchronized tpipes and commit-then-send input/output. The recoverable sequence numbers are also incremented for every input/output message. Resynch support has an added logic to check if the recoverable sequence numbers are properly incremented. If the sequence numbers are not properly incremented, a NAK is sent. Because the resynch is dependent on the recoverable sequence numbers, the resynch must be correct for every input/output. Recoverable sequence numbers apply to transaction pipes, which use them to control resynchronization.

Related concepts:

“Client/server resynchronization with OTMA” on page 830

Related reference:

Chapter 47, “OTMA message prefix,” on page 855

| **Sending type-1 commands from an OTMA client**

| You can submit type-1 IMS commands from an OTMA client application. Some
| restrictions and recommendations apply.

Most commands should be submitted with the send-then-commit (CM1) protocol. However, the following commands require the commit-then-send (CM0) protocol:

- /DBDUMP DATABASE
- /DBRECOVERY AREA | DATABASE
- /START AREA | DATABASE
- /START REGION
- /STOP AREA | DATABASE
- /STOP REGION

The DISPLAY TRANSACTION command behaves differently when it is issued from an OTMA client. The command sends its output directly to the client, not to the IMS master terminal. Depending on the setting of the extended-response-requested flag in the message-control information section of the message prefix, the output is either in an architected format (only supported for this command) or in the standard IMS format.

Recommendation: Because the client must use the commit-then-send (CM0) flow, the output from these commands cannot be tied to the input command. The OTMA prefixes are not replicated (the only field common to both the input and the output is the transaction-pipe name). Therefore, configure the client to submit IMS commands using a transaction pipe that the client reserves for IMS command processing.

Restrictions:

OTMA clients cannot submit IMS commands from the following types of subsystems:

- An IMS Remote Site Recovery (RSR) tracking subsystem
- An IMS Extended Recovery Facility (XRF) alternate subsystem
- A CICS-IMS DBCTL subsystem

Related reference:

 /DISPLAY TRAN command (Commands)

OTMA commit processing

OTMA can control how IMS commits transactions: they can be either commit-then-send or send-then-commit.

Definitions:

- For *commit-then-send* transactions (the IMS standard flow), IMS processes the transaction and commits the data before sending a response to the OTMA client.
- For *send-then-commit* transactions, IMS processes the transaction and sends a response to the OTMA client before committing the data.

Q: What is the major difference between the commit-then-send processing option and the send-then-commit processing option?

A: The commit-then-send processing option commits the transaction output as part of sync-point processing, and then delivers the output to the client later.

The send-then-commit processing option delivers the transaction output first, receives an acknowledgment from the client, and then completes the sync-point processing.

Q: What happened to the commit mode 0 and commit mode 1 processing options?

A: Commit mode 0 is now called “commit-then-send”, and commit mode 1 is called “send-then-commit”. Because the terms “commit-then-send” and “send-then-commit” are more intuitive when referring to these processing options, the terms “commit mode 0” and “commit mode 1” are no longer used.

For an OTMA transaction, a client can receive one of the following from IMS:

- An ACK message for the input, followed by any output messages.
In addition send-then-commit transactions will also receive an ACK message followed by a “deallocate” flow (indicated when the commit-confirmation flag in the message-control information section of the message prefix is set to either Committed or Aborted).
- A NAK message with a sense code.
- A NAK message with the processing flag set to Error Message Follows in the message-control information section of the message prefix. The subsequent message has the same message prefix as the NAK message and has the IMS error message in the application-data section of the message prefix.

Summary of OTMA commit processing

The following table summarizes the differences between commit-then-send and send-then-commit processing.

Several variables are listed in the first column; the differences between processing options are described in the next two columns. Following the table are some usage notes to be aware of.

Table 134. Commit-then-send versus send-then-commit processing

Variables	Commit-then-send	Send-then-commit
Conversational	Client receives a NAK message.	Supported.
Fast Path	Client receives a NAK message.	Supported.
Non-conversational and non-Fast Path transactions	IMS commits after enqueueing the output to the client. The output is delivered later.	IMS sends output to the client and then commits.
Enqueue the input?	Yes.	Yes.
Enqueue the output?	Yes.	No.
Synchronized transaction pipe specified?	Supported.	Client receives a NAK message.
Timeout interval enforced?	No.	Yes, if a timeout value is specified with either synclevel=confirm or synclevel=synchpt.

Notes:

- IMS conversations cannot use the commit-then-send commit mode.

- Send-then-commit input and output is irrecoverable.
- For irrecoverable output (send-then-commit), IMS requests an acknowledgment if the synchronization level is set to Confirm.
- For a recoverable transaction, IMS always requests an acknowledgment for an output message.
- For commit-then-send transactions, IMS always requests an acknowledgment.
- Synchronized transaction pipes can only be used for commit-then-send transactions.
- When a send-then-commit (CM1) input message is sent to a transaction, OTMA treats that transaction as RESPONSE mode even if the transaction is defined as NONRESPONSE. If the application does reply to the IOPCB, the output is send-then-commit. If the application does not reply to the IOPCB and does not complete a program-to-program message switch, then OTMA responds with a DFS2082 RESPONSE MODE TRANSACTION TERMINATED WITHOUT REPLY message.

Related information:

 [DFS2082 \(Messages and Codes\)](#)

Sample OTMA commit processing flows

To show the differences between the two commit modes, the topics listed below show sample flows of data between IMS and clients for each commit mode.

Related reference:

“Sample OTMA message flows” on page 824

Commit-then-send flow

The commit-then-send flow, also known as the IMS standard flow, enqueues IMS output before sending it to the client. Use this flow for standard transaction processing.

To use the standard flow, specify commit-then-send (commit-mode-0) in the state-data section of the message prefix. This sample flow assumes the following:

- The transaction pipe is synchronized. IMS maintains sequence numbers for recoverable input and output for the transaction pipe.
- Acknowledgment is always requested (by both IMS and the client).

If NAK response is received by IMS, then the output is returned to the queue and will be delivered later.

The flow is illustrated in the following figure. Following the figure is a sequential list that provides more details on the flow.




Figure 140. Commit-then-send (IMS standard) flow

The sequence of the flow illustrated in the figure is:

1. Transaction initiated (response required/synchronized tpipe)
2. Transaction is inserted to SMB
3. ACK
4. GU call followed by ISRT to IOPCB
5. Sync Point
6. Output is enqueued to tpipe, and DB is committed
7. Transaction completes
8. Output is sent with response requested
9. ACK
10. Output is dequeued.

An example of the flow of the message activity for a single commit-then-send transaction pipe is shown in the following figure. Following the figure is a sequential list that provides more details on the flow.




Figure 141. Sample message flow for commit-then-send flow

The sequence of flow shown in the figure is:

1. Tran1
2. ACK to Tran1
3. Tran2
4. Output of Tran1
5. ACK to Tran2
6. Tran3
7. ACK to Tran3
8. Tran4
9. ACK to output of Tran1
10. ACK to Tran4.

Related reference:

“Transaction and callout messages” on page 871

Related information:

[DFS2082 \(Messages and Codes\)](#)

Send-then-commit flow

The send-then-commit (commit-mode 1 or CM1) flow sends IMS output before IMS completes synchronization-point (hereafter referred to as sync-point) processing.

To use the send-then-commit flow, specify Commit Mode 1 in the state-data section of the message prefix. This sample flow assumes the following:

- The transaction pipe is not synchronized.
- The synchronization level is specified as None in the state-data section. Therefore, IMS does not request a response (an ACK) when sending output.

The flow is illustrated in the following figure. Following the figure is a sequential list that provides more details on the flow.




Figure 142. Send-then-commit flow

The sequence of flow shown in the above figure is:

1. Transaction initiated
2. Transaction inserted to SMB
3. GU call followed by ISRT to IOPCB
4. Sync point started
5. Output is sent. No response is requested; response is requested only when sync=confirm is specified.
6. Commit confirmed; IMS completed sync point
7. Transaction completes

An example of the flow of the message activity for a single transaction pipe is illustrated in the following figure. Following the figure is a sequential list that provides more details on the flow.




Figure 143. Sample message flow for send-then-commit flow

The sequence of flow shown in the above figure is:

1. Tran1
2. Tran2 request/response
3. Tran3 request/response
4. ACK to Tran3
5. Output of Tran1
6. ACK to Tran2
7. Tran4
8. Output of Tran4
9. Confirm of Tran4
10. Output of Tran2
11. Output of Tran3
12. Confirm of Tran1
13. Confirm of Tran3
14. Confirm of Tran2

As shown in the preceding figure, the client can receive a confirmation for output before receiving the actual output, because z/OS cross-system coupling facility (XCF) does not guarantee that all messages are sent in sequential order. The client must be able to handle this situation during message-receipt processing or by using the XCF Message exit routine.

If a send-then-commit transaction does not run successfully, OTMA sends message DFS2082 to the client and does not perform an insert to the IOPCB.

Send-then-commit flow with confirm

The send-then-commit flow assumes no synchronization for the transactions as they are processed by IMS.

This topic shows a flow in which all transactions are confirmed as they are received (each message requests a response). The sample illustrated in the figure below assumes the following:

- Commit Mode 1 is specified in the state-data section of the message prefix.

- The transaction pipe is not synchronized.
- The Synchronization Level is specified as Confirm in the state-data section.

If NAK is received by IMS, then a user 119ABEND occurs in the application and IMS issues a DFS554 message to the client.

Following the figure below is a sequential list that provides more details on the flow.




Figure 144. Send-then-commit with confirm flow

The sequence of flow shown in the above figure is:

1. Transaction initiated
2. Transaction inserted to SMB
3. GU call followed by ISRT to IOPCB
4. Sync point start
5. Output sent; response requested
6. ACK
7. DB is committed; commit is confirmed; IMS completed sync point
8. Transaction completed

Sample OTMA message flows

The following topics show some sample message flows, and describe how various fields in the message prefix are set.

In the figures, the following abbreviations are used for parts of the message prefix:

MC	Message-control information section
SD	State-data section
SE	Security-data section
US	User-data section
AP	Application-data section

The sample flow diagrams show which parts of the prefix are mandatory for a given message and which are not applicable. Optional fields and prefix sections are enclosed in parentheses.

For transactions submitted by clients, the following principles apply:

- After IMS sends an ACK message to a client, IMS sends a commit confirmation (indicating that the transaction committed successfully or was aborted).
- The commit confirmation terminates a client transaction.

Related reference:

“Sample OTMA messages” on page 884

Client-bid message flow

The following figure shows a client-bid flow, where the client attempts to connect to the server.

This flow can occur when the client has already joined the z/OS cross-system coupling facility group and notices that a server has joined the group. The client-bid flow is:

1. Client-Bid: MC, SD, SE
2. ACK: MC, SD, SE




Figure 145. Client-bid flow

The following table shows the contents of the OTMA message prefix in the first flow of a client-bid exchange: the client-bid request. The table shows the contents of the message control data, state data, and security data sections of the OTMA message prefix.

Table 135. Contents of the OTMA message prefix in the first client-bid flow of a client bid exchange

Message prefix section	Contents of prefix section
Message control data	Architecture level = 1 Message type = command Response flag = response requested Command type = client-bid Prefix flag = state data + security data
State data	The state data format for command messages applies to these fields: Length Member name Originator's token Destination token (DRU exit name) MaxBlocksize Aging value Hash table size
Security data	(Utoken)

The following table shows the contents of the OTMA message prefix in the second flow of a client-bid exchange: the acknowledgement. The table shows the contents of the message control data, state data, and security data sections of the OTMA message prefix.

Table 136. Contents of the OTMA message prefix in the second, acknowledgement flow of a client-bid exchange

Message prefix section	Contents of prefix section
Message control data	Architecture level = 1 Message type = command and response Response flag = ACK Command type = client-bid Prefix flag = state data + security data
State data	The state data format for command messages applies to these fields: Length Member name Originator's token Destination token (DRU exit name) MaxBlocksize Aging value Hash table size
Security data	Utoken

Server-available flow

The following figure shows the message flow of a Server-Available exchange, where the server attempts to connect to the client.

This flow only occurs when the server is already joined to the z/OS cross-system coupling facility (XCF) group and recognizes that a client joins the group. A client should not wait for the server to recognize that it has joined the XCF group; the client should send its client-bid message as soon as it joins the group.

The client should ignore a Server-Available message after it has successfully completed its client-bid request and connected to the XCF group. The flow shown is:

1. Server-available: MC, SD, SE
2. Client-Bid: MC, SD, SE
3. ACK: MC, SD, SE




Figure 146. Server-available flow

The following table shows the contents of the message prefix. The flow step is listed, with the message flow type, message prefix section, and associated contents of the message prefix section for the prefixes MC, SD, and SE. The numbers used to show sequence in both the preceding figure and following table are not part of the actual message prefix.

Table 137. Contents of server-available flow message prefix.

Flow step	Message flow	Message prefix section	Content of prefix section
1	Server Available	MC	Architecture level = 1 Message type = command No response flag Command type = Server Available
		SD	The state data format for command messages applies to these fields: Length Member name Originator's token Destination token
2	Client-bid	MC	Architecture level = 1 Message type = command Response flag = response requested Command type = client-bid Prefix flag = state data + security data
		SD	The state data format for command messages applies to these fields: Length Member name Originator's token Destination token (DRU exit name) MaxBlocksize Aging value Hash table size
		SE	(Utoken)
3	ACK	MC	Architecture level = 1 Message type = command and response Response flag = ACK Command type = client-bid Prefix flag = state data + security data
		SD	The state data format for command messages applies to these fields: Length Member name Originator's token Destination token (DRU exit name) MaxBlocksize Aging value Hash table size
		SE	Utoken

Commit-then-send transaction flow

The following figure shows the flow for a commit-then-send transaction, where the client submits a transaction to the server for processing.

The flow is:

1. Transaction ABC: MD, SD, SE, (US), AP
2. ACK: MC, SD, SE, (US)
3. Transaction output: MC, SD, (US), AP
4. ACK: MC, SD




Figure 147. Commit-then-send transaction flow

The following table shows the contents of the message prefix. The flow step is listed, with the message flow type, message prefix section, and associated contents

of the message prefix section for the prefixes MC, SD, SE, US, and AP. The numbers used to show sequence in both the preceding figure and following table are not part of the actual message prefix.

Table 138. Contents of commit-then-send transaction flow message prefix.

Flow step	Message flow	Message prefix section	Content of prefix section
1	Transaction 'ABC'	MC	Architecture level = 1 Message type = transaction Response flag = response requested Transaction-pipe name Prefix flag = SD/SE/(US)/AP Send-sequence number
		SD	Length Synchronization flag = Commit Mode 0 Synchronization level = Confirm or None (Map name) (Correlator) Length of server user data = 0
		SE	Length (Security flag) Length of fields User ID length (User ID type = 02) (User ID) Profile length (Profile type = 03) (RACF group) Utoken length (Utoken type = 00) (Utoken)
		(US)	This optional section is returned with the transaction output: Length User data
		AP	Length ZZ application data ('ABC' in the example)
2	ACK	MC	Architecture level = 1 Message type = transaction and response Response flag = ACK Transaction-pipe name Prefix flag = SD/SE Send-sequence number
		SD	Length Synchronization flag = Commit Mode 0 Synchronization level = Confirm or None (Map name) (Correlator) Length of server user data = 0
		SE	Length (Security flag) Length of fields User ID length (User ID type = 02) (User ID) Profile length (Profile type = 03) (RACF group) Utoken length (Utoken type = 00) (Utoken)
		(US)	This optional section is returned with the transaction output: Length User data
3	Transaction Output	MC	Architecture level = 1 Message type = data Response flag = response requested Transaction-pipe name Prefix flag = SD/(US)/AP Send-sequence number Server token
		SD	Length Synchronization flag = Commit Mode 0 Synchronization level = Confirm or None (Map name) Server token (Correlator) Length of server user Data (Server user data)
		(US)	This optional section is returned with the transaction output: Length (User data)
		AP	Length ZZ Transaction output data

Table 138. Contents of commit-then-send transaction flow message prefix (continued).

Flow step	Message flow	Message prefix section	Content of prefix section
4	ACK	MC	Architecture level = 1 Message type = data and response Response flag = ACK Transaction-pipe name Prefix flag = SD Send-sequence number
		SD	Length Synchronization flag = Commit Mode 0 Synchronization level = Confirm or None (Map name) Server token (Correlator) Length of server user data (Server user data)

Protecting transactions with OTMA

In a z/OS environment, transaction protection and recovery is managed by z/OS Resource Recovery Services (RRS), part of z/OS Recovery Resource Management Services (RRMS). RRS can apply coordinated changes across multiple mission-critical resources.

OTMA is one of two components that enable IMS to support protected transactions. The second component is APPC/IMS.

To process protected transactions, specify RRS=Y on the control region JCL. IMS then registers as a Resource Manager (RM) with RRMS using the CRGGRM service; it also sets its exits with the Context Services and RRS exit managers using the CRGSEIF service. IMS supports the following RRS exit routines:

- PREPARE
- COMMIT
- BACKOUT
- EXIT_FAILED
- ONLY_AGENT
- SUBORDINATE_FAILED

During initialization, IMS issues message DFS0653I to indicate that it has successfully connected with RRS and that it can now process protected transactions.

Initiating protected transactions from an OTMA client

Perform the following steps to initiate protected transactions from an OTMA client, such as IMS Connect and DB2 for z/OS stored procedures.

1. Specify Synclevel=2 (Syncpt) in the OTMA message prefix.
2. Obtain or reuse a context token. To obtain an RSS context token, use the CTXBEGC service.
3. Set the context token in the OTMA message prefix.
4. Express interest in the unit of recovery (UR) using ATREINT.
5. Send the message to the OTMA.
6. Wait for the output from the IMS transaction.
7. Send an acknowledgment (ACK) to IMS after the output is received.
8. Initiate the z/OS Resource Recovery Services (RRS) commit (ATRACMT) or backout (ATRABCK).

The OTMA client assumes the server distributed syncpoint RM (SDSRM) role. This means that the OTMA client owns the context and is the only RM allowed to initiate or invoke the RRS commit. The flow is similar to that of an APPC/IMS-protected transaction, with the following exceptions:

- The OTMA client initiates the commit using the RRS Commit_Agent_UR service (ATRACMT).
- RRS then directly informs IMS to take a commit. As a result of the previous ATRACMT call, RRS drives the commit exits of all interested RMs.

Processing protected transactions in IMS

IMS receives the protected transaction and extracts the context token from the OTMA message header. IMS saves the context token in its own control block before placing the protected transaction on the IMS message queue.

Also, before placing the protected transaction on the message queue, IMS expresses interest in the context using the Express_Context_Interest service (CTXEINT). IMS will therefore be informed if anything happens to the context while the protected transaction is queued on the message queue.

When IMS schedules the protected transaction into a dependent region, IMS switches the context token to the dependent region TCB using the Switch_Context service (CTXSWCH); it also expresses protected interest in the UR using ATREINT. IMS then presents the protected transaction to an application program that processes that particular transaction. The application contains the business logic (for example, update databases, send messages, and others). After the application completes its work, it reaches a commit point. IMS then sends the application output back to the OTMA client and waits for a commit or backout event from z/OS Resource Recovery Services.

Client/server resynchronization with OTMA

In order to guarantee that client transactions are processed and that they are processed only once, OTMA provides a protocol for synchronizing transactions.

Using a synchronized commit-then-send (Commit Mode 0) transaction pipe, the client and IMS can regain message flow in the event of a client or IMS outage. Resynchronization occurs when either IMS or the client terminates normally or abnormally.

Transaction resynchronization achieves the following:

- Prevents data from being reprocessed
- Detects that data has not been received and causes the client to resend the data
- Detects that resynchronization might not be possible
- Allows the client to decide what actions to take in order to resynchronize

OTMA resynchronization is not symmetrical, and a system's behavior depends on its role: client or IMS. Resynchronization also does not maintain symmetry for send- or receive-sequence numbers. For example, the differences for the input and output sides of an IMS flow are:

Input IMS logs the client sequence numbers when the transaction is enqueued, and from that moment, the client has no control over dequeuing the transaction.

Output

The application output is enqueued to a synchronized transaction pipe, but the output sequence numbers are not logged at that time. Only after sending the output and receiving an acknowledgment from the client does IMS finally dequeue the message and log the incremented sequence numbers.

All output using a synchronized transaction pipe is sequenced. The second output message is not sent until the ACK message from the client is received for the first output message.

Q: Why would I use a synchronized tpipe versus a nonsynchronized tpipe?

A: Use a synchronized tpipe to ensure that client transactions are processed only once in the case of a client or IMS failure. Therefore, synchronized tpipes ensure better transaction recoverability. However, in order to guarantee transaction recovery, you are required to implement resynchronization logic with synchronized tpipes.

Use a nonsynchronized tpipe when the recoverability of a transaction is less of a concern. For nonsynchronized tpipes, the client does not require the resynchronization logic.

Assumptions for OTMA resynchronization

The OTMA resynchronization process is based on several assumptions.

The assumptions include:

- Neither client nor IMS sends an ACK message until it has logged a transaction message.
- The client decides what resynchronization actions IMS should take.
- Both client and IMS can determine whether a transaction and its output messages are recoverable. The client can determine a transaction's recoverability using the architected form of the /DISPLAY TRANSACTION command.
- Recoverable OTMA messages include a value for the recoverable sequence number in the message-control information section of the message prefix. This value is incremented by 1 every time a recoverable message is sent using a tpipe (see "Summary results of IMS transactions and commands" on page 832).
- A 0 (zero) is not a valid recoverable sequence number.
- Recoverable send- and receive-sequence numbers are maintained on a per transaction pipe basis.
- IMS does not support resynchronization for any IMS command input. If the client needs to submit IMS commands using a synchronized transaction pipe, the recoverable sequence number must be set to 0 (zero). If the recoverable sequence number is not set to 0 (zero), IMS rejects the command input with sense code X'0023'.

Recoverable OTMA transactions

The recoverability of OTMA-initiated transactions and commands is determined by several factors.

The factors include:

- Is it a recoverable or unrecoverable transaction?
- Is it a recoverable or unrecoverable command?

- Is the recoverable sequence number 0 (zero) or not?
- Is it a synchronized or nonsynchronized transaction pipe?
- Is it Commit Mode 0 (commit-then-send) or Commit Mode 1 (send-then-commit)?

A recoverable IMS transaction submitted using the send-then-commit transaction flow is not rejected. However, send-then-commit transactions are discarded during IMS restart (they are unrecoverable).

Transactions that use synchronized transaction pipes are recoverable. Input messages are not recoverable for send-then-commit transactions, and requesting an ACK message has no effect on whether a transaction is recoverable. You should request ACK messages for proper synchronization of the synchronized transaction pipe.

Also, when a transaction reaches IMS, its recoverability depends on how it is defined to IMS.

Unrecoverable OTMA transactions

For unrecoverable transactions, the client must know that the transaction is unrecoverable, process it, and then forget about it.

For send-the-commit transactions, output is unrecoverable and not resynchronized. Also, send-then-commit transactions must be associated with nonsynchronized transaction pipes.

Summary results of IMS transactions and commands

The following figure summarizes the results of IMS transactions that a client submits under various processing conditions using a synchronized tpipe.

The summarization differentiates recoverable sequence numbers of zero and non-zero, and shows the differences between recoverable and unrecoverable transactions for commit modes 0 and 1 for both the zero and non-zero sequence.

In the tables, CM0 indicates a commit-then-send (commit mode 0) transaction and CM1 indicates a send-then-commit (commit mode 1) transaction.

Table 139. Results of IMS transactions using a synchronized Tpipe

Recoverable sequence number	CM0 recoverable transaction	CM0 unrecoverable transaction	CM1 recoverable transaction	CM1 unrecoverable transaction
0 (zero)	Client receives ACK message. Output is recoverable, and no input/output recoverable sequence is updated.	Client receives ACK message. Output is not recoverable and no input/output recoverable sequence number is updated.	Client receives NAK message with sense code X'001C'.	Client receives NAK message with sense code X'001C'.

Table 139. Results of IMS transactions using a synchronized Tpipe (continued)

Recoverable sequence number	CM0 recoverable transaction	CM0 unrecoverable transaction	CM1 recoverable transaction	CM1 unrecoverable transaction
Not 0 (zero)	If the recoverable sequence number is valid, client receives ACK message. If it is not valid, client receives NAK message with sense code X'001F'. Transaction and output are recoverable.	Client receives NAK message with sense code X'0023'. Client should set recoverable sequence number to 0 (zero).	Client receives NAK message with sense code X'001C'.	Client receives NAK message with sense code X'001C'.

The following table summarizes the results of IMS transactions that a client submits under various processing conditions using a nonsynchronized tpipe. The summarization differentiates recoverable sequence numbers of zero and non-zero, and shows the differences between recoverable and unrecoverable transactions for commit modes 0 and 1 for both the zero and non-zero sequence.

Table 140. Results of IMS transactions using a nonsynchronized Tpipe

Recoverable sequence number	CM0 recoverable transaction	CM0 unrecoverable transaction	CM1 recoverable transaction	CM1 unrecoverable transaction
0 (zero)	Client receives ACK message. Transaction and output are recoverable.	Client receives ACK message. Transaction and output are not recoverable.	Client receives ACK message. Transaction and output are not recoverable.	Client receives ACK message. Transaction and output are not recoverable.
Not 0 (zero)	Client receives NAK message with sense code X'0023'. Client should set recoverable sequence number to 0 (zero).	Client receives NAK message with sense code X'0023'. Client should set recoverable sequence number to 0 (zero).	Client receives NAK message with sense code X'0023'. Client should set recoverable sequence number to 0 (zero).	Client receives NAK message with sense code X'0023'. Client should set recoverable sequence number to 0 (zero).

The following table summarizes the results of commands that a client issues under various processing conditions using a synchronized tpipe or a nonsynchronized tpipe. The summarization differentiates recoverable sequence numbers of zero and non-zero, and shows the differences between commit modes 0 and 1 for both synchronized and nonsynchronized tpipes.

Table 141. Results of commands that a client issues

Recoverable sequence number	Synchronized tpipe with CM0	Synchronized tpipe with CM1	Nonsynchronized tpipe with CM0	Nonsynchronized tpipe with CM1
0 (zero)	Client receives ACK message. Command output is recoverable and output recoverable sequence number is updated.	Client receives NAK message with sense code X'001C'.	Client receives ACK message. Output is not recoverable.	Client receives ACK message. Output is not recoverable.

Table 141. Results of commands that a client issues (continued)

Recoverable sequence number	Synchronized tpipe with CM0	Synchronized tpipe with CM1	Nonsynchronized tpipe with CM0	Nonsynchronized tpipe with CM1
Not 0 (zero)	Client receives NAK message with sense code X'0023'. Client should set recoverable sequence number to 0 (zero).	Client receives NAK message with sense code X'001C'.	Client receives NAK message with sense code X'0023'. Client should set recoverable sequence number to 0 (zero).	Client receives NAK message with sense code X'0023'. Client should set recoverable sequence number to 0 (zero).

Related concepts:

 Integrity tables (Application Programming)

Related tasks:

“IMS transaction types and transaction states” on page 432

OTMA resynchronization protocol

OTMA resynchronization is based on a series of command exchanges with each client.

The command exchanges for OTMA resynchronization with a client are:

CBresynch (Client_Bid resynch)

CBresynch is sent by the client to request resynchronization with IMS after both the client and IMS have successfully joined the z/OS cross-system coupling facility (XCF) group.

SRVresynch (Server resynch)

SRVresynch must be initiated from IMS to the client after the client has successfully joined the XCF group and issued CBresynch. SRVresynch contains all synchronized tpipe names of which IMS is aware.

REQresynch (Request resynch)

REQresynch must be issued from IMS to the client for each Synchronized tpipe. REQresynch contains the tpipe name, the IMS recoverable send-sequence number for the tpipe, and the IMS recoverable receive sequence number for the tpipe.

REPresynch (Reply resynch)

REPresynch is issued from the client for each tpipe. REPresynch is a reply to the REQresynch request from IMS.

TBresynch (tpipe_Bid resynch)

Tpipe_Bid resynch is issued by the client to initiate resynchronization with IMS for a particular tpipe.

IMS keeps track of the send and receive numbers in the tpipe structure. The send and receive numbers are updated for each input and output message. When resynch occurs, both the client and IMS share their send and receive numbers to verify that both sides are synchronized. The REQresynch command from IMS releases the send and receive numbers from IMS. The client accepts the numbers and does a comparison to the client's send and receive numbers. If the send and receive numbers are not the same, then the client specifies an action to IMS with the REPresynch command. If both sides have the same send and receive numbers, then the resynch completes successfully. If resynch fails, then the failing tpipe is identified and is not used.

Command message exchange for resynchronization must follow the OTMA resynchronization protocol. Normally, the sequence of events that occurs during resynchronization is:

1. The client issues CBresynch when the client attempts to resynchronize with IMS.
2. IMS sends an ACK to acknowledge receipt of CBresynch. From this point on, IMS quiesces any non-resynch type of input or output for all synchronized tpipes. If IMS receives input while resynchronization is in progress for a synchronized transaction pipe (tpipe), the input is rejected with sense code X'0025'.
3. IMS builds the SRVresynch command and sends it to the client. The SRVresynch command lists all synchronized tpipe names of which IMS is aware for that client.
4. The client receives the SRVresynch command and issues an ACK or NAK message to IMS.
5. If IMS receives an ACK message, IMS begins the resynchronization process for each tpipe. IMS sends the REQresynch command that contains the tpipe name, the IMS recoverable send-sequence number for the tpipe, and the IMS recoverable receive-sequence number for the tpipe.
If IMS receives a NAK message from the SRVresynch command, IMS sends the DFS2393 message to the MTO and waits for a client-bid request or a CBresynch command from the client.
6. The client receives the REQresynch request. By comparing the information from the REQresynch request with its own information of the tpipe, the client sends the REPresynch reply to IMS and informs IMS about the tpipe
7. IMS receives the REPresynch reply and takes actions on the tpipe, based on the information from the client. IMS sends an ACK message to the client after it has taken actions dictated by the client. IMS enables the tpipe to handle input and output. If IMS cannot perform what the client has requested, IMS stops the tpipe and sends a NAK message to the client.
8. If more than one tpipe exists, the resynchronization process is repeated in parallel for each tpipe. Other tpipes that are not included in the SRVresynch request can send output in either direction anytime after the client receives the SRVresynch command.

The following figure illustrates the flow of nondeferred resynchronization. Following the figure is a sequential list that provides high-level flow description.




Figure 148. Flow of resynchronization (nondeferred)

1. Client-bid request with resynchronization
2. ACK message

3. SRVresynch command
4. ACK message
5. REQresynch command
6. REPresynch command
7. ACK or NAK message

If the client determines that resynchronization must be deferred for a particular tpipe, the sequence of events for that tpipe differs slightly:

In the REPresynch command, the client can set the “stop and wait for resynchronization” indicator, and can request that IMS defer any input or output while waiting for the TBresynch command from the client. Assuming steps the first four steps have completed, the events following are:

1. IMS sends the REQresynch command that contains the tpipe name, the IMS recoverable send-sequence number and the IMS recoverable receive sequence number.
2. The client receives the REQresynch request. However, due to any product-specific reasons, the client defers resynchronization for this tpipe by sending the REPresynch command with the “stop and wait for TBresynch” indicator on.
3. IMS sends an ACK message to acknowledge receipt of the REPresynch command and waits for TBresynch. Meanwhile, IMS quiesces input and output for the tpipe. If IMS receives any input while waiting for TBresynch, IMS sends a NAK message to the client with sense code X'0025'.
4. The client sends the TBresynch command and requests IMS to resume resynchronization for this tpipe.
5. IMS sends the REQresynch command that contains the tpipe name, the IMS recoverable send-sequence number, and the IMS recoverable receive-sequence number. If the associated tpipe cannot be located using the client's TBresynch command, the client receives a NAK message with sense code X'0025'.
6. The client receives the REQresynch request. By comparing the information from REQresynch request with its own information about the tpipe, the client sends the REPresynch reply to IMS and informs IMS about the tpipe.
7. IMS receives the REPresynch reply and takes actions on the tpipe, based on what the client has requested. IMS sends an ACK message to the client if it has taken actions dictated by the client. Otherwise, IMS sends a NAK message to the client with sense code X'0025' or X'0026'.

The following figure shows the flow of deferred resynchronization. Following the figure is a sequential list that provides high-level flow description.




Figure 149. Flow of resynchronization (deferred)

1. REQresynch command
2. REPresynch command with STOP AND WAIT for TBresynch
3. ACK message
4. TBresynch command
5. REQresynch command
6. REPresynch command
7. ACK or NAK message

Related reference:

"State data section" on page 867

Sample OTMA resynchronization message flow

The following figure shows the flow of messages through a synchronized transaction pipe.

Receive- and send-sequence numbers for each side (IMS and client) are represented by the letters **R** and **S**, which are set in the message-control information section of the message prefix. These numbers apply to the entire message (including multi-segment messages). **R** and **S** are not necessarily related.




Figure 150. Sample OTMA resynchronization message flow

- **1** After the client submits the transaction, IMS enqueues the transaction, and the transaction runs. The receive-sequence number is incremented by 1.
- **2** IMS sends the client an ACK message to acknowledge receiving and enqueueing the transaction.
- **3** IMS enqueues the output and sends the data to the client.
- **4** The client sends an ACK message to IMS to acknowledge receiving the output; however, IMS never receives this ACK to message 15 because of an IMS failure.

Resynchronization proceeds as follows:

- **5** The client sends a client-bid request to IMS to initiate resynchronization.
- **6** IMS sends an ACK message to the client that resynchronization will begin.
- **7** IMS sends the SRVresynch command to the client to begin resynchronization.
- **8** The client sends an ACK message to IMS to acknowledge receiving the SRVresynch command.
- **9** IMS sends the REQresynch command to the client to update the receive- and send-sequence numbers to (9,14).
- **10** The client sends the REPresynch command to IMS to update the receive- and send-sequence numbers to (15,9), and to tell IMS to dequeue the last output message. IMS dequeues message 15 and updates S to 15.

- **11** IMS sends an ACK message to the client.

Sample OTMA resynchronization messages

This topic shows sample OTMA resynchronization messages.

Client-bid request with resynchronization message

The following example shows the OTMA client-bid request with the resynchronization message.

MESSAGE CONTROL INFORMATION:

01102000 0C004040 40404040 4040A0C0	{
00000000 00000000 00000000 00000000	

STATE DATA:

0036D4D8 E7C3C6F6 40404040 40404040	..MQXCF6
40400100 00010002 00010100 00020002	
0002C4C6 E2E8C4D9 E4F00800 00007FFF	..DFSYDRU0...."
FFFF0000 00650056 C3525100 5001A051	C...&...

ACK message to acknowledge receipt of CBresynch

The following example shows the ACK message to acknowledge receipt of CBresynch.

MESSAGE CONTROL INFORMATION:

01308000 0C004040 40404040 4040A0C0	{
00000000 00000000 00000000 00000000	

STATE DATA:

0036D4D8 E7C3C6F6 40404040 40404040	..MQXCF6
40400100 00010002 00010100 00020002	
0002C4C6 E2E8C4D9 E4F00800 00007FFF	..DFSYDRU0...."
FFFF0000 00650056 C3525100 5001A051	C...&...

SRVresynch command message

The following example shows the SRVresynch command message.

MESSAGE CONTROL INFORMATION:

01102000 2C000000 00000000 0000A080	
00000000 00000000 00000000 00010000	

STATE DATA:

000AD1C2 D1F0F0F0 F0C50000 00000000	..JB0000E.....

ACK message to acknowledge receipt of SRVresynch

The following example shows the ACK message sent by client to acknowledge receipt of SRVresynch.

MESSAGE CONTROL INFORMATION:

0130A000 2C000000 00000000 0000A080	
00000000 00000000 00000000 00010000	

STATE DATA:

```
000AD1C2 D1F0F0F0 F0C50000 00000000 |..JBj0000E.....|
```

REQresynch command message

The following example shows the REQresynch command message.

MESSAGE CONTROL INFORMATION:

```
01100000 3000D1C2 D1F0F0F0 F0C5A080 |.....JBj0000E..|
00000000 00000000 00000000 00000000 |.....|
```

STATE DATA:

```
001AD1C2 D1F0F0F0 F0C50000 00020000 |..JBj0000E.....|
00020000 00000000 00000000 00000000 |.....|
```

REPresynch command message

The following example shows the REPresynch command message.

MESSAGE CONTROL INFORMATION:

```
01100000 3400D1C2 D1F0F0F0 F0C5A080 |.....JBj0000E..|
00000003 00000000 00000000 00000000 |.....|
```

STATE DATA:

```
001AD1C2 D1F0F0F0 F0C50000 00020000 |..JBj0000E.....|
00020000 00000000 00000000 00000000 |.....|
```

ACK message for successful resynchronization

The following example shows the ACK message sent by IMS to inform the client that resynchronization on a tpipe successfully completed.

MESSAGE CONTROL INFORMATION:

```
01308000 3400D1C2 D1F0F0F0 F0C5A080 |.....JBj0000E..|
00000003 00000000 00000000 00000000 |.....|
```

STATE DATA:

```
001AD1C2 D1F0F0F0 F0C50000 00020000 |..JBj0000E.....|
00020000 00000000 00000000 00000000 |.....|
```

Related reference:

Chapter 47, “OTMA message prefix,” on page 855

Managing commit-then-send output

OTMA provides several different options for managing commit-then-send (CM0) output that is not immediately returned to hold-queue capable OTMA clients. You can purge the output, reroute the output, or share the output among multiple OTMA clients by using an OTMA super member.

Purging commit-then-send asynchronous output

OTMA clients that support asynchronous commit-then-send message output can have OTMA purge output from a tpipe queue when the I/O PCB output cannot be delivered to the OTMA client.

To specify the purge function, the OTMA client must set the purge flag in the OTMA state data prefix of either the original input message or the NAK response when the output cannot be retrieved.

Purging NAKed output

For OTMA clients that are hold-queue capable, OTMA can purge commit-then-send (CM0) I/O PCB output when OTMA receives a NAK response from the client.

If a NAK message turns on the purge request flag in the OTMA state data prefix, OTMA dequeues the commit-then-send I/O PCB output.

If neither a purge request nor a reroute request is specified on a NAK message, OTMA requeues the I/O PCB output to the asynchronous hold queue of the inputting tpipe. Messages stored in the asynchronous hold queue of a tpipe can be retrieved by a RESUME TPIPE call at a later time.

If both a purge request and a reroute request are specified on a NAK message, OTMA requeues the output to the asynchronous hold queue of the inputting tpipe and issues an error message to the system console.

When a purge request is specified on an initial input message, if OTMA cannot deliver the commit-then-send I/O PCB output to the hold-queue capable client due to the z/OS cross-system coupling facility failure or a client outage, OTMA purges the output.

OTMA does not support the purging of asynchronous output from hold queues.

Rerouting commit-then-send asynchronous output

OTMA clients that support OTMA hold queues can have OTMA reroute I/O PCB output to an alternate hold queue when the output cannot be delivered to the OTMA client.

The OTMA client can specify the reroute function on the following types of messages input to OTMA:

- Input messages for commit-then-send transactions
- Input messages for send-only transactions
- NAK response to commit-then-send output

If a maximum number of tpipes has been defined for an OTMA client (tmember) and the number of tpipes associated with the tmember has reached the maximum, a new tpipe cannot be created to reroute a message. In this case, if a message cannot be rerouted because the maximum number of tpipes has been reached, OTMA deletes the message.

The maximum number of tpipes for a tmember is specified on the MAXTP parameter of the OTMA client descriptor in the DFSYDTx PROCLIB member.

Rerouting asynchronous output for send-only transactions

For send-only transactions from OTMA clients that are hold queue capable OTMA can reroute commit-then-send (CM0) messages on the I/O PCB output queue.

To specify a reroute request for asynchronous commit-then-send output in a send-only transaction message, specify the tpipe name of the alternate asynchronous hold queue to which you want the output rerouted in the user data

section of the OTMA message prefix. For IMS Connect users, see the OMUSR_REROUT_NM field documented in “OTMA user data fields used by IMS Connect” on page 283.

When reroute is specified, OTMA reroute the output to the reroute tpipe instead of the input tpipe for the send-only message.

OTMA also reroutes to the alternate tpipe any IMS DFS messages generated by the send-only transaction.

However, if a send-only transaction requests an ACK or NAK for the input message, OTMA sends back an ACK after the input transaction is enqueued. If the input send-only transaction cannot be enqueued, OTMA sends back a NAK without a DFS message.

Rerouting NAKed output

For OTMA clients that are hold queue capable, OTMA can reroute commit-then-send (CM0) messages on either the I/O PCB tpipe queue or the hold queue when OTMA receives a NAK response from the client.

If a NAK message turns on the reroute flag and specifies a reroute tpipe name in the reroute TPIPE field of the state data prefix, OTMA reroutes any commit-then-send messages on the I/O PCB output queue by requeuing it to the asynchronous output hold queue specified in the reroute TPIPE field.

If both a purge request and a reroute request are specified on a NAK message, OTMA issues an error message to the system console and requeues the output to the asynchronous hold queue of the tpipe used by the original input message.

When a reroute request is specified on an initial input message, if OTMA cannot deliver the commit-then-send I/O PCB output to the hold-queue capable client due to the z/OS cross-system coupling facility failure or a client outage, OTMA reroutes the output to the asynchronous hold queue of the inputting tpipe.

To request in a NAK message that OTMA reroute asynchronous commit-then-send output, set X'20' in byte 5 in the transaction-related state data of the OTMA header and specify the tpipe name of the alternate asynchronous hold queue to which you want the output rerouted at byte 62 of the state data.

Timeout for acknowledgments of commit-then-send output

You can specify a timeout interval for acknowledgments from the OTMA client, after which OTMA removes commit-then-send messages from the tpipe and places it on a default timeout queue or a queue that you specify.

The timeout value you specify applies to acknowledgments for both commit-then-send output and send-then-commit output; however, upon the expiration of the timeout interval, OTMA reroutes commit-then-send output instead of discarding the output as OTMA does with send-then-commit output.

When an acknowledgment for commit-then-send output expires, OTMA reroutes the output to either a specified reroute tpipe hold queue, a specified timeout tpipe hold queue, or the OTMA default timeout tpipe hold queue DFS\$\$TOQ.

Related tasks:

“Timeout interval for acknowledgments to OTMA messages” on page 794

Sharing asynchronous commit-then-send output: the OTMA super member function

Hold-queue-capable OTMA clients, such as IMS Connect, can share asynchronous commit-then-send (CM0) output messages by enabling the OTMA super member function. The OTMA super member function is specifically designed to support multiple instances of IMS Connect in a z/OS Sysplex Distributor environment.

The OTMA super member function manages all of the asynchronous CM0 output of all of its participating OTMA clients by using a common output queue. Any participating hold-queue-capable client can then retrieve the CM0 messages on the super member output queue by issuing its own RESUME TPIPE call, regardless of which client the CM0 output was originally destined for.

The messages on the super member queue are retrieved on a first-in-first-out basis, without regard to which instance of the OTMA client originated the output or which instance of the OTMA client issued the RESUME TPIPE call.

The RESUME TPIPE call from a participating OTMA client does not need to specify anything about the super member. The super member function recognizes the RESUME TPIPE call as coming from a participating OTMA client and returns the output on the common queue.

You can register the participation of an instance of IMS Connect in the super member function by specifying the super member name on the SMEMBER parameter of the IMS Connect HWS configuration statement.

To use the OTMA super member in a configuration that includes multiple IMS systems, you must have shared queues enabled for all the IMS systems. If a shared queues back-end IMS creates ALT-PCB output in the super member-enabled environment, the output can be retrieved from any front-end IMS with an OTMA client in the super member set. All retrieval options of the resume tpipe call are supported in the shared queues environment.

If there is only one IMS system with multiple OTMA clients, shared queues support is not required.

To activate the super member function specify a 1- to 4-character super member name in the SMEMBER parameter of the HWS configuration statement in the IMS Connect configuration member (HWSCFGxx). Super member names must be unique and cannot be the same as existing OTMA member names.

The first instance of IMS Connect that specifies a super member name both activates the super member function and defines the name of the super member for all other instances of IMS Connect. Any instances of IMS Connect activated subsequently that are to use the same super member queue must specify that same super member name as specified by the first instance of IMS Connect.

You can display the name of an activated super member by issuing any one of the following IMS Connect commands:

- The WTOR format command VIEWHWS
- The z/OS MODIFY format command QUERY MEMBER TYPE(IMSCON)

To display additional information about a super member, you can also issue the IMS commands /DISPLAY OTMA and /DISPLAY TMEMBER TPIPE.

When you issue the OTMA commands /START TMEMBER TPIPE, /STOP TMEMBER TPIPE, or /TRACE TMEMBER TPIPE with a regular OTMA member specified, OTMA expands the scope of the command to include any related super member.

You can also issue OTMA commands directly to an existing super member. For example, you can dequeue asynchronous messages from a super member by issuing the command /DEQUEUE TMEMBER *supermember_name* TPIPE *tpipe_name* PURGE.

Asynchronous output messages that are created by ALT-PCB processing are stored in the super member directly. If the input messages are submitted through an IMS Connect that supports the super member, the input parameter lists for the DFSYPRX0 and DFSYDRU0 user exits contain a flag indicating that the message is from a super member supported client. Also, the OTMA state data prefix pointed to by the input parameter list contains the super member name, if any.

Chapter 46. OTMA support for callout requests

IMS application programs running in IMS dependent regions can call out through the Open Transaction Manager Access (OTMA) component of IMS to servers that are outside of the IMS installation to request data or services.

An IMS configuration that supports for such callout requests requires multiple components. The components that are required differ depending on whether the callout requests are processed synchronously or asynchronously.

Callout requests from IMS application programs

OTMA, together with hold-queue-capable OTMA clients such as IMS Connect, supports callout requests from IMS application programs running in IMS dependent regions to data or service providers that are external to the IMS installation.

The external providers can include:

- An Enterprise JavaBeans (EJB) application or a message-driven bean (MDB) application running in a web application server, such as WebSphere Application Server
- A Web service
- A data source or service that is accessed through a user-written IMS Connect client

The request for services or data is referred to as a *callout request*. When an IMS application program makes a callout request, IMS can be viewed as a client in a client-server relationship where the server is the external application to which IMS is making the callout request.

Callout requests that are routed through OTMA can be processed synchronously or asynchronously.

For synchronous callout requests, an IMS application program running in an IMS dependent region issues the DL/I ICAL call and waits in the dependent region to process the response also. When the ICAL call is issued, IMS generates a correlation token for synchronous callout requests and routes the request to an OTMA destination. The correlation token is included with the callout request and must be returned to IMS with the response to route the response back to the requesting IMS application program.

For asynchronous callout requests, an IMS application program running in an IMS dependent region inserts the callout message to an ALTPCB queue and then terminates to free the dependent region. IMS does not generate a correlation token for asynchronous callout requests; consequently, if a response to the callout request is required, the correlation of any must be managed by the IMS application program. When IMS receives a response to an asynchronous callout request, IMS processes the response as a new transaction.

The configuration requirements for synchronous callout requests differ slightly from the requirements for asynchronous callout requests, but the basic components are the same:

- An IMS application that calls out to a source external to IMS for data or services. For synchronous callout requests this same application processes the response. For asynchronous callout requests, if a response is needed, the response is processed either by a different instance of the same application program or a different application program altogether.
- OTMA, which routes the callout request to the appropriate client tpipe queue based on destinations defined by using one of the following methods:
 - OTMA destination descriptor
 - Type-2 commands
 - For asynchronous callout requests, the OTMA Destination Resolution exit routine (DFSYPX0) and the OTMA User Data Formatting exit routine (DFSYDRU0)
- A hold-queue-capable OTMA client, such as IMS Connect. IMS Connect serves as both a TCP/IP gateway to IMS and an interface with OTMA.
- If you are using IMS Connect, an IMS Connect TCP/IP client, such as IMS TM Resource Adapter, IMS Enterprise Suite SOAP Gateway, or a user-written IMS Connect client.
- A data or service provider, such as an EJB application or MDB application in a Java Platform, Enterprise Edition (Java EE) environment, or a web service.

Security for both synchronous and asynchronous callout configurations can be implemented by a security product such as RACF, the OTMA Resume TPIPE Security user exit (OTMARTUX), or both.

Security for both synchronous and asynchronous callout configurations can be implemented by a security product such as RACF, the OTMA Resume TPIPE Security exit routine (DFSYRTUX), or both.

Related concepts:

“OTMA destination descriptors” on page 772

Synchronous callout requests

Synchronous callout requests are processed in real time and travel from the IMS application program running in an IMS dependent region, out to the external data or service provider, and back to the IMS application program while the IMS application program remains scheduled in the IMS dependent region.

Configuring your environment for synchronous callout can be simpler than configuring your environment for asynchronous callout, because IMS manages the correlation of the callout response to the IMS application that made the request.

Restrictions:

- OTMA does not support synchronous callout requests in shared-queues environments that are configured with front- and back-end IMS systems.
- The OTMA Input/Output Edit exit routine (DFSYIOE0) is not supported for either synchronous callout request messages received by OTMA from IMS application programs that issue the DL/I ICAL call or synchronous callout response messages.

OTMA processes synchronous callout request messages as nonrecoverable commit-then-send (CM0) output messages.

The following high-level steps provide an overview of configuring a synchronous callout environment when IMS Connect is used as the OTMA hold-queue capable client:

1. Code an OTMA destination descriptor to route the callout request to the tpipe of the IMS Connect client that is configured to retrieve the callout requests.
2. Restart IMS.
3. If you need to modify an IMS Connect user message exit routine, reassemble and bind the IMS Connect user message exit routine.
4. If you modified an IMS Connect user message exit routine, restart IMS Connect.
5. Configure an IMS Connect TCP/IP client to retrieve the callout request, pass the request to the data or service provider, and return the response to IMS Connect. For information about how to configure your IMS Connect TCP/IP client see the appropriate documentation listed below:
 - IMS TM Resource Adapter overview
 - Callout programming models
 - IMS Enterprise Suite SOAP Gateway overview
 - Configuring user-written IMS Connect clients for synchronous callout requests
6. Code an IMS application program to initiate callout requests by issuing the DL/I ICAL call and to process the response.

When callout service providers connect to IMS through an OTMA client such as IMS Connect, the synchronous callout responses are sent back to OTMA and IMS using the send-only protocol.

IMS provides various sample application programs to test callout support. For more information, see Samples for the callout function (Installation).

After you complete the initial setup of the configuration for callout requests, you can use the DL/I test program (DFSDDLTO) to verify and debug ICAL DL/I call independently of your IMS application programs.

Asynchronous callout request

IMS application programs running in IMS dependent regions can send outbound messages to request services or data from a WebSphere Application Server EJB application or a web service. The request for services or data is referred to as a callout request.

Asynchronous callout requests do not require the IMS application program that issues the request to wait for a response in the dependent region. Upon issuing an asynchronous callout request, the application program can terminate and free the dependent region. Any response to the callout request that is returned to IMS is treated as a new incoming transaction and IMS schedules a new application program instance to process it.

If an asynchronous callout request generates a response, however, the benefit gained by freeing dependent regions might be offset by the additional complexity of managing the response. For asynchronous callout responses, your installation is responsible for developing the method for correlating the response to the original request. For synchronous callout requests, IMS manages correlation for you.

IMS provides various sample application programs to test callout support. For more information, see Samples for the callout function (Installation).

Implementing the asynchronous callout function

Implementing an IMS Configuration that supports the asynchronous callout function involves planning for correlation, application programming, and OTMA routing.

Planning for correlation of asynchronous callout responses

If a callout request generates a response, the response is treated as a new transaction that runs independently from the original transaction that generated the callout request. Consequently, you must correlate the response to the appropriate IMS transaction.

Planning for correlation involves deciding what data to use for correlation and whether or not to store any correlation data in an IMS database until the callout response returns.

You also need to ensure that the tpipe name used by the IMS application in the ISRT *alternate_pcb* call and the tpipe name used by either the EJB or the IMS Enterprise Suite SOAP Gateway connection bundle for the RESUME TPIPE call both point to the same OTMA asynchronous hold queue either by using the same name for each or by coding the appropriate logic in the OTMA destination routing exit routines. The tpipe name is based on the alternate client ID provided by the external application through IMS Connect. The application program issuing the callout request must specify the tpipe name of the asynchronous hold queue in an ISRT *alternate_pcb* call.

Application programming for asynchronous callout requests

An application program issuing an asynchronous callout request must include in the callout request message any data required by the external application it is calling.

If the correlation of a response from the external application is required, the calling application program must also capture any data required for correlation purposes. The application program must also include the tpipe name of the dedicated OTMA asynchronous hold queue for the callout request function.

If any correlation data is included in the callout request message, you must ensure that the external application program tolerates the data and sends it back with the response.

Optionally, the application program might also store data in a database for correlation purposes or for the results of any IMS processing that must be returned with the callout response.

An application program processing a response to a callout request can process any correlation data returned by the external application program, retrieve any correlation data or IMS data stored in the callout database, and return the final data to the appropriate queue, terminal, or user as appropriate.

If you are using IMS TM Resource Adapter, a Java application that processes a callout request externally can be configured to “listen” to the OTMA asynchronous hold queue by issuing a looping RESUME TPIPE call with an appropriate wait time. The Java application program might also be designed to preserve the correlation data and send it back with the callout response.

Related concepts:

“IMS application programs and the asynchronous callout function” on page 851

OTMA routing of asynchronous callout requests

OTMA routing of asynchronous callout requests is required to route the ISRT alternate_pcb call to the appropriate tpipe.

The OTMA routing is performed by the OTMA destination descriptor or the OTMA Destination Resolution exit routine (DFSYPX0) and the OTMA User Data Formatting exit routine (DFSYDRU0).

How the OTMA routing of the callout response is handled depends on your installation. The callout response can be routed back to a user or terminal or to an OTMA asynchronous hold queue.

Related concepts:

“OTMA destination descriptors” on page 772

Related reference:

 OTMA Destination Resolution exit routine (DFSYPX0) (Exit Routines)

 OTMA User Data Formatting exit routine (DFSYDRU0) (Exit Routines)

Initiating asynchronous callout requests from an IMS Connect TCP/IP client

To initiate an asynchronous callout request from an IMS Connect TCP/IP client, the TCP/IP client must submit a Send Only request that invokes an IMS application program that then issues the callout request.

To wait for and retrieve the response to the callout request, the TCP/IP client then must issue a Resume Tpipe call.

IMS TM Resource Adapter and asynchronous callout requests

If you are using the asynchronous callout function with the IMS TM Resource Adapter, your WebSphere Application Server administrator must configure a shareable connection factory to be used by the EJB bean or message-driven bean to retrieve the callout requests from IMS through IMS Connect.

The bean uses the IMS TM Resource Adapter to “listen” to the OTMA asynchronous hold queue by issuing a looping RESUME TPIPE call with a long timeout value. If no messages are on the queue when a RESUME TPIPE call is issued, the IMS TM Resource Adapter remains in a wait state. If the RESUME TPIPE call times out before a callout request arrives on the asynchronous hold queue, control is returned to the bean, which can then issue another RESUME TPIPE call to continue listening or break the loop to perform error checking.

Also, your Java application programmers must modify the bean to specify an alternate client ID that identifies the OTMA tpipe on which the callout request is queued.

Optionally, you can have your bean process the callout request itself, or you can have the bean pass the callout request on to a web service.

If you are using the IMS TM Resource Adapter with a z/OS Sysplex Distributor that distributes input to IMS through multiple instances of IMS Connect, the callout requests must be made using a TMEMBER for which the OTMA super member function is enabled.

For information about coding the IMS TM Resource Adapter and EJB applications, including sample application code, see IMS TM Resource Adapter overview and Callout programming models.

SOAP Gateway and asynchronous callout requests

IMS Enterprise Suite SOAP Gateway includes support for the asynchronous callout function, which enables IMS application programs to make asynchronous callout requests to web services through the SOAP Gateway.

When enabled for asynchronous callout requests, the SOAP Gateway listens for new callout requests by continuously issuing the RESUME TPIPE call to IMS Connect. If a callout request message is on the tpipe queue, OTMA sends the callout request to IMS Connect, IMS Connect processes the message, converting the message to XML if necessary, and then sends the message to the SOAP Gateway.

If you are using the SOAP Gateway with a z/OS Sysplex Distributor that distributes input to IMS through multiple instances of IMS Connect, the callout requests must be made using a TMEMBER for which the OTMA super member function is enabled.

If the callout request message requires conversion to XML, XML conversion support must be enabled in IMS Connect and the callout request message must specify the XML adapter and the converter in its OTMA header of the callout request message. If the callout request message does not specify an XML adapter and a converter, IMS Connect does not convert the message to XML.

You can specify an XML adapter and converter for callout request messages by using the OTMA destination descriptor. For example:

```
D DESCNAME TYPE=IMSCON TMEMBER=HWS1 TPIPE=PIPENAME  
D DESCNAME ADAPTER=HWSXMLA0 CONVRTR=CNVNAMED
```

In the example, if an IMS application program inserts a message to the ALT IOPCB using the descriptor name DESCNAME, OTMA adds the adapter and converter names to the OTMA header of the callout request message and then queues the message to the tpipe PIPENAME.

IMS Connect support for XML conversion with the asynchronous callout function requires a callout XML converter for each web service that you want your IMS application program to call. Generate the callout XML converters by using Rational Developer for System z Version 7.1.1 or later.

For more information about the SOAP Gateway, see the SOAP Gateway documentation in the IBM Knowledge Center.

Related concepts:

“Overview of IMS Connect XML Conversion Support” on page 178

“OTMA destination descriptors” on page 772

IMS application programs and the asynchronous callout function

You might need to modify or create IMS application programs to support the callout function if you need IMS to process a response to the callout request, a terminal requires an acknowledgment after submitting a transaction, or a response-mode transaction initiates a program-to-program switch.

If you do not need a response to the callout request, the application program needs to specify only the appropriate OTMA destination name in an insert call to an alternate PCB. No modification of the application program is necessary.

If, however, you need a response to the callout request, you must modify the IMS application programs to correlate the response to the original input transaction. The method of correlation is ultimately up to you and your installation, but it will very likely require the application program to include in some type of data to be used for correlation purposes in the outgoing callout request message and then look for that same data in the response message.

You might also have the IMS application program that issues the callout request store correlation data or other data in an IMS database for retrieval when the callout response is returned. In a shared queues environment, such a database is also useful if the response transaction might be scheduled on a different IMS system. If a reply must be sent back to either an LTERM or an IMS Connect TCP/IP client, the LTERM name or TMEMBER and TPIPE name should also be stored in the database.

The following code samples shows an example of a COBOL copybook data structures that might be used for the callout function and the callout response.

The callout request data structure includes a correlator field, CORRID, to send correlation data to the EJB:

```
01 CALLREQ
   02 LL          PIC S9(4) COMP.
   02 ZZ          PIC S9(4) COMP.
   02 CORRID      PIC (8).
   02 ACCTNUM     PIC (20).
```

When the response is returned, the CORRID field of the COBOL copybook data structure contains the same correlation data, which the IMS application can use to relate the response to the initial request, if needed:

```
01 CALLRESP
   02 LL          PIC S9(4) COMP.
   02 ZZ          PIC S9(4) COMP.
   02 CORRID      PIC (8).
   02 ACCTBAL     PIC (20).
```

Avoiding hung application programs and terminals with asynchronous callout requests

Generally, terminals do not implement any timeout mechanisms and if a response is not returned to the terminal, the terminal hangs.

A callout request could result in a hung terminal in at least two possible cases: the application program issuing a callout request terminates without responding to the terminal or a response-mode transaction performs a program-to-program switch without adequate correlation of the response to either the callout request or the inputting terminal.

The first case might occur if, for example, a terminal waits for a response after submitting a transaction that initiates a callout request, and the IMS application program processing that transaction terminates after inserting the callout request to the ALTPCB, the terminal might hang waiting for a response. To avoid this, IMS application programs that issue a callout request should be modified to send an acknowledgment to the inputting terminal.

In the second case, a terminal might hang in the following scenario: an application program, PGMA, processes a response-mode transaction that waits for a response after initiating a program-to-program switch to application program, PGMB, which in turn issues a callout request that generates a response to IMS. Then, if a third application program, PGMC, processes the response, but does not notify PGMA, PGMA hangs waiting for a response. To avoid this, PGMA should be modified to support the correlation of the response through PGMB and PGMC, so that PGMC can pass the response to PGMA for return to the initiating client terminal. If a callout request will not generate a response, response-mode transactions should not be used with applications program such as PGMA.

Correlating responses to asynchronous callout requests

If the external application or web service returns a response message to IMS as a result of a callout request, the IMS application programs processing the request and response will need to correlate the response message to the callout request for proper processing of the response message.

You might also need to correlate the response message to the terminal or TCP/IP client application that originated the callout request.

Correlating a response to an asynchronous callout request:

The IMS application programs are responsible for correlating the response message to the initial asynchronous callout request. An IMS application can define a data element in the callout request, such as a message identifier or a unique callout request ID, that can be used to correlate a response with the initial input message.

After a response has been correlated to its initiating callout request, the response can then be correlated to the original inputting terminal or IMS Connect TCP/IP client.

Routing asynchronous callout request responses to the inputting terminal:

To route a response message to a terminal, the IMS application that issues the asynchronous callout request can capture the LTERM name of the terminal.

After it receives the response message, the IMS application that processes the response can then identify the originating terminal by the LTERM name. The LTERM name can either be included as part of the data of the callout request and the response messages or the LTERM name can be saved in a temporary database that is accessible to the IMS application that processes the response.

If a temporary database is used, the data used for the correlation of the response to the callout request must be included with the LTERM name in the database. After the response is returned, the IMS application can return the response to the initiating terminal by making a CHNG call and ISRT the response to the LTERM name captured by the IMS application program that issued initial callout request.

Routing asynchronous callout responses to an IMS Connect TCP/IP client:

To route the response to the TCP/IP client, the IMS application program that issues the asynchronous callout request message must capture the IMS Connect TMEMBER and OTMA TPIPE names used by the initiating TCP/IP client.

Upon receipt of the response to the callout request, the IMS application program that processes the response uses the TMEMBER and TPIPE names to return the response to the correct TCP/IP client. To successfully route the response to the correct TMEMBER, you will also need to use either the OTMA destination descriptor or the OTMA DFSYPRX0 and DFSYDRU0 exit routines.

To make the TMEMBER and TPIPE names available to the IMS application program that processes the response, the names can be included as part of the data of the callout request and response messages or the names can be saved in a temporary database that is accessible to the IMS application program that processes the response.

If a temporary database is used, the IMS application program processing the response must use further correlation to retrieve the correct TMEMBER and TPIPE names when the response transaction is returned. When correlation is complete, the IMS application programs sends the response back to the original IMS Connect TCP/IP client by making a CHNG call and an ISRT using the TMEMBER and TPIPE names.

Chapter 47. OTMA message prefix

OTMA messages have a prefix that conforms to a format that is mapped by the DFSYMSG macro in the IMS.ADFSMAC data set.

The maximum length for a message prefix is 4096 bytes; this length does not include the application data.

The OTMA message prefix can contain the following sections:

1. Message control information section
2. State data section
3. Security data section
4. User data section
5. Application data section

The following table shows the sections of the OTMA message prefix and lists some of the key fields within each section. In this table, the term “Server” refers to IMS.

Table 142. OTMA message prefix segments and their key fields

Message control information	State data	Security data	User data	Application data
Tpipe name	Destination override	User ID		
Message type	Map name	Utoken		
Sequence numbers	Sync flags	Security flags		
Processing flag	Sync_level			
Response indicator	Commit mode tokens			
Chaining indicator	Server state			

Message-control information section

For every OTMA message, you must provide message-control information in the first section of the OTMA message prefix.

The message-control information section is mapped by the TMAMCTL DSECT of the DFSYMSG macro.

The following table is a summary of the content of the message-control information section of the message prefix (column 1 from the table in Chapter 47, “OTMA message prefix”). The summary includes byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 143. Format of the message-control information section of the OTMA message prefix.

Byte	Length	Content	Value	Meaning
0	1	Architecture level	X'01'	OTMA release level. This is mandatory for all messages.

Table 143. Format of the message-control information section of the OTMA message prefix (continued).

Byte	Length	Content	Value	Meaning
1	1	Message type: A specification in this field is mandatory for all messages.	X'80'	<i>Data:</i> Server output data (output from an IMS application program). This data is not a transaction.
			X'40'	<i>Transaction:</i> Transaction or IMS command input to the server. The actual transaction name is specified in the application-data section of the prefix.
			X'20'	<i>Response:</i> A response message.
			X'10'	<i>Protocol command:</i> An OTMA protocol command (not an IMS command). See "Protocol command type" below at byte offset 4.
			X'08'	<i>Commit Confirmation:</i> Commit complete. Used by the server to notify the client of sync point completion. Only used for send-then-commit transactions. See "commit-confirmation flag" below.
2	1	Response flag: Values in the response flag field are mutually exclusive.	X'80'	<i>ACK:</i> Positive acknowledgment.
			X'40'	<i>NAK:</i> Negative acknowledgment. A NAK can be accompanied by an additional sense code.
			X'20'	<i>Response Requested:</i> A response is requested for this message.
			X'10'	<i>Extended Response Requested:</i> Requests architected transaction or command attributes to be returned to the client.
			X'08'	<i>Response to a synchronous callout request:</i> This message is a response to a synchronous callout request that was issued by an IMS application program.
			X'04'	<i>The input transaction expired before GU call:</i> Set by the server, this flag indicates that the input transaction timed out on the IMS input queue before an IMS application could issue a GU call to retrieve it.
			X'02'	<i>Support for delayed ACK or NAK response:</i> The OTMA client sets this flag to instruct OTMA to send a NAK to the client if OTMA receives a late or invalid ACK/NAK from the client.
			X'01'	<i>Return input message on transaction expiration:</i> Set by the client to request that, if a transaction times out before processing by IMS, OTMA returns the original input message to the client instead of message DFS3588I.
3	1	Commit-confirmation flag	X'80'	<i>Committed:</i> Server committed successfully.
			X'40'	<i>Aborted:</i> Server aborted commit.
			X'08'	<i>Aborted due to the OTMA timeout condition:</i> Indicates that an ACK or NAK response from the OTMA client was not received before the timeout limit was reached.

Table 143. Format of the message-control information section of the OTMA message prefix (continued).

Byte	Length	Content	Value	Meaning
4	1	Protocol command type	X'04'	<i>Client-Bid</i> : Sent by a client to the server. The response-requested flag and the appropriate state data fields (for example, Member Name) must also be set.
			X'08'	<i>Server Available</i> : Sent by the server to a client. The appropriate state data fields (for example, Member Name) must also be set.
			X'0C'	<i>CBresynch</i> : Sent by a client to the server to request a resynchronization. This client-bid request with resynchronization to follow is optional, and causes the server to send an SRVresynch command to the client.
			X'14'	<i>Suspend Processing for All Tpipes</i> : The server sends this command to suspend all message activity with the client.
			X'18'	<i>Resume Processing for All Tpipes</i> : The server sends this command to resume message processing with the client.
			X'1C'	<i>Suspend Input for Tpipe</i> : The server sends this command when it is overloaded.
			X'20'	<i>Resume Input for Tpipe</i> : The server sends this command when it is ready for client input (following a Suspend Input for tpipe command).
			X'24'	<i>Resume output for tpipe</i> : Sent by a client to the server to request queued tpipe output be resent.
			X'28'	<i>Resume output for the hold queue for tpipe</i> : Sent by a client to the server to request messages from the hold queue for tpipe.
			X'29'	<i>Cancel resume tpipe request</i> : Sent by a client to the server to cancel a pending resume tpipe request. The resume tpipe request must be identified by specifying its resume tpipe token in byte 4 of the OTMA state data.
			X'2A'	<i>No messages on tpipe hold queue</i> : Sent by the server to the client when a resume tpipe request is received and the tpipe hold queue is empty.
			X'2C'	<i>SRVresynch</i> : Sent by the server to a client who has sent a CBresynch. This command identifies all synchronized tpipes within the server.
			X'30'	<i>REQresynch</i> : Sent by the server to a client to specify the state of a synchronized tpipe.
			X'34'	<i>REPresynch</i> : Sent by a client to the server to indicate the type of resynchronization to be performed by the server.
			X'38'	<i>TBresynch</i> : Sent by a client to the server to initiate resynchronization for a particular tpipe.
			X'3C'	<i>Server state</i> : Sent by the server to the client with the current status of server processing. The client can use this information to redirect transactions to a different server if server processing slows.

Table 143. Format of the message-control information section of the OTMA message prefix (continued).

Byte	Length	Content	Value	Meaning
5	1	Processing flag	X'80'	<i>Resume Tpipe token:</i> For resume tpipe requests only, indicates that a resume tpipe request includes a token that uniquely identifies the request. OTMA uses the token to queue and process multiple resume tpipe requests in order.
				The Resume Tpipe token is also specified by the client when the client cancels a resume tpipe request.
			X'40'	<i>Synchronized Tpipe:</i> Input and output sequence numbers are maintained for the transaction pipe.
			X'20'	<i>Asynchronous output:</i> The server is sending unsolicited queued data messages.
			X'10'	<i>Error Message Follows:</i> An error data message follows. Set by the server when sending a NAK.
			X'08'	<i>Message in Hold Queue:</i> One or more messages in the hold queue for tpipe. Not available for shared queues.
			X'01'	<i>Error Message Sent:</i> An error message was sent as a CM1 response. A CM1 transaction generated an error message as the only response.
6		Tpipe name		OTMA identification and processing control token. This name is used to override the LTERM name on the I/O PCB for an IMS application program.
14	8	Chain flag This flag is mandatory for multi-segment messages.	X'80'	<i>First-In-Chain:</i> The first segment of a multi-segment message. A message of only one segment is indicated by setting both the first-in-chain and last-in-chain flags.
			X'40'	<i>Middle-In-Chain:</i> Part of a multi-segment message.
			X'20'	<i>Last-In-Chain:</i> The last segment of a multi-segment message.
			X'10'	<i>Discard Chain:</i> Discard the current chain of message segments.
15	1	Prefix flag The value of this field indicates which sections of the message prefix are attached to this message.	X'80'	<i>State Data:</i> The state-data section is included with the message. State data section is mandatory for each message.
			X'40'	<i>Security:</i> The security section is included with the message.
			X'20'	<i>User Data:</i> The user-data section is included with the message. This data is specified by an OTMA client.
			X'10'	<i>Application Data:</i> The application-data section is included with the message.
16	4	Send-sequence number		The sequence number for the transaction pipe. Incremented on every send for each transaction pipe.
20	4	Sense code		
		ORG to byte offset 20 (TMAMCSNS)		
20	2	Sense code		Accompanies a NAK message.
22	2	Reason code		Accompanies a NAK message.
		ORG to byte offset 20 (TMAMCSNS)		

Table 143. Format of the message-control information section of the OTMA message prefix (continued).

Byte	Length	Content	Value	Meaning
20	4	Aging value		Specifies the accessor environment element (ACEE) aging value, in seconds. IMS creates an ACEE if the age of the current ACEE is greater than this value. The aging value specifies how often the cached user ID ACEE should be refreshed. The aging value flag, X'40', must also be set in byte 5 of the State Data. This value does not apply to an ACK or NAK message. The minimum value for caching support is 300 seconds (5 minutes). If the aging value specified is less than the minimum, IMS always creates a non-cached ACEE.
24	4	Recoverable sequence number		The recoverable sequence number for the transaction pipe. Incremented on every send of a recoverable message using a synchronized transaction pipe. Required for resynchronization only.
28	2	Segment sequence number		Sequence number for segments of a multi-segment OTMA message.
30	1	CM1 ACK Timeout	1 to 255	Specifies a message level timeout value in seconds for this message.
31	1	Reserved		Applies only to send-then-commit messages when synclevel=confirm or synclevel=syncpt. Sets a message level timeout value for this transaction only. The value cannot be greater than the timeout value set in OTMA; otherwise, the value is ignored. The time out value is set in OTMA either by the /START TMEMBER command or the OTMA client descriptor.
ORG to byte offset 24 (TMAMCRSQ)				
24	8	Resume tpipe requester ID		

Related reference:

“Explanation of OTMA message-control information fields”

Explanation of OTMA message-control information fields

The fields in the message-control information portion of the OTMA message prefix are used to specify the characteristics of an OTMA message, such as its type, its contents, its processing requirements, and so forth.

This topic provides explanations for the fields in the message-control information section of the message prefix.

Architecture level (TMAMCALV)

Specifies the OTMA architecture level. The client specifies an architecture level, and the server indicates in the response message which architecture level it is using. The architecture levels used by a client and a server must match.

Mandatory for all messages.

Message type (TMAMCMGT)

Specifies the message type. Every OTMA message must specify a value for the message type. The values are not mutually exclusive. For example,

when the server sends an ACK message to a client-submitted transaction, both the transaction and response flags are set.

Data (TMAMCDTA - X'80')

Specifies server output data sent to the client. If the client specifies synchronization level Confirm in the state-data section of the message prefix, the server also sets Response Requested for the response flag. If the client does not specify a synchronization level, the server uses the default, Confirm.

Transaction (TMAMCTXN - X'40')

Specifies client input data to the server.

Whether the server replies with an ACK or NAK message depends only on whether Response Requested is also set for the response flag.

Response (TMAMCRSP - X'20')

Specifies the message type as response message, and is set when the message response flag specifies Response Requested.

If this flag is set, the response flag specifies either ACK or NAK.

The send-sequence numbers must match for the original data message and the response message. Chained transaction input messages to the server must always request a response before the next transaction (for a particular transaction pipe) is sent.

Command (TMAMCCMD - X'10')

Specifies an OTMA protocol command. OTMA commands must always specify Response Requested for the Response flag.

Commit confirmation (TMAMCCMT - X'08')

Specifies that commit is complete. This is sent by the server when a sync point has completed, and is only applicable for send-then-commit transactions. This flag can also be set by an OTMA client to inform the OTMA server to end the IMS conversational transaction.

Response flag (TMAMCRSI)

Specifies either that the message is a response message or that a response is requested.

Acknowledgments to transactions include attributes (for that transaction) in the application-data section of the message prefix only if the transaction specifies Extended Response Requested.

ACK (TMAMCACK - X'80')

Specifies a positive acknowledgment.

NAK (TMAMCNAK - X'40')

Specifies a negative acknowledgment.

See the sense code field for more information on the reason for the NAK.

Response requested (TMAMCRRQ - X'20')

Specifies that a response is requested for this message. This can be set for message types of Data, Transaction, or Command.

When sending send-then-commit IMS command output, IMS does not request an ACK regardless of the synchronization level.

Extended response requested (TMAMCERQ - X'10')

Specifies that an extended response is requested for this message.
Can be set by a client only for transactions (or for transactions that specify an IMS command instead of a transaction code).

If this flag is set for a transaction, IMS returns the architected attributes for that transaction in the application-data section of the ACK message.

If this flag is set for a command, IMS returns the architected attributes in the application-data section of the ACK message. This flag can be set for the IMS commands /DISPLAY TRANSACTION and /DISPLAY TRANSACTION ALL.

Response to a synchronous callout request (TMAMSYRP - X'08')

Specifies that this message is a response to a synchronous callout message issued by an IMS application program running in an IMS dependent region.

When the flag for a synchronous callout response is set, most hold-queue capable clients, such as IMS Connect, also set the send-only message flag (TMAMHSOM - X'80') in the client flag field (TMAMHCFL) of the state data section of the message prefix.

Support for delayed ACK or NAK response (TMAMDACK - X'02')

Specifies that the client asks OTMA to return a NAK response to the client if OTMA receives a late or invalid acknowledgement from the client. For example, OTMA returns a NAK to the client when OTMA receives an acknowledgement after either the ACK timeout interval expires or a /STO TMEMBER TPIPE command has been issued clear the WAIT status of the tpipe. The NAK returned to the client by OTMA includes an X'2B' sense code.

Commit-confirmation flag (TMAMCCCI)

Specifies the success of a commit request. Sent by the server to the client in a commit-confirmation message. These messages are only applicable for send-then-commit transactions, and are not affected by the synchronization-level flag in the state-data section of the message prefix.

Committed (TMAMCCTD - X'80')

Specifies that the server committed successfully.

Aborted (TMAMCABT -X'40')

Specifies that the server aborted the commit.

IMS ready to commit (TMAMCCTD - X'20')

Specifies that the server is ready to commit pending a commit notification from the client.

Aborted due to the OTMA timeout condition (TMAMCTMO -X'08')

Indicates that an ACK or NAK response from the OTMA client was not received before the timeout limit was reached.

Command type (TMAMCTYP)

Specifies the OTMA protocol command type.

IMS MTO commands are specified in the application-data section of the message.

Client-bid (TMAMCBID - X'04')

Specifies the first message a client sends to the OTMA server. This command must also set the response-requested flag and the

security flag in the message-control information section of the message prefix. The appropriate state-data fields (for example, Member Name) must also be set.

The security-data prefix must specify a Utoken field so the OTMA server can validate the client's authority to act as an OTMA client.

Because the server can respond to the client-bid request, this message should not be sent until the client is ready to start accepting data messages.

Server available (TMAMCAVL - X'08')

Specifies the first message the server sends to a client. It is sent when the server has connected to the z/OS cross-system coupling facility (XCF) group before the client has connected. The client replies to the Server Available message with a client-bid request. The appropriate state data fields (for example, Member Name) must also be set.

If the client connects first, it is notified by XCF when the server connects, and begins processing with a client-bid request.

CBresynch (TMAMCRSN - X'0C')

Specifies a client-bid message with a request by the client for resynchronization. This command is optional and causes the server to send an SRVresynch message to the client. The CBresynch command is the first message that a client sends to the OTMA server when it attempts to resynchronize with IMS and existing synchronized tpipes exist for the client. Other than the CBresynch message indicator in the message prefix, the information required for the message prefix should be identical to the client-bid command.

If IMS receives a client-bid request from the client and IMS is aware of existing synchronized tpipes, IMS issues informational message DFS2394I to the MTO. IMS resets the recoverable send- or receive-sequence numbers to 0 (zero) for all the synchronized tpipes.

Suspend processing for all tpipes (TMAMCSPA - X'14')

Specifies that the server is suspending all message activity with the client. All subsequent data input receives a NAK message from the server. Similarly, the client should send a NAK message for any subsequent server messages.

Clients can suspend processing for a particular transaction pipe by submitting a /STOP TPIPE command as an OTMA message.

Resume processing for all tpipes (TMAMCRSA - X'18')

Specifies that the server is resuming message activity with the client.

Clients can resume processing for a particular transaction pipe that has been stopped by submitting a /START TPIPE command as an OTMA message.

Suspend input for tpipe (TMAMCSPN - X'1C')

Specifies that the server is overloaded and is temporarily suspending input for the transaction pipe. All subsequent client input receives NAK messages for the transaction pipe specified in

the message-control information section of the message prefix. A response is not requested for this command.

This command is also sent by IMS when the master terminal operator enters a /STOP TPIPE command.

Resume input for tpipe (TMAMCRSM - X'20')

Specifies that the server is ready to resume client input following an earlier Suspend Input for tpipe command. A response is not requested for this command.

This command is also sent by IMS when the IMS master terminal operator issues a /START TPIPE command.

Resume output for tpipe (TMAMCRTP - X'24')

Specifies one or multiple tpipe names to the OTMA server. All queued output on the tpipes will be resent again.

Resume output for all tpipes (TMAMCRAT - X'26')

Sent by client to request that OTMA resume sending output messages from all tpipes associated with the client. Does not apply to tpipe hold queues.

Resume output for the hold queue for tpipe (TMAMCRHQ - X'28')

Specifies that a client is requesting to retrieve messages from the hold queue for tpipe. There are command options to retrieve messages.

Cancel resume tpipe request (TMAMCDRH - X'29')

Sent by a client to the server to cancel a pending resume tpipe request. The resume tpipe request must be identified by specifying its resume tpipe token in byte 4 of the OTMA state data.

No messages on tpipe hold queue (TMAMCMMSG - X'2A')

Sent by the server to the client when a resume tpipe request is received and the tpipe hold queue is empty.

SRVresynch (TMAMCSRS - X'2C')

Specifies the server's response to a client's CBresynch command. This command specifies the states of synchronized transaction pipes within the server (the send- and receive-sequence numbers).

This command is sent as a single message (with single or multiple segments), and an ACK is requested.

REQresynch (TMAMCRQS - X'30')

Specifies the send-sequence number and the receive sequence for a particular tpipe. REQresynch is sent from IMS to a client.

REPresynch (TMAMCRPS - X'34')

Specifies the client's desired state information for a tpipe. A client sends the REPresynch command to IMS in response to the REQresynch command received from IMS.

TBresynch (TMAMCTBR - X'38')

Specifies that the client is ready to receive the REQresynch command from IMS.

Resource state (TMAMMNTR - X'3C')

Sent by server to notify the client of the state of OTMA resources. The client can use this information to redirect transactions to a different server if server processing slows.

|
|
|
|

Processing flag (TMAMCPFG)

Specifies options by which a client or server can control message processing.

Resume tpipe token (TMAMQRTP - X'80')

For resume tpipe requests only, indicates that a resume tpipe request includes a token that uniquely identifies the request. OTMA uses the token to queue and process multiple resume tpipe requests in order.

The Resume Tpipe token is also specified by the client when the client cancels a resume tpipe request.

Synchronized tpipe (TMAMCSYP - X'40')

Specifies that the transaction pipe is to be synchronized. Allows the client to resynchronize a transaction pipe if there is a failure. Only valid for commit-then-send transactions.

This flag causes input and output sequence numbers to be maintained for the transaction pipe. All transactions routed through the transaction pipe must specify this flag consistently (either on or off).

Asynchronous output (TMAMCASY - X'20')

Specifies that the server is sending unsolicited queued output to the client. This can occur when IMS inserts a message to an alternate PCB.

Certain IMS commands, when submitted as commit-then-send, can cause IMS to send the output to a client with this flag set. In this case, the OTMA prefixes contain no identifying information that the client can use to correlate the output to the originating command message. These command output data messages simply identify the transaction-pipe name. IMS can also send some unsolicited error messages with only the transaction-pipe name.

Error message follows (TMAMCERR - X'10')

Specifies that an error message follows this message. This flag is set for NAK messages from the server. An additional error message is then sent to the client.

The asynchronous-output flag is not set in the error data message, because the output is not generated by an IMS application.

Message in the hold queue (TMAMCQUE - X'08')

Specifies that one or more messages exist in the hold queue for the tpipe to be delivered. This flag is always on for an IMS output message that has been sent from the hold queue for tpipe. Therefore, this flag can be used to determine whether there is any message in the hold queue for an IMS output message that has been sent from the regular queue for tpipe.

To determine whether the IMS output message is sent from the regular queue or from the hold queue, check the "From Hold Queue" flag in the Server State of the State Data.

To retrieve one or more messages from the hold queue, issue the "Resume Output for the Hold Queue for tpipe" protocol command.

Error message sent (TMAMCER3 - X'01')

An error was detected while processing a CM1 transaction and the

application did not issue an ISRT to the IOPCB. Therefore, the DFS message is the only response for this CM1 transaction. The OTMA client, for example IMS Connect or WebSphere MQ, must deliver the DFS message to the client application even if OTMA sets the deallocate abort flag in the commit confirmation message.

Tpipe name (TMAMCTNM)

An 8-byte character field that specifies a transaction-pipe name. For IMS, this name is used to override the LTERM name on the I/O PCB. This field is applicable for all transaction, data, and commit-confirmation message types. It is also applicable for certain response and command message types.

Chain flag (TMAMCCHN)

Specifies how many segments are in the message. This flag is applicable to transaction and data message types, and it is mandatory for multi-segment messages.

First-in-chain (TMAMCFIC - X'80')

Specifies the first segment in a chain of segments which comprise a multi-segment message. Subsequent segments of the message only need the message-control information section of the message prefix. Other applicable prefix segments (for example, those specified by the client on the transaction message) are sent only with the first segment (with the first-in-chain flag set).

If the OTMA message has only one segment, the last-in-chain flag should also be set.

Middle-in-chain (TMAMCMIC - X'40')

Specifies a segment that is neither first nor last in a chain of segments that comprise a multi-segment message. These segments only need the message-control information section of the message prefix.

Restriction: Because the client and server tokens are in the state-data section of the message prefix, they cannot be used to correlate and combine segmented messages. The transaction-pipe name and send-sequence numbers can be used for this purpose; they are in the message-control information section of the message prefix for each segment.

Last-in-chain (TMAMCLIC - X'20')

Specifies the last segment of a multi-segment message.

Discard chain (TMAMCCAN - X'10')

Specifies that the entire chain of a multi-segment message is to be discarded. The last-in-chain flag must also be set.

Prefix flag (TMAMCPFL)

Specifies the sections of the message prefix that are attached to the OTMA message. Every message must have the message-control information and state-data sections, but any combination of other sections can be sent with an OTMA message.

State data (TMAMCSTD - X'80')

Specifies that the message includes the state-data section of the message prefix.

Security data (TMAMCSEC - X'40')

Specifies that the message includes the security-data section of the message prefix.

User data (TMAMCUSR - X'20')

Specifies that the message includes the user-data section of the message prefix.

Application data (TMAMCAPP - X'10')

Specifies that the message includes the application-data section of the message prefix.

Send-sequence number (TMAMCSSN)

Specifies the sequence number for a transaction pipe. This sequence number is updated by the client and server when sending messages or transactions.

Recommendation: Increment the number separately for each transaction pipe.

This number can also be used to match an ACK or NAK message with the specific message being acknowledged.

Sense code (TMAMCSNS)

Specifies a 4-byte sense code that accompanies a NAK message. TMAMCSNS has two parts, a 2-byte sense code (TMAMCSNC) and a 2-byte reason code (TMAMCRSC). For an explanation of the codes returned in this field, see *IMS Version 11 Messages and Codes, Volume 2: Non-DFS Messages*.

Sense code (TMAMCSNC)

Specifies a 2-byte sense code that accompanies a NAK message. For an explanation of the codes returned in this field, see *IMS Version 11 Messages and Codes, Volume 2: Non-DFS Messages*.

Reason code (TMAMCRSC)

Specifies the 2-byte reason code that accompanies a NAK message. This code can further qualify a sense code.

Userid aging value (TMAMAGNG)

Specifies a 4-byte aging value in seconds for the input user ID. This field is different from the aging value specified in the OTMA client-bid command for OTMA connection. The aging value in the client-bid command sets the default aging value for all the OTMA user IDs; however, the Userid Aging Value overrides the default aging value for a specific user ID. If the Userid Aging Value is less than 300 seconds (5 minutes), IMS always creates a non-cached accessor environment element (ACEE).

Recoverable sequence number (TMAMCRSQ)

Specifies the recoverable sequence number for a transaction pipe. Incremented on every send of a recoverable message using a synchronized transaction pipe. Both the client and the server increment their recoverable send-sequence numbers and maintain them separately from the send-sequence number. Required for resynchronization only.

Segment sequence number (TMAMCSEQ)

Specifies the sequence number for a segment of a multi-segment message. This number must be updated for each segment, because messages are not necessarily delivered sequentially by XCF.

This number must have a value of 1 if the message has only one segment.

Resume tpipe requester ID (TMAMRTID)

The user ID of the client that is issuing the resume tpipe request.

Related concepts:

“Client/server resynchronization with OTMA” on page 830

Related reference:

“OTMA resynchronization protocol” on page 834

“State data section”

“Security data section” on page 880

“User data section” on page 882

“Application data section” on page 883

“Transaction and callout messages” on page 871

“Resume output for the hold queue for tpipe” on page 879

State data section

The state data is mandatory for any OTMA message. It immediately follows the message-control information section in the message prefix. It contains transaction-related information.

The state-data section has different formats for transaction-related information and for OTMA protocol commands. The state data section can be followed by a security-data section.

Server-Available and Client-Bid commands

The state data for the Server-Available and Client-Bid commands section of the OTMA message prefix is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

The following table summarizes the format for state data for command messages, including byte offset, length, content, hexadecimal value, the meaning, and usage comments.

Table 144. Server-Available and Client-Bid command format.

Byte	Length	Content	Value	Meaning
0	2	Length		Specifies the total length of the state-data section of the message prefix, including the length field.
2	16	Member name		Specifies the z/OS cross-system coupling facility (XCF) member name of originating server.
18	8	Originator's token		Specifies the XCF-member token of the originator (either client or server) of the message.
26	8	Destination token		Specifies the XCF-member token of the destination (either client or server) of the message.
Note: The following fields are present only for client-bid commands.				
34	8	DRU exit name		Specifies the name of the OTMA Destination Resolution exit routine.
42	2	MaxBlocksize		Specifies the maximum block size for XCF conversations between the server and the client. This field is optional

Table 144. Server-Available and Client-Bid command format (continued).

Byte	Length	Content	Value	Meaning
44	1	Client-Bid flag	X'80'	<i>Hold Queue</i>: Specifies that a hold queue for a tpipe is needed. The hold queue can hold commit-then-send output that is receives a NAK response, as well as alternate PCB output for an OTMA client. The output messages in the queue will not be delivered until the client requests that those messages be delivered. Use of the hold queue is optional. Without the hold queue, a regular queue for the tpipe will be used to hold and deliver all the Commit-then-send output messages. The default is no hold queue for a tpipe. OTMA resynchronization protocol currently does not support the hold queue.
			X'40'	<i>Reserved</i> : reserved for OTMA Callable Interface.
45	1	Additional client-bid flags	X'80'	<i>Message flood control</i> : specifies that OTMA use the maximum number of input messages provided by the client-bid command at byte 62 for message flood control.
			X'20'	<i>Member timeout</i> : specifies that OTMA use the member timeout value for send-then-commit messages provided by the client-bid command at byte 65.
			X'08'	<i>Super member flag</i>: activates the super member function for this member or client. To use this optional super member function, the member or client needs to set this flag and to specify the super member name, see the super member name at byte 54 below.
			X'04'	<i>Synchronous callout support</i> : specifies support for synchronous callout messages that are issued by IMS applications.
			X'02'	Set CM0 timeout queue name.
46	4	Aging value		Specifies the access control environment element (ACEE) aging value in seconds. If the age of an existing ACEE for a user ID is greater than this value, IMS creates a new ACEE. If the aging value specified is less than the supported minimum value, IMS creates a non-cached ACEE. The minimum value for caching support is 300 seconds (5 minutes).
50	4	Hash table size		Defines the size of the IMS OTMA hash table for processing multi-segment messages input. The hash table is used to correctly chain all the input segments together. The suggested value is X'00000065'.

Table 144. Server-Available and Client-Bid command format (continued).

Byte	Length	Content	Value	Meaning
54	4	Super member name		Specifies a 1-4 byte super member name. When the super member function is activated, all of the asynchronous output messages for hold queue capable member or client, such as IMS Connect, will be shared by a set of clients specifying the same super member name.
58	2	Offset to callout token		Specifies the offset in the user data section at which this client expects the synchronous callout token to be placed by OTMA.
60	2	<i>Reserved</i>		
62	2	Message flood threshold value	0-9999	Specifies the maximum number of concurrent input messages from this member or client. If 0 is specified, the OTMA default is 5000. Values between 1 and 200 are treated as 200.
64	1	<i>Reserved</i>		
65	1	Timeout value for acknowledgments from client		Specifies the member level timeout value in seconds for both commit-then-send (CM0) and send-then-commit (CM1) OTMA messages. For CM1 messages, this timeout value applies only when synclevel=confirm or synclevel=syncpt. Valid timeout intervals are 1 to 255 seconds. If no time out value is specified, the OTMA default is 120 seconds. The timeout value specified in this field cannot be greater than the timeout value set OTMA; otherwise, the value is ignored. The time out value is set in OTMA either by the /START TMEMBER command or the OTMA client descriptor.
66	8	CM0 timeout queue name		Specifies an 8-byte name for a tpipe timeout queue for commit-then-send (CM0) output.

Related reference:

“OTMA state data fields used by IMS Connect” on page 271

SRVresynch command

The state data for the SRVresynch command portion of the OTMA message prefix is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

The SRVresynch command is sent by IMS to pass all its known synchronized tpipe names to the client. If the command data can not fit into a single buffer, chained multi-segment buffers will be sent instead. The following table summarizes the format of state data for the SRVresynch command. The summary includes byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 145. SRVresynch command format

Byte	Length	Content	Description
0	2	Length	Length of the state-data section, including the length field.
2	8	Tpipe name	The transaction pipe name. The tpipe name can be repeated as necessary.

REQresynch command

The state data for the REQresynch command portion of the OTMA message prefix is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

The REQresynch command is used by IMS to pass the send-sequence number and the receive sequence for a specific tpipe to the client. The following table summarizes the format of state data for the REQresynch command. The summary includes byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 146. REQresynch command format

Byte	Length	Content	Description
0	2	Length	Length of the state-data section, including the length field.
2	8	Tpipe Name	The transaction pipe name.
10	4	Send-sequence number	IMS recoverable send-sequence number for the transaction pipe.
14	4	Receive-sequence number	IMS recoverable receive-sequence number for the transaction pipe.
18	1	Tpipe flag 1	Reserved for future use.
19	1	Tpipe flag 2	Reserved for future use.
20	6	RESERVED	

REPresynch command

The state data for the REPresynch command portion of the OTMA message prefix is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

The REPresynch command is sent by the client in reply to the REQresynch request from IMS. It contains the desired state information for a tpipe. The following table summarizes the format of state data for the REPresynch command. The summary includes byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 147. REPresynch command format.

Byte	Length	Content	Value	Description
0	2	Length		Length of the state-data section, including the length field.
2	8	Tpipe name		The transaction pipe name.
10	4	Send-sequence number		Client recoverable send-sequence number for the transaction pipe.
14	4	Receive-sequence number		Client recoverable receive-sequence number for the transaction pipe.

Table 147. REPresynch command format (continued).

Byte	Length	Content	Value	Description
18	1	Tpipe flag 1	X'00'	<i>Continue</i> : the last message has been received and IMS continues processing for this synchronized transaction pipe.
			X'04'	<i>Dequeue Last Output</i> : IMS can dequeue the last output message. The recoverable send-sequence number is updated.
			X'08'	<i>Reset Sequence Numbers</i> : IMS resets the recoverable send-sequence number and the recoverable receive-sequence number, as passed in this command.
			X'0C'	<i>Stop Tpipe</i> : IMS stops this synchronized transaction pipe.
			X'10'	<i>Stop Tpipe and wait for TBresynch</i> : IMS stops this synchronized transaction pipe and waits for TBresynch from the client.
19	1	Tpipe flag 2		Reserved.
20	6	Reserved		

TBresynch command

The state data for the TBresynch command portion of the OTMA message prefix is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

The TBresynch command is sent by the client to IMS if the client decides it is ready to receive REQresynch from IMS. The TBresynch command can be issued in the following two situations:

- The client has received an ACK message after sending REPresynch with “stop and wait for TBresynch” to IMS.
- The client may request a TBresynch with IMS at any time after the initial nondeferred resynchronization has completed for this tpipe.

The following table summarizes the format of state data for the TBresynch command. The summary includes byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 148. TBresynch command format

Byte	Length	Content	Description
0	2	Length	Length of the state-data section, including the length field.
2	8	Tpipe name	The transaction pipe name.

Transaction and callout messages

The state data for the transaction-related and synchronous callout request information in the OTMA message prefix is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

The following table summarizes the format of state data for transaction-related and synchronous callout request information. The summary includes byte, length, content, hexadecimal value, and descriptions of the use of each field.

Table 149. State data format for transaction-related information.

Byte	Length	Content	Value	Description
0	2	Length		Specifies the total length of the state-data section of the message prefix, including the length field.
2	1	Server state: specifies the mode in which the transaction is running.	X'80'	<i>Conversational State:</i> for a conversational transaction. Both the server and the client set this flag when sending conversational data.
			X'40'	<i>Response Mode:</i> for a response-mode transaction. Set by the server when sending output. This state has little significance for an OTMA server, because OTMA does not use sessions or terminals.
			X'20'	<i>From Hold Queue:</i> specifies that the output message was sent from the IMS hold queue for the tpipe. The server initially sets this flag when sending a commit-then-send output message. The client also needs to set this flag when sending the subsequent ACK or NAK to IMS.
			X'10'	<i>Hold queue capable:</i> asynchronous output can be sent to hold queue. Set by the server.
			X'08'	<i>Rerouted message:</i> output has been rerouted. Set by the server when processing a rerouted output.
			X'04'	<i>CM1 time out:</i> a time out occurs for this CM1 output. Set by the server.
			X'02'	<i>Transaction expiration: length of time:</i> When X'02' is specified, the transaction expiration function is enabled and submitted transactions expire after a specified length of time. The length of time is specified in the TMAMOSXP field.
			X'01'	This value has different meanings for different message types. <i>Transaction expiration: point in time:</i> On transaction messages from the client, X'01' enables the transaction expiration function and indicates that the submitted transaction expires at a specified point in time. Specify the point in time in STCK format in the user data prefix. Specify the byte offset of the time specification at byte 22 (the TMAMOSXP field) of the state data for transaction messages. <i>Resume tpipe token included:</i> For messages sent by the server from a tpipe hold queue, X'01' signifies that the output message includes a resume tpipe token at byte 48 (TMAMRTOK) in the OTMA state data prefix.

Table 149. State data format for transaction-related information (continued).

Byte	Length	Content	Value	Description
3	1	Synchronization or callout flag: specifies either the commit mode of a transaction or that a message is a synchronous callout request.	X'40'	<i>Commit-then-send:</i> a commit-then-send transaction. The server commits output before sending it to the client.
			X'20'	<i>Send-then-commit:</i> a send-then-commit transaction. The server sends output to the client before committing it.
			X'08'	<i>Synchronous callout message:</i> a synchronous callout message originating from an IMS application that is running in an IMS dependent region.
			X'02'	<i>Need dump for transaction expiration:</i> Set by client to request a symptom dump and a DFS554A message if this transaction expires on the IMS input queue before an application issues a GU call to retrieve it.
			X'01'	<i>Send SQ BE ALTPCB msg to FE IMS::</i> Set by the client to request that the output generated on a back-end IMS in a shared queues environment be returned to the client via the front-end IMS system.
4	1	Synchronization level: specifies the transaction synchronization level; that is, the way in which the client and an IMS application program interact for program output messages. The default is Confirm, in which IMS always requests a response when sending commit-then-send output to a client.	X'00'	<i>None:</i> specifies that the programs participate in coordinated commit processing on resources updated during the conversation under the z/OS Resource Recovery Services recovery platform. The server application program does not request an ACK message when it sends output to a client. None is only valid for send-then-commit transactions.
			X'01'	<i>Confirm:</i> specifies that synchronization is requested. The server sends transaction output with the response flag set to Response Requested in the message-control information section of the message prefix. Confirm can be used for either commit-then-send or send-then-commit transactions.
			X'02'	<i>Syncpt:</i> specifies that the programs participate in coordinated commit processing on resources updated during the conversation under the RRS recovery platform. A conversation with this level is also called a protected conversation and the resources updated under this conversation use the two-phase commit protocol.

Table 149. State data format for transaction-related information (continued).

Byte	Length	Content	Value	Description
5	1	Client flags: specifies optional processing requested by the client.	X'80'	<i>Send Only Message:</i> this is a send only message; the response is placed on the hold queue. This flag is valid only if the client requested a hold queue during open.
			X'40'	<i>Aging Value or NAK message for callout</i> If specified for the aging value, this flag instructs IMS to accept the accessor environment element (ACEE) aging value specified at byte 20 of the message control information. If specified on a NAK response to a synchronous callout message, OTMA keeps the callout message on the queue until it is retrieved by another resume tpipe call or the callout message times out.
			X'20'	<i>Reroute Request:</i> reroute requested. Setting this flag reroutes CM0 output to the destination that is specified in the Destination Override field.
			X'10'	<i>Purge CM0 IOPCB output:</i> when CM0 IOPCB output cannot be delivered to the client, delete it. Set by the client in the input message.
			X'08'	<i>Enable message level time out:</i> enable the timeout value specified in the message-control information section. This timeout value only applies to this particular CM1 message. If there is no message level time out specified, the member level timeout value is used.
			X'04'	Obsolete. EWLM is no longer supported by IMS. If this flag is specified, it is ignored by OTMA.
			X'02'	<i>Ignore PURG Calls:</i> for commit then send (CM0) messages that generate multiple PURG calls in the TP PCB, when this optional flag is set, IMS ignores the PURG calls so that the OTMA client application receives one response CM0 message with multiple output segments. When this flag is not set, if there are multiple ISRT and PURG calls on the TP PCB for a CM0 message, the OTMA client application receives multiple response messages. This flag is applies only to CM0 input messages. If this flag is specified in a CM1 input message, IMS resets the flag and ignores it for any remaining program-to-program switches and ALT-PCB processing.
			X'01'	<i>Send only ordered:</i> identifies the protocol as send only with serial delivery

Table 149. State data format for transaction-related information (continued).

Byte	Length	Content	Value	Description
6	8	Map name		Specifies the formatting map used by the server to map input or output data streams (for example, 3270 data streams). Although OTMA does not provide MFS support, you can use the map name to define the output data stream. The name is an 8-byte MOD name that is placed in the I/O PCB. IMS replaces this field in the prefix with the map name in the I/O PCB when the message is inserted. The map name is optional.
14	16	Server token		For CM1 messages, the server sets the server token for correlation purposes. After receiving a CM1 output message from the server, the client must return the server token on the following acknowledgment (ACK or NAK) or conversational iteration.
Alternate mapping for bytes 14 through 29 when server token is not specified or not relevant				
14	4	Super member name		On output from ALTPCB, specifies the super member name for the DFSYPRX0 and DFSYDRU0 exit routines
18	2	Transaction expiration value or offset		If the server state flag is set to X'02' (TMAMTXP2), this field contains the length of time in seconds after which the input transaction expires. If the server state flag is set to X'01' (TMAMTXP1), this field contains the offset to the point in time specification, in STCK format, in the user data prefix. This field cannot be used for CM1 conversational iterations.
20	2	Offset to correlation token for synchronous callout messages		Specifies the offset in the user data section at which the client can find the correlation token in a synchronous callout message. This value should match the offset specified by the client in the TMAMOSYN field of the state data in the client bid message.
22	8	Resume tpipe user ID		
End of alternate mapping for bytes 14 through 29				
30	16	Correlation token		A client token to correlate input with output. This token is optional and is not used by the server. Recommendation: Use this token to help clients manage their transactions.
46	16	CM1 context ID		On input transaction messages sent by the client, specifies the RRS token that is used with SYNCLVL=02 and protected conversations.
Alternate mapping for bytes 46 through 61 when CM1 context ID is not used				
46	8	Resume tpipe token		On output messages sent by the server from the tpipe hold queue, this field contains the resume tpipe request token.

Table 149. State data format for transaction-related information (continued).

Byte	Length	Content	Value	Description
54	8	Time received from XCF		In the message returned to a client after a transaction has expired, this field contains the time that the expired transaction was received by IMS from XCF.
End of alternate mapping for bytes 46 through 61				
62	8	Destination override		Specifies an LTERM name used to override the LTERM name in the IMS application program's I/O PCB. This override is used if the client does not want to override the LTERM name in the I/O PCB with the transaction-pipe name. This optional override is not used if it begins with a blank.
70	2	Server user data length		Specifies the length of the server user data, if any. The maximum length of the server user data is 256 bytes. The server user data length is not included in the length calculation.
72	*	Server user data		Specifies any data needed by the server. If included in a transaction message by the client, it is returned by the server in the output data messages. Variable length. Optional.

Server state protocol command

The state data for the server state protocol command is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

The server state protocol command is used by OTMA both to notify the OTMA client about the current state or a change in the state of IMS processing and as a heartbeat message, which is issued at 60 second intervals.

The server state protocol command is identified by X'3C' in the command type field (TMAMCTYP) of the message control information section of the OTMA header.

The following table summarizes the format of state data for the server state protocol command. The summary includes byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 150. Server state protocol command format

Dec offset	Hex offset	Length	Field name	Value and description
0	X'00'	2	TMAMHLEN	Length of the state-data section, including the length field itself.

Table 150. Server state protocol command format (continued)

Dec offset	Hex offset	Length	Field name	Value and description
2	X'02'	2	TMAMRSIM_STATUS	The state of OTMA resources in the server. X'03' - Normal state IMS is available and processing OTMA messages normally. X'02' - Degraded state IMS is processing OTMA messages slowly. OTMA issues a degraded state protocol command when one or more conditions indicate that IMS is not processing OTMA messages as quickly as it should. X'01' - Unavailable state IMS can no longer accept OTMA transactions for processing. OTMA issues the unavailable state protocol command to alert the OTMA client that one or more severe conditions prevent IMS from processing OTMA messages.
4	X'04'	1	TMAMRSIM_SVRFLG1	Flags for resource in the first group of unavailable resources. X'80' - TMAMRSIM_S1SHTDN The IMS server is shutting down and no longer available.
5	X'05'	1	TMAMRSIM_SVRFLG2	Reserved
6	X'06'	1	TMAMRSIM_SVRFLG3	Reserved
7	X'07'	1	TMAMRSIM_SVRFLG4	Flags for resource in the fourth group of unavailable resources. X'01' - TMAMRSIM_S4FLOOD The server is flooded with OTMA messages that are waiting to be processed and is no longer available.
8	X'08'	1	TMAMRSIM_WRNFLG1	Flags for resources in the first group of degraded resources. X'80' - TMAMRSIM_W1FLOOD Global flood warning for all OTMA clients X'40' - TMAMRSIM_W1MTP Warning for OTMA clients that specify a maximum allowable number of tpipes: the total number of tpipes has reached the global warning threshold defined by the highest value specified on the MAXTP parameter of any OTMA client descriptor. X'20' - TMAMRSIM_W1MTPF The MAXTP limit defined by the DFSOTMA client descriptor entry was reached. Subsequent requests for new tpipes from all OTMA clients are rejected. An exception to this is when MAXTPBE=NO is defined for a backend IMS system for application GU processing in a shared queues environment. X'10' - TMAMRSIM_W1MTP80 The number of tpipes in the IMS system reached eighty percent of the global tpipe limit that is defined by the MAXTP parameter in the DFSOTMA system client descriptor.
9	X'09'	1	TMAMRSIM_WRNFLG2	Reserved
10	X'0A'	1	TMAMRSIM_WRNFLG3	Reserved

Table 150. Server state protocol command format (continued)

Dec offset	Hex offset	Length	Field name	Value and description
11	X'0B'	1	TMAMRSIM_WRNFLG4	Flags for resources in the fourth group of degraded resources. X'08' - TMAMRSIM_W5MTP The number of tpipes for this OTMA client has reached the maximum allowable number of tpipes set for this client on the MAXTP parameter of the OTMA client descriptor. No new tpipes can be created for this OTMA client until the number of tpipes drops. X'04' - TMAMRSIM_W4MTP The number of tpipes for this OTMA client has reached 80% of the maximum allowable number of tpipes set for this client on the MAXTP parameter of the OTMA client descriptor. X'02' - TMAMRSIM_AWE Message AWE reaches 80% flood X'01' - TMAMRSIM_W4FLOOD Flood warning for this client only. The number of OTMA messages waiting to be processed on the server is at eighty percent of the maximum allowable number defined for the server.
12	X'0C'	1	TMAMRSIM_NRSFLGS	Other flags for non-resource related indicators. X'80' - TMAMRSIM_HB60S Identifies this message as a heartbeat message. The server is available and resource usage is within normal limits. Heartbeat messages are sent every 60 seconds.
13	X'0D'	3	Reserved	
16	X'10'	16	TMAMRSIM_SRVNAME	The 16 character z/OS cross-system coupling facility (XCF) member name of the OTMA server.
32	X'20'	16	TMAMRSIM_CLTNAME	The 16 character XCF member name of the OTMA client.
48	X'30'	20	Reserved	
68	X'44'	12	TMAMRSIM_UTC	UTC time for this message

Related concepts:

"OTMA resource monitoring" on page 786

Related reference:

"Message-control information section" on page 855

Resume output for Tpipe

The state data for resume output for Tpipe portion of the OTMA message prefix is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

This command is sent by the client to force any queued output to be resent again. The number of tpipes and tpipe names are needed in the command.

If the hold queue for the tpipe exists and holds messages, those messages will also be sent to the client. The following table summarizes the format of state data for Resume Output for tpipe. The summary includes byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 151. Resume output for Tpipes command format

Byte	Length	Content	Description
0	2	Length	Length of the state-data section.
2	2	Tpipe count	Number of tpipe names in the command.
4	8	Tpipe name	The transaction tpipe name.
Different tpipe names can be added as necessary.			

Resume output for the hold queue for tpipe

The state data for resume output for the hold queue for tpipe portion of the OTMA message prefix is mapped by the TMAMHDR DSECT of the DFSYMSG macro.

An OTMA client sends the command to inform IMS to deliver one or all queued messages on the hold queue for tpipe. If this command is not issued, messages are held in the hold queue. However, the option that is specified in the command can be used to request how IMS holds and delivers messages. One of the four options in the following table can be specified in the State Data. If the client or z/OS cross-system coupling facility returns a NAK message to IMS, the current option is reset to No-Auto, which is the default.

The following table summarizes the format of state data for Resume Output for the hold queue tpipe. The summary includes, as appropriate, byte, length, content, hexadecimal value, and the description.

Table 152. Resume output for the hold queue for tpipes command format.

Byte	Length	Content	Value	Description
0	2	Length		Length of the state-data section.
2	1	Delivery option	X'00'	<i>No-Auto</i> : exhaust all the messages in the queue only when the command is issued. This is the default.
			X'01'	<i>One Only</i> : deliver one message in the queue when the command is issued.
			X'02'	<i>Auto</i> : exhaust all the messages in the queue. After that, automatically deliver messages when they are queued.
			X'04'	<i>Auto-One</i> : deliver one message automatically when a message is available in the queue. The message might already be in the queue or it might be delivered later. After the message is delivered, this option is reset to No-Auto.
3	1	Callout mode	X'80'	RESUME TPIPE call retrieves only synchronous callout messages.
			X'40'	RESUME TPIPE call retrieves both synchronous callout messages and asynchronous messages.
4	8	Resume tpipe token		OTMA clients, such as IMS Connect, generate a resume tpipe token to uniquely identify RESUME TPIPE requests.
				OTMA uses the token to queue and process in order multiple resume tpipe requests from the same client.
				The token is also used when the client cancels a resume tpipe request.

Cancel resume output for tpipe hold queue request

The OTMA cancel resume output for tpipe hold queue request protocol command is sent by the client to cancel a resume tpipe request that was previously submitted by the client.

Upon receiving the request to cancel a resume tpipe, OTMA uses the resume tpipe token to locate the request to be canceled. If found, the resume tpipe request is discarded.

Table 153. Format of the cancel resume output for tpipe hold queue request protocol command

Byte	Length	Content	Description
0	2	Length	Length of the state-data section.
2	2	Reserved	Reserved field
4	8	Token	The resume tpipe token of the resume tpipe request to be canceled.

No messages on tpipe hold queue

The OTMA server sends the no messages on tpipe hold queue protocol command to inform the OTMA client that the tpipe hold queue does not contain a message or response for the current resume tpipe request.

OTMA issues this protocol command in when any of the following events occur:

- When the option of the "Resume Output for the Special Queue for Tpipe" is TMAMCRHQ_ONE and there is no IMS message for the client.
- When the option of the "Resume Output for the Special Queue for Tpipe" is TMAMCRHQ_NOAUTO and there is no IMS message for the client.
- When the option of the "Resume Output for the Special Queue for Tpipe" is TMAMCRHQ_NOAUTO and OTMA flushes all the existing messages in OTMA queue.

Table 154. Format of no messages on tpipe hold queue protocol command

Byte	Length	Content	Description
0	2	Length	Length of the state-data section.
2	2	Reserved	Reserved field
4	8	Token	The resume tpipe token of the resume tpipe request for which there are no messages.

Security data section

The security-data section is mandatory for every transaction or command, and is optional for OTMA protocol commands.

The security data portion of the OTMA message prefix is mapped by the TMAMSEC DSECT of the DFSYMSG macro.

The following table is a summary of the content of the security-data section of the message prefix. The summary includes, as appropriate, byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 155. Content of security data fields.

Byte	Length	Content	Value	Description
0	2	Length		Length of the security-data section, including the length field.
2	1	Security flag	N	<i>No Security</i> : no RACF checking is done. It is assumed that the user ID and password are already verified.
			C	<i>Check</i> : RACF checks transactions and commands. Transaction and command authorization RACCHECKs are performed (TCLASS and CCLASS).
			F	<i>Full</i> : RACF checks transactions, commands, and regions. Transaction, IMS command, and MPP region authorization RACCHECKs are performed.
3	1	Reserved		
	1	Utoken length		Length of Utoken plus the length of Utoken Type. Length does not include length field itself.
	1	Utoken type	X'00'	Type of data to follow.
	*	Utoken		The user token. Variable length, from 1 to 80 bytes.
	1	User ID length		Length of the user ID plus the length of the User ID Type. Length does not include length field itself.
	1	User ID type	X'02'	Type of data to follow.
	*	User ID		The user ID. Variable length, from 1 to 8 bytes.
	u	Profile length		Length of the profile plus the length of the Profile Type. Length does not include length field itself.
	1	Profile type	X'03'	Type of data to follow.
	*	Profile		The SAF profile. Variable length, from 1 to 8 bytes.

Related reference:

“Explanation of OTMA security data fields”

Explanation of OTMA security data fields

The following information provides additional detail on the content of the security-data section of the message prefix.

Length

Specifies the length of the security-data section of the message prefix, including the length field.

Security Flag

Specifies the type of security checking to be performed. It is assumed that the user ID and password are already verified.

No Security

Specifies that no security checking is to be done.

Check Specifies that transaction and command security checking is to be performed.

Full Specifies that transaction, command, and MPP region security checking is to be performed.

Reserved

After the reserved field, the following three fields can be omitted or appear in any order. Each field has the following structure:

- Length field
- Field type
- Data field

The length field is not calculated in the length calculation. The actual length of the user ID or profile should not be less than the value specified for the length of each field.

Utoken Length

Specifies the length of the user token plus the length of the user token type.

Utoken Type

Specifies that this field contains a user token.

Utoken

Specifies the user token. The user ID and profile are used to create the user token. The user token is passed along to the IMS dependent region.

If the client has already called RACF, it should pass the Utoken with field type X'00' so that RACF is not called again.

User ID Length

Specifies the length of the User ID plus the User ID type.

User ID Type

Specifies that this field contains a user ID.

User ID

Specifies the actual user ID.

Profile Length

Specifies the length of the profile plus the length of the profile type.

Profile Type

Specifies that this field contains a profile.

Profile

Specifies the system authorization facility (SAF) profile. For RACF, this is the group name.

User data section

The user-data section of the OTMA message prefix is variable length and follows the security-data section. It can contain any data.

The user data portion of the OTMA message prefix is mapped by the TMAMUSR DSECT of the DFSYMSG macro.

The following table is a summary of the content of the user-data section of the message prefix. The summary includes, as appropriate, byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 156. Content of user data fields

Byte	Length	Content	Description
0	2	Length	Length of the user-data section, including the length field.
2	*	User data	The user data. Optional; variable length.

Related reference:

“Explanation of OTMA user data fields”

Explanation of OTMA user data fields

The following information provides additional detail on the content of the user-data section of the message prefix.

Length

Specifies the length of the user-data section of the message prefix, including the length field. The maximum length of the user data is 1024 bytes.

User Data

Specifies the optional user data. This data is managed by the client, and can be created and updated using the DFSYDRU0 exit routine. The server returns this section unchanged to the client as the first segment of any output messages.

Reroute Tpipe

Specifies the reroute tpipe name for an input transaction submitted from an OTMA hold queue capable client, such as IMS Connect. This optional field starts at the offset X'5C' from the beginning of user data prefix and is used in the following two scenarios:

- When OTMA cannot deliver a CM0 IOPCB output message to the client, OTMA uses the reroute tpipe name, if specified, to reroute the output.
- When a send-only input transaction specifies the reroute tpipe name, any IOPCB output message which is the result of the Send-Only transaction will be queued to the reroute tpipe.

Application data section

The application-data section of OTMA messages is variable length and follows the user-data section of the message prefix.

You include IMS commands and transactions in the application-data section. The data in this section is unchanged by the receiver (server or client), and is transmitted directly to the server application program or to the client application program.

The application data portion of the OTMA message prefix is mapped by the TMAMAPP DSECT of the DFSYMSG macro.

The following table is a summary of the content of the application-data section of the message prefix. The summary includes, as appropriate, byte, length, content, hexadecimal value, the meaning, and includes usage comments.

Table 157. Application data

Byte	Length	Content	Description
0	2	Length	Length of the application-data section. The length includes the length field itself. The maximum length is 32 KB (32767 bytes).
2	2	ZZ	Application data IMS ZZ fields.
4	*	Application data	The optional application data. Multiple send requests might be required for a server output segment. For a client's transaction, the transaction code is specified in the first 8 bytes of the data area following the LLZZ. For transactions specified with MULTSEG, the standard IMS LLZZ format is required for each segment. The transaction code is only required in the first segment. Variable length. The maximum length of application data is 32 KB-4.

Sample OTMA messages

The following three sample OTMA messages are intended to show what OTMA messages look like when fully constructed, including the parts of the message prefix. The examples are not necessarily related to each other.

OTMA client-bid message

The following figure shows an OTMA client-bid message. The total length of the state-data section plus the security-data section of the message prefix is X'8C' bytes.

MESSAGE CONTROL INFORMATION:

01102000 04004040 40404040 4040A0C0
00000000 00000000 00000000 00000400

| ff{
| |

STATE DATA + SECURITY DATA:

0036C3D3 C9C5D5E3 F1404040 40404040
40400100 00010003 00020100 00010003
0001C4C6 E2E8C4D9 E4F02000 00007FFF
FFFF0000 00650056 C3525100 50018059
15569555 55555555 55555555 B7B686B0
81A61515 1B1B1B1B 1B1B1B1B B7B686B0
81A61515 55555555 55555555 8C918CA4
15151515 55555555 55555555 09151515
15151515 15151515 15151515

| ..CLIENT1
|
| ..DFSYDRU0....".
| C...&...
| ..n.....%&f[
| aw.....%&f[
| aw....._j_u
|
| |

OTMA transaction message

The following figure shows an OTMA transaction message. The total length of the state-data, security-data, and application-data sections of the message prefix is X'D6' bytes.

MESSAGE CONTROL INFORMATION:

01402000 0000E3D7 C9D7C5F1 4040A0D0

| TPIPE1 ff} |

00000001 00000000 00000001 00010000	
STATE DATA + SECURITY DATA + APPLICATION DATA:	
00480020 0100E3C5 E2E3D4C1 D7400000	TESTMAP ..
00000000 00000000 00000000 0000C9D4	IM
E2F0F0F0 F0F10000 00000000 00004040	S00001.....
40404040 40404040 40404040 40404040	
40404040 40400000 0056C652 51005001	F...&.
80465551 95555555 55555555 55555555	n.....
55555555 55555555 55555555 55555555	
55555555 55555555 55555555 555586A3	ft
A781B0B7 A4155555 55555555 5555B1B7	xa[%u.....&%
8CB6A5A5 A415B7BD B7A41515 15150038	%&vvu.%"%u.....
0000C1D7 D6D3F1F8 4040E2C1 E840C8C5	..APOL18 SAY HE
D3D3D640 40404040 40404040 40404040	LLO
40404040 40404040 40404040 40404040	
40404040 4040	

OTMA response message

The following figure shows an OTMA response message. The total length of the state-data, security-data, and application-data sections of the message prefix is X'EE' bytes.

MESSAGE CONTROL INFORMATION:

01A08000 0000E3D7 C9D7C5F1 404080D0	.ff....TPIPE1 .}
00000001 00000000 00000001 00010000	

STATE DATA + SECURITY DATA + APPLICATION DATA:

00480020 0100E3C5 E2E3D4C1 D740AB7F	TESTMAP %"
28EB9FAD 9A024040 40404040 4040C9D4	.. Y.. IM
E2F0F0F0 F0F10000 00000000 00004040	S00001.....
40404040 40404040 40404040 40404040	
40404040 40400000 0056C652 51005001	F...&.
80465551 95555555 55555555 55555555	n.....
55555555 55555555 55555555 55555555	
55555555 55555555 55555555 555586A3	ft
A781B0B7 A4155555 55555555 5555B1B7	xa[%u.....&%
8CB6A5A5 A415B7BD B7A41515 15150050	%&vvu.%"%u.....&
0300D6E4 E3E2C5C7 40D5D67E F0F0F0F0	..OUTSEG NO=0000
F140E2D7 C5C3C9C6 C9C5C440 E2C5C7E2	1 SPECIFIED SEGS
C9E9C57E F0F0F0F8 F06B40E2 C5C7D5D6	IZE=00080, SEGNO
7EF0F0F0 F0F34040 40404040 40F67EC3	=00003 6=C
D6D340F6 F04040F7 7EC3D6D3 40F7	OL 60 7=COL 7..

Chapter 48. OTMA Callable Interface

The IMS OTMA Callable Interface (C/I) provides a high-level interface that allows application programs on other z/OS subsystems access IMS applications through OTMA.

The OTMA C/I API consists of API calls that are available to a C/C++ program. The API calls are used to join the IMS/OTMA z/OS cross-system coupling facility (XCF) group, to connect to IMS, to allocate communication sessions, to send IMS transactions/commands, to receive output from IMS, to close communication sessions, and to leave the XCF group.

The OTMA C/I API calls and sample OTMA C/I application programs are documented in *IMS Version 11 System Programming APIs*. Codes returned by the OTMA C/I are documented in *IMS Version 11 Messages and Codes, Volume 4: IMS Component Codes*

The following figure provides an overview of the OTMA C/I. Shown from left to right, in a sample z/OS environment, are sample C and C++ API calls (for example, OTMA_OPEN). The API calls pass through an object stub, the SVC interface routine, the API (for example, DFSYOPEN), and finally through the XCF group to IMS OTMA.

OS/390 environment




Figure 151. OTMA Callable Interface overview

The program that invokes the API calls can be running from an authorized or unauthorized library in problem or supervisor state. DFSYC0, a C header file, is provided to define the API calls. DFSYCRET, a load module, contains the entry points for each API call and is linked with the application program. OTMA C/I uses BPE SVC Services to process the API call.

Benefits: A key benefit of OTMA C/I is that it is easy to use.

Other reasons to use OTMA C/I are that it:

- Extracts out the details of OTMA and XCF
- Submits IMS transactions and commands

- Enables programs running from other z/OS subsystems to connect to multiple IMS systems
- Calls the APIs from an authorized or unauthorized library
- Connects to all IMS OTMA releases

OTMA C/I initialization

OTMA C/I provides a stand-alone program, DFSYSVIO, that must be run to initialize the OTMA C/I after either an IPL of z/OS or the application of maintenance to the OTMA C/I.

DFSYSVIO invokes DFSYSVC0, one of the OTMA C/I modules. DFSYSVC0 loads and registers the SVC services by an authorized address space running on the same z/OS image as the application programs accessing it.

You must add an entry in the z/OS program properties table (PPT) for the OTMA Callable Interface initialization program. The steps for doing this are:

1. Edit the SCHEDxx member of the SYS1.PARMLIB data set.
2. Add the following entry to the SCHEDxx member:

```
PPT PGMNAME(DFSYSVIO) /* PROGRAM NAME = DFSYSVIO */
      CANCEL /* PROGRAM CAN BE CANCELED */
      KEY(7) /* PROTECT KEY ASSIGNED IS 7 */
      SWAP /* PROGRAM IS SWAPPABLE */
      NOPRIV /* PROGRAM IS NOT PRIVILEGED */
      DSI /* REQUIRES DATA SET INTEGRITY */
      PASS /* CANNOT BYPASS PASSWORD PROTECTION */
      SYST /* PROGRAM IS A SYSTEM TASK */
      AFF(NONE) /* NO CPU AFFINITY */
      NOPREF /* NO PREFERRED STORAGE FRAMES */
```

3. Take one of the following actions to make the SCHEDxx changes effective:
 - Re-IPL the z/OS system.
 - or
 - Issue the MVS SET SCH= command.

Related Reading: For additional reading about updating the program properties table, see *z/OS MVS Initialization and Tuning Reference*.

A sample JCL procedure for running DFSYSVIO is as follows:

```
//OTMAINIT PROC RGN=3000K,SOUT=A
//*
//IEFPROC EXEC PGM=DFSYSVIO,
//          REGION=&RGN
//*
//STEPLIB DD DISP=SHR,UNIT=SYSDA,
//          DSN=IMSVS.SDFSRESL
//*
//SYSPRINT DD SYSOUT=&SOUT
//SYSUDUMP DD SYSOUT=&SOUT
```

After OTMA C/I is installed and initialized, you can use it to connect to any IMS OTMA release.

OTMA C/I security

To protect z/OS cross-system coupling facility groups from any non-authorized caller use IMSXCF.OTMACI, a RACF resource, defined in the RACF facility class for the OTMA C/I.

When the RACF resource is defined, RACF RACHECK is invoked before OTMA C/I performs a XCF JOIN. This method protects the access to XCF, the XCF group, and the member. This RACF checking is performed only when a non-authorized caller is using OTMA C/I.

OTMA C/I restrictions

Certain restrictions apply to the OTMA Callable Interface (C/I).

These restrictions include:

- C/I must be installed in a z/OS environment before it can be invoked. If C/I is not installed and invoked, an F92 abend occurs when `otma_create` or `otma_open` is issued. If C/I is not properly installed, a DFS3911E error message occurs.
- Application program languages other than C and C++ are not currently supported by OTMA C/I.
- All OTMA calls must be made in the same state (PSW key, supervisor or problem state, authorized or non-authorized) as the `otma_open` call. For example: If you were authorized when you did the `otma_open` call, you must be authorized for all subsequent calls.
- The resynchronization feature of IMS OTMA is not supported.
- IMS command `/SECURE OTMA PROFILE`, is not currently supported.

Timing out OTMA C/I sessions after `otma_send_receive` API calls for CM1 transactions

The OTMA C/I does not include a timeout function for the `otma_send_receive` API call and does not support the OTMA message-level transaction expiration function; however you can still implement a timeout function.

To implement a time out function for OTMA C/I client sessions use one of the following methods:

- Coding the `EXPRTIME` parameter on the `TRANSACT` system definition macro when you define the IMS transactions that will be invoked by the OTMA C/I client.
- Including a time out function with your OTMA C/I client.
- By both coding the `EXPRTIME` parameter and including a time out function in your OTMA C/I client.

Recommendation: Use the `EXPRTIME` parameter to implement a timeout function for transactions submitted through OTMA C/I instead of including a time out function with your OTMA C/I client. If the OTMA C/I client implements the time out function, although it is unlikely, a z/OS X'0C4' abend can occur if the OTMA C/I client releases the session storage after OTMA has received a response from IMS, but before OTMA posts the wait ECB.

If your OTMA C/I client includes a time out function, you can reduce the possibility of a z/OS X'0C4 abend by:

- When coding the OTMA C/I client time out function, include a delay in processing after the OTMA C/I client issues the `otma_free` call and before it frees the session storage.
- Specify a shorter time out value on the `EXPRTIME` parameter of the `TRANSACT` system definition macro for IMS transactions invoked by the OTMA C/I client.

Related reference:

TRANSACT macro (System Definition)

Chapter 49. OTMA architected transaction attributes

When you issue an IMS /DISPLAY TRANSACTION command from OTMA, the output is in the form of an OTMA message, returned to the client in the application-data section of the message prefix.

The following table shows the Attributes Segment for a given transaction and describes the syntax of architected command output. The description includes byte, length, content, hexadecimal value, the meaning, and includes usage comments where appropriate.

Table 158. Transaction attributes segment

Byte	Length	Content	Value	Meaning	Comments
0	2	Length		Length of the Transaction Attributes Segment (LL).	This length includes the length field itself.
2	2	ZZ			
4	8	Transaction code		The 8-byte IMS transaction code.	
12	1	Transaction type flag 1			The flag 1 values are mutually exclusive.
		<i>Valid</i>	X'00'	A valid OTMA transaction.	
		<i>CPIC</i>	X'04'	A CPI-C (APPC) transaction.	
		<i>FPX</i>	X'08'	A Fast Path-exclusive transaction.	
		<i>FPP</i>	X'0C'	A Fast Path-potential transaction.	
		<i>MSC</i>	X'10'	An MSC remote transaction.	
		<i>Invalid Syntax</i>	X'FE'	Syntax error.	The data area contains the text of the error.
		<i>Invalid</i>	X'FF'	Transaction not found, or invalid.	
13	1	Transaction type flag 2			The flag 2 values are not mutually exclusive.
		<i>Response</i>	X'80'	An IMS response mode transaction.	
		<i>Conversation</i>	X'40'	An IMS conversational transaction.	
		<i>Update</i>	X'20'	The transaction has update capability.	
		<i>Irrecoverable</i>	X'10'	The transaction is defined as irrecoverable.	

Table 158. Transaction attributes segment (continued)

Byte	Length	Content	Value	Meaning	Comments
		<i>Multi-Segment</i>	X'08'	The transaction has multiple segments.	
		<i>Uppercase</i>	X'04'	Uppercase translation requested.	
14	1	Transaction status			The indicated values are not mutually exclusive.
		<i>STOP</i>	X'80'	The transaction input queuing is stopped.	One of the following IMS commands was issued for the transaction: • /STOP • /PURGE
		<i>OLC</i>	X'40'	The transaction input queuing is stopped.	One of the following IMS commands was issued for the transaction: • /MOD PREPARE • /MOD COMMIT
		<i>NOSCH</i>	X'20'	IMS scheduling is stopped.	One of the following IMS commands was issued for the transaction: • /PSTOP • /LOCK
15	1	Reserved			
16	2	Length		Length (LL) of error text, if any.	If there is error text, it replaces all subsequent sections of the message.
18	*	Error text		Text of the error message.	Variable length. The error test section is applicable only if transaction type flag 1 is set to Invalid Syntax (X'FE')
16	8	PSB name		IMS PSB name.	Only present if there is no error text.
24	1	Class		SMB message class for IMS scheduling.	
25	1	Current priority		Current SMB priority.	
26	1	Normal priority		Normal SMB priority.	
27	1	Limit priority		SMB limit priority.	
28	2	Enqueue count		Number of messages enqueued.	
30	2	Dequeue count		Number of messages dequeued.	
32	2	Enqueue limit		Enqueue limit count.	
34	2	Processing limit count		Processing limit count.	
36	2	Output max segment length		The maximum output segment length.	
38	2	Output limit of message segments		The output limit of message segments.	

Table 158. Transaction attributes segment (continued)

Byte	Length	Content	Value	Meaning	Comments
40	2	Parallel limit		The PARLIM value from TRANSACTION statement.	
42	1	Region count		The number of regions in which the transaction is currently scheduled.	

Part 11. SLU P and Finance Communication

These topics discuss in detail SLU P and Finance Communication, including an overview of SLU P and Finance communication, the IMS facilities used for SLU P and Finance communication, network operation with SLU P and Finance communication, and SLU P network protocols.

Chapter 50. Overview of SLU P and Finance Communication

This topic and the topics below present an administrative overview of SLU P and Finance Communication Systems, and explains how IMS implements this architecture.

User-written application programs for IBM Finance Communication Systems can be defined to operate in two different ways with IMS—as an IBM 3600/4700 Finance Communication Controller or as a secondary logical unit type P (SLU P). These systems are described as `UNITYPE=FINANCE` and `UNITYPE=SLUTYPEP`, respectively, on the `TYPE` macro statement.

Definitions:

- The term *SLU P* is used in the following topics when describing support that is applicable to both system types.
- The terms *Finance* and *SLUTYPEP* are used in the following topics when it is necessary to distinguish between the systems.

A major difference between SLU P and Finance systems is the level of MFS support available to workstations in the programmable control unit.

- Finance—Workstations can be identified as displays, printers, passbook printers, and ATMs (automated teller machines), such as the 4730 systems. MFS can provide detailed support for display paging, printer page formatting, and ATMs.
- SLU P—Each LU in the controller has a program name and its unique device characteristics are unknown to IMS. The programmable controller is responsible for device control and formatting. The MFS administrator and the program administrator for the remote controller need to establish data field structures for exchanged messages. MFS accepts the input structure and rearranges it as required for the IMS application program. At output, MFS accepts application program-provided data and converts it to the correct format for transmission to the controller. This MFS facility is called distributed presentation management (DPM).

Related tasks:

“ETO and 3600/Finance and SLU P” on page 130

The IMS-SLU P network

The three major elements of an IMS-SLU P network are the controller and the terminals executing the remote program, the communication link, the central processor executing the host IMS system.

Each element includes programming to perform a part of the total data processing performed by the system. IMS resides in the host processor and communicates with an application program in the SLU P system's controller through the communication link. The controller's application program monitors the terminals attached to the controller.

IMS supports 4700 terminals attached to an IBM 4701/4702 Finance Communication Controller. IMS supports the IBM 4730 Personal Banking Machine, both directly attached to a 37x5 as a SLU P and attached to a 4701/4702 as part of the SLU P system.

Related reading: For a list of 4700 terminals, see *IBM 4700 Finance Communication System: System Summary*.

IMS also supports 3600 terminals attached to either an IBM 3602 or 4701/4702 as part of a Finance system. (IMS does not distinguish between 3600- and 4700-series terminals.)

Examples of other IBM products that connect to IMS with the SLU P protocols are the Series/1, the IBM 3650, and the IBM 8100.

Each of these terminals, or workstations, can be defined to IMS as a component of the appropriate SLU P system using the COMPT or ICOMPT keyword on the TERMINAL macro statement.

System configuration

Before configuring a SLU P system, you must identify the financial operations that are required of the planned system. After you have done this, you can define to IMS the configuration of the logical units and components.

The SLU P system is regarded by IMS as a subsystem and consists of one or more logical units.

Definition: When configuring the system, you can define one or more *logical units*. These logical units can consist of devices, storage, and application programs. These logical units are referred to as *logical workstations* in Finance Communication System publications.

SLU P and Finance workstations

A workstation can consist of one or more terminals. Each workstation performs a specific type of financial operation, such as teller operation (deposits and withdrawals), report printing, or cash dispensing.

The workstation, as defined to IMS, does not need to reflect the actual devices attached to it at the 4701/4702, but the workstation must reflect the view of the workstation as seen by the IMS system.

A sample configuration is available from the IMS library (IMS.ADFSMAC; member name=DFSCP360).

System controller application program

The operation of a workstation is controlled by a user-written application program that resides in the SLU P system's controller. The application program can be designed to control one or more workstations and needs only to perform the functions required by its workstations.

The functions available to the application program include:

- Reading and writing to the terminals associated with the workstation
- Editing and verifying the data received from the terminal
- Reading and writing to the controller diskette
- Reading and writing to the host processor
- Editing and verifying data received from the host processor

- Operating offline when the host processor or IMS is unavailable

Writing the controller application program with MFS and XRF

If your system uses MFS and XRF, there are certain actions you must take when writing a controller application program.

When writing the controller application program for a system that uses MFS, do each of the following:

Note: Do not specify the WL (warning line) option on the 4700 controller's TERMINAL macro instruction. Specifying the WL option can cause an unexpected Finance system data exception error to occur on an LWRITE instruction.

- Specify the PS (page size) option on the 4700 controller's TERMINAL macro only if EJECT is specified on the IMS MFS DEV statement. If PS is not specified, the EJECT from IMS results in a new line function. If the “end-of-page” condition occurs on a workstation, an unexpected Finance system data exception error can occur. An EJECT from IMS resets the line count, resulting in an “end-of-page ” condition. If the CFOLD option is specified on the 4700 controller's TERMINAL macro instruction, the formats defined for MFS should not cause spacing for the center fold on a passbook. The SLU P system performs this function automatically.
- Check to see if the IMS MFS DFLD statement specifies ATTR=YES. If it does, any attempt to print an underscore (X'6D') can result in either of the following:
 - A data check on a 3610/3612
Specify a print position for the EBCDIC value (X'6D') when the OUTRTBL statement specifies other than the 128-character set for the 3610/3612 device.
 - Printing of a blank on a 3618 device
Specify a print position for the EBCDIC value (X'6D') when the OUTRTBL statement specifies the 48-character set for a 3618 device.
- Check to see if the MFS DEV statement option of EJECT, BGNMXG, BGNPP, or ENDPP is used to send eject characters (X'0C') within one IMS output transmission (other than the last character). If so, an “intervention-required” condition occurs on the passbook printer after the eject. In this situation, your application program in the controller must be capable of issuing multiple LWRITE instructions to the passbook printer to print the IMS output transmission. Each LWRITE should include the data up to and including an eject character, or up to the end of the transmission.

Related reading: For more information on coding the controller's macros, see *IBM 4700 Finance Communication System: System Summary*.

Considerations for controller application programs for XRF systems

If your SLU P system operates in an XRF complex, your controller application program must be able to handle messages lost during an XRF takeover.

Many messages that are “in-flight” when the alternate system takes control are recovered with no action from the application program or terminal operator. In this case, the takeover is transparent.

When the takeover is not transparent, message DFS3861I is issued in one of the following combinations:

- Message DFS3861I only—If the output message is recoverable, this indoubt message is issued. If irrecoverable, this message is not issued. If you are receiving output from IMS, and the new active IMS system cannot determine if you received the message, IMS requests an exception response but is not receiving one from your program.
- A CANCEL command followed by message DFS3861I—multi-segment output is in progress (last-in-chain not yet sent). If the in-flight message is recoverable, you should receive the in-flight message again. If the in-flight message is irrecoverable, the next in-flight message (if one is on the queue) is issued.
- An exception response followed by message DFS3861I—your input message is lost. Resend your last input message.

Converting controller application programs from Finance to SLU P

All current functions available for logical units defined to IMS as Finance, except the SCAN/NOSCAN option and Finance MFS, have similar functions for logical units defined as SLU P.

IMS does not scan output for the control character sequence when sending output to a secondary logical unit type P. Therefore, the implied option is NOSCAN. If an application program uses MFS, it must be converted to use the distributed presentation management (DPM) option, which divides responsibility for message formatting between MFS and a user program residing within the logical unit. With the use of the DPM option, physical terminal characteristics are not defined to MFS. MFS formats and presents data to the user program component of the SLUTYPEP. The user program must complete the formatting, if necessary, and present the data to a physical device.

To execute correctly when defined to IMS as SLU P, existing remote user-written programs that run in a Finance controller must be converted as follows:

- Any input message headers must be converted to SLU P message headers.
- Any input message indicating only begin-bracket must also indicate change-direction.
- If the TERMINAL macro SCAN option is specified, the Finance controller application program must be converted to detect and process the same options provided by the SCAN option.
- If the controller application program uses MFS, the remote application must be converted to use the device-independent MFS distributed presentation management option.
- When converting from Finance-specific MFS to DPM, the MFS format descriptions must be redefined and recompiled using the IMS MFS Language utility.

VTAM facilities used

The physical transmission of data between IMS and a SLU P system is controlled by VTAM.

Related reading: For more information on the communication facilities that govern data transmission between IMS (a VTAM application program) and the controller, see *Network Program Products General Information* and *z/OS Communications Server: SNA Programming*.

The VTAM facilities that IMS uses, particularly for a SLU P system, include:

- Connection, disconnection, and establishing logon mode
- Messages and responses
- Definite response 1 and definite response 2 (All data transmitted between IMS and a workstation must request a definite response 1, definite response 2, definite response 1 and 2, or exception definite response 1 or 2. VTAM commands must request definite response 1.)
- Sequencing and chaining
- Quiescing
- Facilities for ensuring orderly communication, including the use of brackets and change-direction indicators. In SNA protocol, a bracket is one or more chains of request units (RU) and the responses that are exchanged between two LU-to-LU half sessions that represent a transaction. In VTAM, a change-direction indicator means that the sender has finished sending and is prepared to receive.
- Sequence-number recovery
- Receiving input, sending messages
- Unconditional bracket termination
- Primary Error Recovery Procedure (ERP)
The ERP is terminated by an IMS output error message sent to the workstation.

IMS also supports the class of service (COS) and session outage notification (SON) facilities.

Related concepts:

“Using SON/COS support in IMS” on page 430

Related reference:

“Error handling” on page 955

VTAM commands and indicators used with SLU P

VTAM commands and indicators (communication control information) are necessary for data transmission between an IMS application program and another VTAM logical unit.

The following table shows which VTAM commands and indicators IMS sends to and receives from the SLU P controller. Results are unpredictable if any unsupported commands or indicators are used.

With the exception of the request-recovery command and the change-direction indicator, which are described following the table, the commands and indicators operate as described in *z/OS Communications Server: SNA Programming*.

Recommendation: During a SLU P session, every input message must have either a change-direction (CD) or end-bracket (EB); otherwise, the session is terminated.

Table 159. VTAM commands and indicators sent and received by IMS

VTAM command/indicator	IMS sends to controller	IMS receives from controller
INDEPENDENT OF SESSION		
Initiate Session		X
Procedure Error	X	
Terminate Session		X

Table 159. VTAM commands and indicators sent and received by IMS (continued)

VTAM command/indicator	IMS sends to controller	IMS receives from controller
SESSION CONTROL		
Bind	X	
Clear	X	
Request-Recovery (RQR)		X
Set-and-Test-Sequence Numbers (STSN)	X	
Start Data Traffic (SDT)	X	
Unbind	X	
NORMAL (synchronous) FLOW		
Begin-Bracket (BB)	X	X
BID	X	
CANCEL	X	X
CHASE		X
End-Bracket (EB)	X	X
Logical Unit Status (LUS)		X
Quiesce Complete (QC)	X	
Ready-to-Receive (RTR)		X
Change Direction (CD)		X
EXPEDITED (asynchronous) FLOW		
Quiesce-at-end-of-chain (QEC)		X
Release Quiesce (RELQ)		X
Request-Shutdown (RSHUTD)		X
Shutdown (SHUTD)	X	
Shutdown Complete (SHUTC)		X
Signal		X

Request-recovery command

Upon receipt of the VTAM Request-Recovery (RQR) command, IMS issues a CLEAR. The SLU P application programmer must exercise care when sending this command while a recoverable IMS output message is in progress.

To send an RQR command to IMS, the application program must first respond to the output operation that is in progress. If the application program does not send the reply but instead sends only the RQR command, the controller automatically sends the outstanding DR2 reply at the LEXIT or next LREAD statement. If IMS receives the RQR command before the DR2 reply, the CLEAR generated by IMS can cause the DR2 reply to be lost. The status of the current output message is then unpredictable.

Restriction: Session initiation and resynchronization caused by an SNA Request-Recovery (RQR) command is not allowed when the node is in response mode and the response reply message is not yet available for output; that is, the input response mode transaction is still queued or in the process of execution.

Response-mode transactions are not recoverable or restartable prior to the application sync point; therefore, session input acknowledgment does not occur until input processing is complete.

Related concepts:

“Message resynchronization” on page 922

Change-direction indicator

IMS supports use of the change-direction indicator on the last chain element of an input message or a VTAM command. This allows device support to be more consistent with VTAM operation. The operation of IMS is not affected by the indicator.

Establishing connection and specifying logon modes

When establishing connection using the VTAM OPNDST macro instruction, IMS specifies session parameters that define the rules that a logical unit must follow when communicating with IMS.

IMS examines either a user-supplied or a VTAM default set of session parameters, and overlays only those parameters on which IMS has dependencies. The remaining bytes are not changed. You can include user data with the BIND parameters.

User data included within the BIND is used by IMS as input to the signon process and is similar in function to the IMS /SIGN command. This is required for ETO Finance terminals and SLU P terminals, but is optional for static terminals.

If you do not specify a mode name on the TERMINAL macro statement during IMS system definition or on the logon descriptor, you can supply a mode name using any of the following methods:

- On the VTAM network operator VARY command
- On the IMS OPNDST command
- On a request for system initialization by another logical unit

If you specify a mode name on the IMS TERMINAL macro statement or on the logon descriptor, it is used, unless you override it using the IMS /OPNDST command or the VTAM VARY command.

If you do not supply a mode name during IMS system definition using a logon descriptor, on an IMS command, or using CINIT, VTAM assumes a default mode name.

Related reading: For more information on establishing logon mode tables and defining logon mode table entries, see *z/OS Communications Server: SNA Resource Definition Reference*.

You can override the use of the default logon mode table entry in one of two ways:

- At system definition, you can specify the logon mode entry on the TERMINAL macro using MODETBL= keyword.
- You can specify on the MODETBL keyword of the /OPNDST command that the logon mode table entry replace the table entry defined at system definition.

IMS overrides the session parameters for function management, transmission services, and primary and secondary network protocol.

Related reference:

-  Bind parameters for SLU P and LU 6.1 (System Programming APIs)
-  TERMINAL macro (System Definition)
-  /OPNDST command (Commands)

Establishing connection with the XRF complex

To establish a session with IMS in an XRF complex, your logon request must include either the USERVAR name you defined in the USERVAR tables or the MNPS ACB name shared by both of the IMS systems in the XRF complex.

In the case of XRF with USERVAR, VTAM matches the USERVAR logon message to the IMS application that is currently active. For XRF with MNPS, VTAM links directly to the active IMS's MNPS ACB.

If you use XRF with USERVAR, SLU P terminal programs cannot use the APPLID name in the PLUNAME field of the BIND to reestablish a session with IMS. Instead, SLU P terminal programs must use the USERVAR in the user data field of the BIND. This restriction does not apply to XRF with MNPS, because this type of XRF system does not use a USERVAR.

SLU P systems are normally defined as class-1 terminals; in most cases in an XRF complex, when an alternate IMS takes over processing from an active IMS, the takeover is accomplished without losing class-1 terminal sessions. The terminals might or might not be aware of the takeover.

If you are connecting the 4730 Personal Banking Machine to an XRF IMS as a SLU P device, use the following machine data configuration options:

- Invalid message: 10 or 11
- Multiple commands: 10 or 11

The first digit (1) prevents the 4730 from sending an exception response to certain network messages, including the DFS3861 message sent after an XRF takeover. Exception responses to these messages terminate the SLU P session between the 4730 and IMS. The second digit (0 or 1) indicates whether errors should be logged.

Related reading: For more information on the 4730, see *IBM 4730 Personal Banking Machine Operations Support Manual*.

Related concepts:

“Considerations for controller application programs for XRF systems” on page 901

-  XRF variations: MNPS and USERVAR (System Administration)
-  Specifying MNPS to manage terminal sessions for XRF (Operations and Automation)

Related reference:

-  Finance communication system bind parameters (System Programming APIs)
-  Defining IMS.PROCLIB members for XRF (System Administration)
-  DFSHSBxx member of the IMS PROCLIB data set (System Definition)

Bracket and send/receive management

The change-direction and bracket indicators are used to begin, end, and control the direction of synchronous transmissions between a SLU P system and IMS.

IMS support of these indicators is summarized in the following table (X = supported).

Table 160. Use of VTAM bracket and change-direction indicators

Synchronous transmissions	BB	EB	BB/EB	BB/CD
Received by IMS with data	X		X	X
Sent by IMS with data		X	X	
Received by IMS with data flow synchronous commands (normal flow)		X		

IMS does not support the absence of bracket indicators or the use of end-bracket-only on inbound synchronous data. IMS does not send output synchronous data with begin-bracket (BB) only, change-direction (CD) only, or BB/CD. Because IMS always ends a bracket (unconditional bracket termination) on input synchronous data, it is unnecessary to send either EB or CD on output synchronous commands. Input synchronous commands can specify EB or CD, but must be consistent with the current bracket and send/receive states; otherwise, the session is terminated. IMS does not send VTAM indicators with data flow synchronous commands (normal flow).

Chapter 51. IMS facilities used for SLU P and Finance

The following topics describe in detail the IMS facilities that support the Systems Network Architecture (SNA) environment.

Component definition

IMS considers a *workstation* to be one physical terminal. If a workstation is made up of more than one device, each physical device that makes up the workstation must be defined as a component of that workstation.

Each component can have an IMS logical terminal (LTERM) associated with it. One LTERM is associated with each component. A user-written MPP can address specific components of a workstation through the appropriate LTERM.

IMS assumes that all input is from the first LTERM in the component list that passes the necessary operational and security checks. Message switches, broadcast messages for specific logical terminals, and data replies to transactions are directed to the component associated with the specified output LTERM.

LTERM naming

To facilitate the use of the CHNG call, define a convention for naming LTERMs. One method is to set the LTERM name to be a combination of the workstation and component identifiers.

Example: IMS considers the 4704 keyboard and display system components as one component, the 4706 magnetic stripe reader as another component, and the 4710 receipt printer as yet another component. Thus, workstation 100 can have three components: WS100DS (4704), WS100MS (4706), and WS100RP (4710). Such a standard permits an MPP to interrogate the I/O PCB (LTERM name field) to identify the workstation and then to specify the proper alternate PCB for output using the CHNG call.

For a multiple-component terminal, such as a SLU P, you must ensure that IMS creates enough queues to service the terminal. You can use an exit routine or a specific user descriptor for that unique logical unit to ensure that IMS creates the required number of queues.

Output component selection

IMS system definition allows a workstation to have a maximum of four output components.

If more than four output components are required for a workstation, or multiple LTERMs are not desired, a user-defined MPP-to-workstation protocol and data format must be provided.

Workstation components are assigned an identification number: X'01', X'02', X'03', and X'04'. For Finance terminals, the identification number is based on the order in which they are defined. For SLU P terminals, the identification number defaults to Comp¹, (Program1, Basic). IMS nonqueued system messages have no specific destination and are directed to the first component.

An IMS Transaction Input edit routine can be used to append the station's node name to the input message. The MPP can reference the node name to determine which LTERM name to use for output. For terminals created using the Extended Terminal Option (ETO) feature, append the user name to the input message because the user could be signed off.

To return a reply message to a workstation operating in terminal response mode, or to a conversational transaction, the MPP must insert the reply to the I/O PCB. However, for some stations, that reply message might be intended for a component other than the one associated with the LTERM specified in the I/O PCB. Because the destination of the I/O PCB cannot be modified by the CHNG call, an alternate PCB type is required. The alternate PCB type must fulfill these response requirements and also be modifiable. The PCB type, called the response alternate PCB, can be defined using PSBGEN. This PCB can be used instead of the I/O PCB to return reply messages to conversational transactions and to stations operating in terminal response mode. It can also be defined as modifiable. This allows the CHNG call to be used with this PCB to select the appropriate destination for the reply. When the CHNG call is used, the LTERM specified in the call must be assigned to the same physical terminal as the LTERM specified in the I/O PCB. If this is not the case, a status code is returned on the ISRT call.

Related tasks:

Chapter 6, "Administering the Extended Terminal Option," on page 87

Input component determination

Proper relationships between input and output components can be established using the NAME macro during IMS system definition, or by using an ETO user descriptor and the Signon exit routine.

This relationship allows the terminal to specify its input component and causes output to be returned to a component that is defined during IMS system definition. Proper definition and use of input components can reduce or eliminate the need for LTERM naming conventions, MPP change calls, and inserts to alternate PCBs.

A user can establish a connection for SLU P usage by defining a component for each of the operators, devices, or processing requirements controlled by a remote application program.

IMS performs input component selection for the remote logical units and assumes that all input is from the first LTERM in the list that passes the necessary operational and security checks. For input from SLU P, input component selection is performed based on the component indicated in an optional input message header. If no function management header is received, IMS assumes the input is to be associated with the LTERM for component 1. Message switches, broadcast messages for specific LTERMs, and data replies to transactions are directed to the component associated with the specified output LTERM.

Terminal-response mode

Terminal-response mode is a mode of operation that can be defined for transactions, users, and terminals attached to IMS.

When a SLU P station is operating in terminal-response mode, all operations are stopped between the workstation and IMS from the time IMS receives a transaction until IMS receives acknowledgment that the reply message has been received by the workstation. For normal IMS output messages, this acknowledgment is the

receipt of the DR2 response requested for recoverable transaction output. For Fast Path output messages, this acknowledgment is the receipt of the next input message or an RTR command. For MFS-paged output messages, this acknowledgment is the receipt of the next input message, MFS NEXTMSG or NEXTMSGP control commands, or a requested DR2 response (if not Fast Path output) on the last page of the message. Output caused by a nonrecoverable transaction requests an exception DR2 response and does not require a response from the workstation. Terminal-response mode can reduce the processing required by the controller application program.

Terminal-response mode can be selected by an installation during IMS system definition or with an ETO user descriptor. Your system can be defined as forced, negated, or transaction-dependent:

- If it is defined as forced, every input transaction is in terminal-response mode.
- If it is defined as negated, no input transaction is in terminal-response mode.
- If it is defined as transaction-dependent, terminal response mode is determined on a transaction-by-transaction basis. The use of transaction-dependent terminal-response mode can make the controller application programs more complex, because they must handle communication protocols resulting from both forced and negated terminal-response mode environments.

Terminal-response mode can only be invoked by valid transactions, not by message switches, IMS commands, VTAM commands and indicators, or MFS control requests. Transactions that are unacceptable (for example, because of a security violation or an invalid transaction code) cannot invoke terminal-response mode.

When a workstation is operating in terminal-response mode, the following processing occurs after the workstation has established a session with IMS:

1. The workstation sends an input transaction.
2. IMS places the workstation in terminal-response mode.
3. IMS passes the transaction to a message processing program (MPP).
4. The MPP processes the transaction.
5. The MPP returns a reply, using either the I/O PCB or an response alternate PCB.
6. IMS returns a DRx response, if one is requested.
7. IMS sends the reply to the workstation.
8. The workstation returns a DR2 response, if one is requested.
9. IMS removes the workstation from terminal-response mode.

While a workstation is in terminal-response mode, IMS accepts no input from it, and sends no output to it other than the reply from the MPP. If the MPP abends while processing the input transaction, IMS might send an exception response and the associated IMS error message. Any output that is not in reply to the transaction that initiated terminal-response mode is held in IMS's output queue. After IMS removes the workstation from terminal-response mode, it sends the unsolicited output.

When Fast Path is used, the following processing occurs:

1. The workstation sends an input transaction.
2. IMS places the workstation in terminal-response mode.
3. IMS passes the transaction to a Fast Path message processing program (IFP).
4. The IFP processes the transaction.

5. The IFP returns a reply, using either the I/O PCB or a response alternate PCB.
6. IMS returns DRx, if one is requested.
7. IMS sends the reply to the workstation.
8. An exception DR2 response is required on output; therefore, the workstation sends the next input message or RTR command. If IMS has output for the workstation on the message queue, a definite response is required. The workstation sends DR2.
9. IMS removes the workstation from terminal-response mode.

The master terminal operator can reset a workstation in Fast Path terminal-response mode before the response is returned by issuing the /STOP NODE and /START NODE commands in sequence from the master terminal.

Defining a workstation for terminal-response mode

You must understand the operation sequences of terminal response mode before defining a workstation to operate in terminal-response mode.

Less processor time is required for stations that operate in terminal response mode, and the controller application program that controls such a workstation might also be less complex.

The following considerations also apply:

- Typical data collection applications cannot be performed from stations that operate in terminal-response mode. In this mode, a reply from the MPP is required for each transaction before IMS accepts another transaction. Waiting for this response extends the time of the data entry process.
- IMS does not send a response to an input message until the output reply is available. Thus, when the response is returned, it indicates that the input has been processed and the application has reached a sync point.
- While a workstation is in terminal-response mode, IMS does not attempt to obtain any more input from that station. Master terminal operator intervention is required if an error prevents creation or transmission of a reply. Some conditions can prevent a reply being sent to the workstation. These include:
 - The LTERM stopped.
 - IMS was unable to schedule an MPP (database stopped, MPP stopped).
 - An MPP logic error caused no reply to be returned except for EMH (expedited message handler), which generates a zero-length reply.

If any of these conditions occurs, the workstation is temporarily inoperative. Before the workstation can be used again, the master terminal operator must diagnose and correct the error.

- A response message remaining on the IMS output queue or inserted by the user-MPP after session termination is re-sent after initiation of the next session. If the BID option is specified, this message is preceded by the VTAM BID command. Both the begin-bracket and end-bracket indicators are sent with the message.

System analysts must evaluate these factors for their own system application programs and operating environment.

Related concepts:

“Terminal-response mode” on page 910

Output messages sent while in a between-brackets state

If an IMS output message is sent while the workstation is not protected from output and in a between-brackets state, IMS sends one of two things.

IMS sends either:

- The message with both the begin- and end-bracket (that is, bids for bracket with data) if the workstation is defined with the NOBID option
- The BID command to request permission to begin a bracket if the workstation is defined with the BID option.

A workstation can accept the bracket with any requested DR1 or DR2. A DR1 to the bid causes IMS to send the output message with both begin- and end-brackets.

A workstation can reject the bracket with an appropriate exception DR1 or DR2.

When a workstation is ready for output after rejecting a bracket, it sends the ready-to-receive (RTR) command to request output from IMS.

The following steps show how IMS handles output for a workstation that is in terminal-response mode and that is defined with the BID option when the session is between-brackets:

1. When IMS receives a message switch for a workstation defined with the BID option, IMS places the message in the output queue.
2. After removing the workstation from terminal-response mode, IMS sends a BID command to notify the workstation that output is pending while between-brackets. IMS requests a DR1 response to the BID command.
3. The workstation returns the DR1 response if it is ready to receive the output.
4. IMS sends the output when it receives the DR1.
5. The workstation returns a DR2 response, if requested.
6. If the workstation is not ready to receive the output, it returns an exception DR1 response indicating "bid rejected". If the output is recoverable, IMS returns the message to the queue and waits until the workstation indicates it is ready to accept the output. If the exception response indicates "RTR will not follow" and the output is nonrecoverable, IMS discards it. If the exception response indicates "RTR will follow," IMS returns the message to the queue (regardless of recoverability) and waits for a VTAM ready-to-receive (RTR) command.
7. When the workstation is ready to receive the output, the workstation sends the VTAM RTR command and requests a DR1 response.
8. If the output message that caused the bid is still available, IMS returns the DR1 response, followed by the output message.
9. The workstation returns a DR2 response (12).

While IMS waits for the RTR command from the workstation, the workstation can perform any type of processing desired. IMS accepts and responds to transactions during this time, but does not transmit any unsolicited output until it receives the RTR command.

If IMS no longer has output available to send when the workstation sends RTR (perhaps because of an intervening /ASSIGN or /DEQUEUE command), or if the data was irrecoverable and IMS discarded it, IMS returns an exception DR1 response and an error message indicating no output is available.

Related concepts:

“Display screen protection for finance stations” on page 917

Related reference:

“Controller or station-detected errors” on page 956

Designing for output messages sent while in between-brackets state

When IMS is between-brackets, it sends the BID command to bid the workstation before actually sending an output message.

If the bid is rejected, the recoverable output message remains queued if the exception sense data indicates an RTR is forthcoming; otherwise, the message is discarded.

If IMS receives an exception response to an irrecoverable output message (inquiry-only transaction), the message is lost, because IMS dequeues the message immediately upon sending it to VTAM.

Related concepts:

“Output messages sent while in a between-brackets state” on page 913

Related reference:

“Controller or station-detected errors” on page 956

IMS Message Format Service

The following topics describe some IMS MFS facilities that specifically apply to the SLU P system.

Related concepts:

 [IMS Message Format Service \(Application Programming\)](#)

 [MFS message formats \(Application Programming APIs\)](#)

Designing MFS for the workstation environment

Both input and output data can be processed by MFS. The exit for the Input Transaction edit routine is available to provide additional editing capabilities.

The availability of MFS to workstations is defined on an individual basis by the system administrator during IMS system definition. When MFS is defined as available, each input or output message can optionally be processed using MFS. The format of messages to be processed by MFS is defined using the IMS-supplied MFS Language utility.

Operations using MFS can be quite different from operations using the IMS Basic Edit facilities. Device formats and operator procedures should be designed carefully, with the objectives of easy use and high-operator productivity.

MFS provides two levels of message formatting: device-level message formatting and distributed presentation management (DPM). Device-level message formatting is for Finance logical units. DPM support is for SLU P logical units.

To prevent null screens, a null record (data length = 0) produced by MFS as the result of a user-specified MFS format description is not sent by IMS.

MID/MOD chaining

MFS input formatting for a workstation occurs when a message input descriptor (MID) name is provided with an input message.

The MFS message output descriptor (MOD) can supply a MID name to be used for formatting the next input message. It is the responsibility of the workstation to supply this MID name when it sends the input message. This results in supplying the MID name.

MID/MOD chaining can be accomplished with little or no intervention by the workstation operator by observing the following procedure during the design of SLU P system application programs:

1. Remove the MID name from the received output message header and save it for use on the next input message.
2. Display or print the output message.
3. Get the next operator input.
4. Add the MID name saved in step 1 to the transaction.
5. Send the transaction to IMS.

Related tasks:

“Activating MFS input formatting for Finance workstations” on page 942

 MID/MOD chaining (Communications and Connections)

MFS output formatting for the SLU P system

MFS can be used to format pages, produce standard headings and footings, and provide forms control. Without MFS, the message processing program (MPP) or the controller application program must provide these functions.

Using MFS can also provide MPP device independence and allow SLU P components to be used for low-volume applications invoked from any of several terminal types. This device independence allows the controller application program to concentrate on high-volume applications, reducing its complexity and maintenance work.

The availability of MFS to a workstation is defined on a station-by-station basis by the system administrator during IMS system definition or ETO logon descriptor definition.

MFS message recovery

Input messages processed by MFS are edited before the message is queued and logged. Output messages are edited immediately prior to transmission; therefore, for recovery purposes, an input message is formatted by MFS and an output message is formatted by the IMS application program.

MFS control functions (Finance)

The operators of devices using MFS can use a number of control functions.

The control functions available to the operators of devices using MFS are:

Page Advance (NEXTPP)

Transmit the next physical page of the current message, if one exists.

Logical Page Advance (NEXTLP)

Transmit the first or only physical page of the next logical page of the current message.

Logical Page Requests

PAGEREQ=nn where nn is the number of the logical page desired. If the request is valid, transmit the data fields defined in the specified logical page of the current message.

Message Advance Protect (NEXTMSGP)

Discontinue printing or displaying the current output message and begin transmitting the next message in the output queue. If none exists, return an exception DR2 to notify the operator.

Message Advance (NEXTMSG)

Discontinue printing or displaying the current output message and begin transmitting the next message, if any.

These control functions can be indicated to MFS by including them in input function management headers or by defining an operator control field within the input data to MFS.

MFS control functions (SLU P)

The following device-level MFS control functions are available to logical units defined as SLU P.

SLU P terminals process the DPM-formatted output messages that specify paging option on the MFS device format. The paging option can be OPTIONS=DPAGE or OPTIONS=PPAGE.

Page Advance (NEXTPP)

If a current message with OPTIONS=PPAGE is defined, transmit the data fields defined in the next presentation page. If a current message with OPTIONS=DPAGE is defined, transmit the data fields defined in the next logical page.

Logical Page Advance (NEXTLP)

Transmit the data fields defined in the next logical page of the current message if one exists and if either OPTIONS=PPAGE or OPTIONS=DPAGE is specified.

Logical Page Request

PAGEREQ=nn, where nn is the number of the logical page desired. If the request is valid, transmit the data fields defined in the specified logical page of the current message.

Message Advance Protect and Message Advance (NEXTMSGP and NEXTMSG)

Discontinue transmitting the current output message and begin transmitting the next message in the output queue.

The control functions of NEXTPP, NEXTLP, NEXTMSGP, and NEXTMSG can be specified in input function management headers or by defining an operator control field within the input data to MFS. Logical page requests can only be entered within the input data or using the operator control field.

MFS paging and BID options

The BID option provides output protection but has a high performance cost (an extra line flow for each message).

When the BID option is specified and the output occurs while between-brackets, IMS waits for a positive response to BID before sending output.

To avoid sending BID while sending MFS-paged output, each input paging request should indicate begin-bracket and change-direction. Each output page then contains only end-bracket and no BID results, because the page is sent while in an in-brackets state.

Display screen protection for finance stations

When a Finance station is defined with the NOBID option, IMS provides support similar to the screen protection function for the IBM 3270 Information Display System.

IMS does not transmit two consecutive output messages to a display component without an intervening input message from the component. This gives the user the opportunity to view and respond to one message before another is displayed.

When IMS transmits an output message to a display component, it marks the component as protected, unavailable for output. IMS does not transmit another output message to it until an input message is received from that component. The input message can be an IMS transaction, a message switch, an IMS command, the VTAM RTR command, or an MFS control request. Upon receipt of one of these, IMS changes the display component's status to unprotected.

When a workstation is defined with the BID option, consecutive messages are not transmitted to a station. Input from the workstation or a positive response to the BID command from IMS must be received after one transmission before the next message can be sent.

When a workstation uses MFS, the screen is protected on a physical-page basis. MFS control requests or input data are used to request additional screens of data.

If a workstation does not use MFS, the screen is protected on a message-by-message basis. The input of any IMS transaction, message switch, or command causes the screen to be unprotected. The VTAM ready-to-receive (RTR) command can be used to request the next output message. If no message is available, or if the node is in a status that does not allow output to be sent (output stopped or quiesced), IMS returns an exception DR1, followed by an error message indicating that no output is available.

If the RTR command is received during the transmission of MFS-paged output, IMS returns an exception DR1 response, followed by an error message indicating an invalid paging request, because IMS cannot determine the MFS control function to be performed.

After an exception response is received by IMS for a current IMS output message to a display component, the screen is unprotected. After an exception response is sent by IMS, any defined display component is automatically marked protected and unavailable for output.

Related reference:

“MFS control functions (Finance)” on page 915

Extended output component protection (SLU P)

When a SLU P is defined with the NOBID option, IMS provides support similar to the screen-protection function for the IBM 3270 Information Display System and display-screen protection for Finance components.

IMS does not transmit two consecutive output messages to an output-protected component without an appropriate intervening input message from the LU. This gives the LU an opportunity to process one message, according to user-defined procedures, before another is sent to the same component. SLU P support allows the output of all components to be protected.

IMS marks the component as protected and unavailable for output under the following conditions:

- When IMS transmits a message to a component defined as PROGRAM2 on the IMS TERMINAL macro
- When IMS transmits a message (or a page of a message) formatted by MFS with a paging option defined

While a component is protected, IMS does not transmit another output message to the component until an input message from the logical unit resets the component to unprotected. Input messages that can reset a component's status to unprotected are:

- An IMS transaction with an FM header
- A message switch with an FM header
- An IMS command with an FM header
- An MFS control request with an FM header indicating a specific component to be unprotected
- Any of these, either without an FM header or with an FM header indicating component zero (which resets protection on all components of a terminal)
- An RTR command

After the component's status is reset to unprotected, IMS sends any available output. If a component is defined as PROGRAM2 and is not described by a device format with paging OPTIONS=DPAGE or OPTIONS=PPAGE, output to the component is protected on a message-by-message basis. The input messages are used to reset output component protection.

If a SLU P component is defined to use MFS DPM and the device format used is defined with the paging option OPTIONS=DPAGE or OPTIONS=PPAGE, the output component is set to protected when a logical page or presentation page is sent to the component. MFS control requests or input data are used to reset output component protection and request additional logical or presentation pages from MFS.

During the transmission of MFS-paged output, the only allowable input is:

- Input containing no message header (resets all components' protection).
- Input containing a header indicating the specific component to which MFS is currently paging. Otherwise, the session is terminated, because IMS cannot simultaneously send output to more than one component of a terminal.

If a SLU P terminal is defined with the BID option, output component protection occurs after every message. In this mode of operation, consecutive messages are

not transmitted to a terminal. Input from the terminal must be received; otherwise, if additional output is queued, a BID command is sent after one transmission before the next message is sent.

If a terminal does not use MFS, screen protection is performed on a message-by-message basis. The input of any IMS transaction, message switch, or command causes the screen to be unprotected.

If the ready-to-receive (RTR) command is received and no messages are available, or if the terminal is in a status that does not allow output to be sent (output stopped or quiesced), IMS returns an exception DR1, followed by an error message indicating no output is available. If the RTR command is received during the transmission of MFS-paged output, IMS returns an exception DR1 response, followed by an error message, indicating invalid paging request because IMS cannot determine the MFS control function to be performed.

After any exception response is sent by IMS, any display component defined for the logical unit is automatically marked protected and unavailable for output.

After an exception response is received for a current IMS output message to a component defined as PROGRAM2 or for an MFS DPM-paged output message, the component to which the output was sent is left unprotected.

After any exception response is sent by IMS, all components defined as PROGRAM2 are automatically marked protected and unavailable for output.

Related reference:

“MFS control functions (SLU P)” on page 916

Input and output editing options (SLU P)

The type of editing subparameter in the COMPTn parameter of the system definition TERMINAL macro and the Extended Terminal Option (ETO) logon descriptor can be indicated at a component-by-component level for input to and output from components defined as SLU P.

The four types of editing provided are:

BASIC

Requests that no deblocking be performed on input to IMS. MFS cannot be used on input or output.

BASIC-SCS1

Requests that deblocking occur when an SNA character string (SCS)-defined new line (NL-X'15') or form feed (FF-X'0C') character is sent to IMS. MFS is not used on input or output.

MFS-SCS1

Requests that deblocking occur when an SCS-defined new line (NL-X'15') or form feed (FF-X'0C') control character is sent to IMS. MFS-SCS1 formats can be used for both input and output.

DPM-An

Allows messages to be formatted using MFS distributed presentation management (DPM). No deblocking on input to IMS is performed. 'n' is a number from 1 to 15.

For DPM, the device type symbolic name is specified to MFS as DPM-Xn, where X is A or B. Message formatting is specified on the TERMINAL

macro for logical units defined as SLU P by using device type symbolic names of the form DPM-An. DPM-Bn device type symbolic names refer to logical units defined for Intersystem Communication (ISC).

If BASIC or SCS1 is specified, each message received or sent by IMS as one or more related transmissions forms a VTAM chain:

- For BASIC, each input transmission to IMS, less any supplied input FM header, is treated as an IMS segment to be presented to MFS or directly to the IMS message processing program.
- For SCS1, an IMS input segment is created at each SCS-defined new line or form feed control character. IMS segments can be created from a portion of, from all of, or from more than one spanning transmission.

For either BASIC or SCS1, each output segment is sent by IMS in a single transmission, unless it and any IMS-appended FM header are larger than the logical unit's buffer size as indicated to IMS by the OUTBUF parameter on the system definition TERMINAL macro or on the ETO logon descriptor. In this case, IMS sends the segment in as many transmissions as required. Each transmission, except possibly the last transmission of a segment, is the maximum-defined output buffer size.

When DPM or MFS-SCS1 is specified, input and output messages can be formatted by MFS DPM or MFS-SCS1, respectively, on a message-by-message basis. Messages that are not to be formatted are edited as previously described for BASIC or SCS1.

Related concepts:

 MFS message formatting functions (Application Programming APIs)

 IMS Message Format Service (Application Programming)

Related tasks:

Chapter 6, “Administering the Extended Terminal Option,” on page 87

Use of responses or brackets to acknowledge recoverable input

To facilitate message resynchronization, IMS allows an input update or a recoverable-inquiry transaction to optionally request a definite response.

IMS allows this request for a definite response if the user specifies the `OPTIONS=OPTACK` for the workstation.

With this operand specified, IMS can acknowledge input with the next output through use of the input and output bracket indicators.

Restriction: `OPTIONS=ACK` is not supported for SLU P terminals created using the ETO feature.

If the `OPTACK` option is defined, performance can be improved if the workstation requests begin-bracket, change-direction, and exception DR1 or DR2 (rather than DR1 or DR2) on recoverable input to IMS. The next output from IMS acknowledges the input by indicating end-bracket only. The type of output messages that are sent from IMS depend on the type of recoverable input, the defined response mode, and the availability of output.

Contention can occur when a session is in a between-brackets state. In this case, IMS sends an unsolicited output message indicating both BB and EB or a BID

command at the same time that the workstation is sending an input message indicating begin-bracket and change-direction. This output message does not acknowledge the input message and can be either accepted or rejected by the workstation. Rejecting an IMS nonrecoverable output message can result in losing the message. If contention occurs when IMS is sending a recoverable IMS output message (output requesting DR2 response), the workstation must send either an exception or a definite response before IMS can receive the input message.

Related concepts:

“Message resynchronization” on page 922

“Output messages sent while in a between-brackets state” on page 913

Related tasks:

Chapter 6, “Administering the Extended Terminal Option,” on page 87

Related reference:

“Fast Path messages with Finance and SLU P” on page 923

Chapter 53, “SLU P message protocols,” on page 939

“MFS control functions (Finance)” on page 915

 **TERMINAL** macro (Communications and Connections)

Message recovery

The process of message recovery is accomplished within IMS by using the IMS system log and the checkpoint and restart routines.

Recovery is possible, because IMS has direct control of the communication link. With a SLU P system, however, this control is shared with the controller and the controller application program. If a network failure, such as a processor failure or an IMS abend, occurs while the controller application program is receiving a message from IMS, the traditional methods of IMS recovery cannot always detect the lost message. The IMS recovery methods are extended for the SLU P systems to include the controller application program for message resynchronization.

The controller application program participates in message resynchronization so that a lost message condition, if any, can be detected and corrected. Failure of the controller application program to perform its responsibilities during message resynchronization results in a loss of message integrity in the system.

Message resynchronization is necessary for a workstation when a session with the workstation is terminated between the time it sends a recoverable message and the time it receives a reply for that message. The type of resynchronization depends on how the message is defined.

For all messages except response mode, response conversational mode, and Fast Path, the input becomes recoverable when it has been successfully enqueued and made available for scheduling. During message resynchronization, IMS indicates the last successfully received recoverable message.

For response mode, response conversational mode, and Fast Path, the input is not made recoverable and restartable until the application reaches its first sync point. If the IMS system fails before this sync point, the message is considered to be nonrestartable. During message resynchronization, IMS indicates whether the message has reached a sync point or needs to be reissued.

If a session with a workstation is terminated after IMS sends a message to the workstation but before IMS receives the response, message resynchronization is necessary for this workstation. The output message for which no response is received must remain associated with this workstation until message resynchronization determines whether the workstation received the message. If the /ASSIGN command is used to move the message to a different workstation, message resynchronization is no longer possible.

Message resynchronization

The purpose of message resynchronization is to guarantee the integrity of messages across sessions.

Message resynchronization occurs at the start of a session unless IMS is cold started. Finance and SLU P sessions warm start using the control blocks created from the original descriptor, even though that descriptor might have been changed or deleted after IMS created the control blocks. Only a cold start ensures that the control blocks that are created represent the new or updated descriptor.

When message resynchronization is necessary because of a network failure, the resynchronization must complete successfully before IMS permits normal data transmission. To initiate message resynchronization, IMS sends the VTAM set-and-test-sequence-numbers (STSN) command. The LU must respond to this command. To be able to respond properly, a copy of the following must be maintained:

- The sequence number of the last request unit (RU) of the last inbound sync-point message that was sent by the LU. The inbound sync-point message is one of the following:
 - The last successful recoverable input message (the one that requested a DR1 or DR2) if the ACK option was defined to IMS for this LU
 - The last input message, if the OPTACK option was defined for this LU
- The sequence number of the last request unit (RU) of the last outbound sync-point message received by the LU. The outbound sync-point message is the last successful recoverable output message. Recoverable output messages are those requesting DR2 responses.

If Fast Path is used, each output message is recoverable and requests an exception DR2. These Fast Path-recoverable messages are identified by a flag within the output message header, rather than by an explicitly requested DR2 response.

Optionally, the LU can maintain a copy of the last inbound recoverable message sent by the terminal. If this is done, the SLU P system can retransmit, after resynchronization, any message not received by IMS.

Restriction: Session initiation and resynchronization caused by an SNA Request-Recovery (RQR) command is not allowed when the node is in response mode and the response reply message is not yet available for output; that is, the input response mode transaction is still queued or in the process of execution. Response-mode transactions are not recoverable or restartable prior to the application sync point; therefore, session input acknowledgment does not occur until input processing is complete.

Session initiation or resynchronization results in session termination while the response mode transaction is in this 'indoubt' state, and IMS is not able to indicate that the associated input sync-point request was committed or backed out. This

temporary error condition is indicated to the master terminal operator by message DFS2081I. An IMS /DISPLAY command can be used to determine when the response mode reply message is available and session initiation can again be attempted. A display status of RESP-INP indicates input is still in process; RESP indicates input processing is complete and output is available for transmission.

Finance and SLU P in an XRF complex

Finance and SLU P sessions in an XRF complex are handled as Class 1, 2, or 3 terminals.

Class 1 attempts to provide a maximum level of transparency and performance by tracking the active IMS session with a standby session on the XRF backup system. Class 2 and 3 support involves terminating the active session if the IMS system fails. At takeover, the alternate IMS system automatically reinitiates the session and automatically signs on the original active system user if necessary. Class 3 support requires manual session restart and signon after XRF takeover.

For information about XRF from a z/OS perspective, see *z/OS Communications Server: SNA Network Implementation Guide*.

Related concepts:

 [Extended Recovery Facility Overview \(System Administration\)](#)

Related tasks:

[“Establishing connection with the XRF complex” on page 906](#)

Fast Path messages with Finance and SLU P

When using Fast Path, you must specify certain options on the TERMINAL system definition macro.

The following options must be specified in the TERMINAL macro:

OPTACK

Allows input messages containing begin-bracket and change-direction indicators to be acknowledged by an end-bracket indicator on the next output message.

FORCERESP or TRANSRESP

Allows input of a response mode message to IMS.

To obtain best performance, Fast Path input messages should not request definite responses but should be coded with begin-bracket, change-direction, and exception DR2 response.

FPACK/NFPACK

Determines if special Fast Path output protocols should be used.

Fast Path transactions are single-segment transactions and are defined as recoverable. Fast Path transactions must be defined as response mode.

Fast Path output messages (Finance)

When defined with the system definition TERMINAL macro option FPACK or an ETO logon descriptor, Fast Path output messages are sent requesting an exception DR2 response.

Fast Path output is sent requesting DR2 response when:

- Queued output is available to be sent to any nondisplay component.
- The system definition `TERMINAL` macro option `BID` or an ETO logon descriptor option `BID` is defined, and queued output is available to be sent to a display component.
- The system definition `TERMINAL` macro option `NOBID` or an ETO logon descriptor option `NOBID` is defined, and queued output is available to be sent to a display component when the Fast Path output reply message is directed to a nondisplay component. The Fast Path output does not result in the display component being screen protected.

Unlike non-Fast Path output, the output message is left outstanding (not dequeued), and the workstation remains in response mode until the workstation sends data, a ready-to-receive (RTR) command, or the DR2 response to cause the message to be dequeued. The RTR command should be considered when no input is to be generated for an abnormally long time (for example, when the terminal operator plans to leave the terminal). The input message from the workstation, the RTR command, or the DR2 response acknowledges that the preceding output has been received and is recoverable; therefore, IMS might dequeue the output message, process any input message, or send any available output.

The system definition `TERMINAL` macro option `NFPACK` or an ETO logon descriptor indicates that the Fast Path output exception DR2 and next input acknowledgment protocol should not be used. In this case, Fast Path output messages are always sent requesting standard recoverable output message response protocols (DR2 response).

When the DR2 response is received, IMS dequeues the output message, removes the terminal from response mode, and sends any available queued output. Queued output is not sent to a display component if the Fast Path output reply message is sent to the same component because of the IMS screen protection support.

Related concepts:

“Display screen protection for finance stations” on page 917

Fast Path output messages (SLU P)

When defined with the system definition `TERMINAL` macro option `FPACK` or an ETO logon descriptor, Fast Path output messages are sent requesting an exception DR2 response when certain events occur.

Fast Path output messages request an exception DR2 response when any of the following occurs:

- Queued output messages are available to be sent to a component defined for the workstation as `PROGRAM1`.
- The system definition `TERMINAL` macro option `BID` or an ETO logon descriptor option `BID` is defined, and queued output is available to be sent to a component defined for the workstation as `PROGRAM2`.
- The system definition `TERMINAL` macro option `NOBID` or an ETO logon descriptor option `NOBID` is defined, and queued output messages are available to be sent to a program. This component is not output protected due to a previous output or this Fast Path message. (That is, multiple components might be protected or reset selectively as defined in this chapter.)

Unlike non-Fast Path output, the output message remains outstanding (not dequeued) and the terminal remains in response mode until the terminal sends data, a ready-to-receive (RTR) command, or the DR2 response to cause the message

to be dequeued. The RTR command should be considered when no input is to be generated for an abnormally long time (for example, when the terminal operator plans to leave the terminal). The input message from the terminal, the RTR command, or the DR2 response acknowledges that the preceding output has been received and is recoverable; therefore, IMS might dequeue the output message, process any input message, or send any available output.

The system definition `TERMINAL` macro option `NFPACK` or an ETO logon descriptor indicates that the Fast Path output exception DR2 and next input acknowledgment protocol should not be used. In this case, Fast Path output messages are always sent requesting standard recoverable output message response protocols (DR2 response).

When the DR2 response is received, IMS dequeues the output message, removes the terminal from response mode, and sends any available queued output for `PROGRAM1` components or non-output protected `PROGRAM2` components defined for the workstation. Output protection is set at the component level during output for `PROGRAM2` components and can be selectively reset for one or more components based on subsequent input.

Also use the `NFPACK` option for all 4730 terminals. This allows any asynchronous output messages to be sent immediately following the acknowledgment of the Fast Path output reply.

Related concepts:

“Extended output component protection (SLU P)” on page 918

Fast Path message resynchronization

If Fast Path is defined, sequence number management and message resynchronization might require change.

Related concepts:

“Message resynchronization” on page 922

Chapter 52. Network operation for SLU P and Finance

The following topics describe how to start an IMS network, how to initiate sessions, the different transaction types, message switching, and IMS commands.

Starting an IMS network

Before a session with IMS can be established for any workstation VTAM, NCP, controllers, and logical units must be active.

Specifically, the following components must be active:

- VTAM and NCP must be active.
- Controllers and logical units that are not activated automatically by VTAM must be activated with a VARY command by the VTAM network operator.
- Controllers must be powered on, and appropriately configured, initialized, and activated for VTAM and NCP.

Making IMS ready

IMS must be made ready to receive VTAM logon (session initialization) requests. Ready IMS by issuing the IMS /START command with the DC keyword.

The /START DC command tells VTAM to pass any queued VTAM logon requests (and any logon requests for workstations known to VTAM as belonging to IMS) to IMS.

The /START DC command activates the following processes:

- Initiates IMS Transaction Manager processing
- Opens the VTAM access method control block
- Enables the IMS VTAM logon exit

Any logon requests received by VTAM before the IMS /START DC command is issued, but after the IMS VTAM access method control block (ACB) has been opened, are queued in VTAM until the /START DC command processing is completed. If VTAM is active when IMS is initialized, and the DFSDCxxx PROCLIB member keyword VACBOPN=INIT, then the IMS VTAM ACB is opened. If DFSDCxxx PROCLIB member keyword VACBOPN=DELAY, then the IMS VTAM ACB open is delayed until the /START DC command is processed.

Session initiation (starting workstations)

A *session* is the logical connection of a workstation to a VTAM application program, such as IMS or the system utilities. A session must be established before data can be transmitted between a workstation and IMS. Message resynchronization is performed during session initiation, unless IMS is cold started.

Session initiation can be requested in one of the following ways:

- The workstation requests session initiation by sending the INITIATE-SELF command. VTAM verifies the command and passes the request to IMS. If the terminal is requesting a session with an IMS XRF complex, the IMS USERVAR name or the MNPS ACB should be used for the application name.

- The z/OS VTAM network operator requests session initiation on behalf of the workstation by using the VTAM VARY command with the LOGON option. VTAM processes the request and passes it to IMS.
- VTAM passes a logon request to IMS for each workstation that is defined to VTAM as belonging to IMS.
- The IMS master terminal operator requests session initiation for a workstation by entering the IMS /OPNDST command.

Regardless of how session initiation is requested, identical processing occurs when IMS receives the request.

Session-initiation transmission sequence

The following list shows the sequence of transmissions that occurs when a workstation requests session initiation.

The numbers relate to major events in the sequence.

1. The controller sends the INITIATE-SELF command. VTAM checks the validity of the command, looking for syntax errors that might have been introduced by either the controller or NCP.
2. If the command is not valid or if VTAM does not find a match when verifying the initiation request, VTAM returns an EXC/DR1 response. EXC/DR1 terminates the initiation request. The SLU P system can retry the initiation request or notify the operator to begin corrective action.

If the command is valid, VTAM returns a DR1 response. VTAM then verifies the content of the resource field in the INITIATE-SELF command. This field must contain the name of the VTAM application program with which the workstation wants to enter a session. VTAM compares the resource field content to its list of active application programs (those programs with an open VTAM ACB).

3. If a match is found, VTAM passes the logon request to the specified application program (IMS).

In an XRF environment, VTAM routes the request to the currently active system.

When IMS receives the request, it compares the workstation node name supplied by VTAM to the node names defined during IMS system definition.

4. If IMS does not recognize the workstation, it issues a VTAM CLSDST command, which causes VTAM to send a procedure error command. A procedure error terminates the session initiation request. The system notifies the workstation operator that the session is denied.

If IMS recognizes the workstation, it issues a VTAM OPNDST command, which causes VTAM to send a BIND command. The BIND command contains the name of the application program (IMS) that issued the OPNDST command.

In an XRF complex that uses USERVAR instead of MNPS, the BIND command from VTAM contains the USERVAR segment in the user data field and the active APPLID in the PLUNAME field.

IMS supplies a set of BIND parameters, which define the communication rules and protocol that must be followed. IMS ignores the parameters of the mode table entry supplied by the workstation or network operator.

The workstation must respond to the BIND command.

Related reading: For more information on the contents of the BIND data, see “Bind Parameters for SLU P and LU 6.1” in *IMS Version 11 System Programming APIs*.

5. If the workstation cannot begin a session, it returns an EXC/DR1. This can occur when either VTAM or IMS is requesting session initiation and the workstation is currently unable to communicate (for example, because it is involved in offline processing).

If the workstation can begin a session, it returns a DR1. When IMS receives the DR1, it performs message resynchronization, if necessary.

6. If message resynchronization is not necessary, or when it is complete, IMS sends a VTAM start-data-traffic (SDT) command.

The workstation is in session and can transmit data to IMS.

Related reference:

 Bind parameters for SLU P and LU 6.1 (System Programming APIs)

Controller application program involvement in message resynchronization

The controller application program must participate in message resynchronization and, in order to do so, must have maintained copies of the sequence numbers of the last recoverable message that was successfully sent to IMS and the last recoverable message that was successfully received by the workstation.

Design considerations

The system application analyst must determine where to maintain the sequence numbers necessary for set-and-test-sequence-numbers (STSN) processing.

Three options are available:

- The controller disk or diskette.

The disk or diskette in the controller is the most reliable method and does not require involvement of the workstation operator if retransmission of a message is required. Either the permanent file or the transient files of the disk or diskette can be used. The REPLACE instruction should be used when writing to the disk or diskette file. REPLACE causes the data to be physically written to the disk or diskette and not just buffered in the controller's control storage. Disk or diskette storage assures sequence number retention across a power failure or a controller failure that can destroy the contents of controller's control storage. However, disk or diskette access and data transfer time can add significant overhead to each recoverable transaction.

- The controller's control storage.

If the controller's control storage is used to retain sequence numbers, a power loss or controller failure would destroy the contents of the controller's storage, making message resynchronization impossible. By not performing message resynchronization, recoverable input and output messages can be lost or duplicated.

- A workstation output device.

Message sequence numbers can be stored on a workstation output component such as the display or the journal printer. However, this method requires workstation-operator intervention to retrieve the numbers during STSN processing.

Sequence number management

An understanding of how sequence numbers are handled during normal message transmission is prerequisite to understanding how and why the SLU P system performs the functions described in this topic.

The management of sequence numbers is shared between the controller, VTAM, and the VTAM application program (IMS). The controller assigns sequence numbers to messages originated by its workstations. VTAM assigns VTAM sequence numbers to messages originated by IMS. To ensure recoverability, IMS maintains a separate copy of the sequence numbers associated with the messages it sends.

Sequence numbers are assigned to each workstation. For each workstation, the controller tracks the last sequence number that it assigns (called the *last-assigned value*), and tracks the last sequence number that it receives (called the *last-received value*). Similarly, IMS tracks the last sequence number that VTAM assigns (called the *last-assigned value*), and tracks the last sequence number it receives (called the *last-received value*).

When the controller issues a request to transmit data, it updates its last assigned value for the requesting workstation, appends the new number to the data, and sends the message. When VTAM receives the message, it merely passes it on to IMS. IMS, in turn, adds 1 to its last received value for that workstation and compares this value with the sequence number of the message it just received. If the two numbers match, IMS processes the message as required. If the two numbers do not match, IMS issues a VTAM CLSDST macro instruction to terminate the session.

When IMS has a message to transmit, it updates its copy of VTAM's last assigned value and sends the message to VTAM. VTAM updates its last assigned value, appends the new number to the message, and sends the message. The controller removes the appended sequence number, updates its last received value, and compares this value with the sequence number that accompanied the message. If the two numbers match, the controller sends the message on to the application program. If the two numbers do not match, the controller returns an exception DR2 to indicate a sequence error. When IMS receives the sequence error indication, it issues a VTAM CLSDST macro instruction to terminate the session.

Set-and-Test-Sequence-Numbers (STSN)

Message resynchronization is initiated by IMS when it sends the set-and-test-sequence-numbers (STSN) command to the workstation.

The STSN command contains a 5-byte data field:

Byte 0 Action code

Bytes 1, 2

Controller sequence number of the last inbound sync-point message IMS received from the workstation

Bytes 3, 4

VTAM sequence number of the last outbound sync-point message IMS sent to the workstation

IMS uses the action code to ask the SLU P system to verify the controller and VTAM sequence numbers. The bits of the action code byte are:

Bits 0, 1

Refer to the controller sequence-number field

Bits 2, 3

Refer to the VTAM sequence-number field

Bits 4, 5, 6, and 7

Reserved

The following values are acceptable for bits 0, 1, 2, and 3 of the action code:

00 IGNORE

Not used by IMS.

01 SET

Set the appropriate sequence number to the value indicated in the sequence-number field.

10 INVALID

Not used by IMS. The SLU P system must return its version of the sequence number in the command response.

11 SET AND TEST

Set the appropriate sequence number to the value indicated in the sequence-number field. The SLU P system must indicate in the command response whether the sequence number values are acceptable.

IMS uses the SET option for the controller sequence number. For the VTAM sequence number, IMS uses either SET or SET AND TEST. SET is used when no acknowledgment to a recoverable output message is outstanding from the previous session. SET is also used when the IMS master terminal operator enters the /DEQUEUE or /ASSIGN commands that cause the last recoverable message to be removed from the queue. SET AND TEST is used if IMS sends a recoverable message to the workstation but does not receive the required acknowledgment prior to session termination. The response sent by the SLU P system to the STSN command indicates whether the workstation received the message.

When the SLU P system receives the STSN command, it must be able to:

- Verify the controller sequence number and arrange to retransmit a message to IMS if required.
- Verify the VTAM sequence number and inform IMS whether the number is acceptable.
- Return a DR1 to IMS and, if required, return a 5-byte data response to the STSN command.

To verify the controller sequence number, the SLU P system compares the number provided by IMS with the number it maintained from the previous session. An equal compare indicates that IMS received all messages sent by the workstation. If the IMS-provided number is smaller, IMS did not receive the last inbound sync-point message sent by the workstation. These two numbers can differ by more than 1 if intervening irrecoverable or chained messages were sent by the workstation when the workstation was defined with the ACK option. A message that was not received by IMS should be retransmitted after message resynchronization is complete. If a copy of the message was maintained, that copy could be sent. Otherwise, the SLU P system can inform the workstation operator that the last inbound recoverable message was not received by IMS.

To verify the VTAM sequence number, the SLU P system compares the number provided by IMS with the number it maintained from the previous session. An equal compare indicates that the workstation has received all messages. If the IMS-provided number is unequal, the workstation did not receive the last outbound recoverable message.

The response to the STSN command indicates the results of the synchronization tests to IMS. The SLU P system must return a DR1 and, optionally, 5 bytes of data. If both the controller and VTAM sequence numbers are acceptable, only the DR1 is required. If one or both of the sequence numbers are not acceptable, a DR1 and the 5-byte data response are required. The format of the data response has the same format as the STSN command:

Byte 0 Action code

Bytes 1, 2

Controller sequence number of the last inbound recoverable message the workstation sent to IMS

Bytes 2, 3

VTAM sequence number of the last outbound recoverable message the workstation received from IMS

The SLU P system uses the action code to indicate the test results. The bits of the action code byte are:

Bits 0, 1

Refer to the controller sequence-number field

Bits 2, 3

Refer to the VTAM sequence-number field

Bits 4, 5, 6, and 7

Reserved

The following values are acceptable for bits 0, 1, 2, and 3 of the action code:

00 RESET

The sequence number in the command is unacceptable and should be reestablished at its previous value. This code should never be returned to IMS.

01 TEST POSITIVE

The sequence number in the command is acceptable. This code must be returned in response to the SET option. If the SET AND TEST option is specified and the SLU P system finds the sequence number acceptable, TEST POSITIVE should be returned.

10 INVALID

The SLU P system has detected an error in the sequence number.

11 TEST NEGATIVE

The SLU P system does not agree with the sequence number presented with the SET AND TEST option.

TEST POSITIVE is the only acceptable response to the SET option. TEST POSITIVE and TEST NEGATIVE are acceptable responses to the SET AND TEST option. INVALID should be used if the SLU P system detects a “should-not-occur” condition when comparing the sequence numbers. An example of this condition is a controller sequence number specified in the command that is higher than the application program's copy of the sequence number.

If IMS receives a TEST NEGATIVE response to the SET option, a RESET response, or an INVALID response, it terminates the session.

When returning any action code except TEST POSITIVE, the SLU P system should use the STSN response to return its version of the sequence numbers. The sequence

numbers can be valuable information during problem determination and debugging. Any value can be returned, because IMS does not use the returned values. A TEST NEGATIVE response is all that is required to indicate that the sequence number is not acceptable.

When IMS receives the STSN response, it sends the start-data-traffic (SDT) command to the workstation. The SDT command allows normal message transmission to begin. If IMS receives a TEST NEGATIVE response to the SET AND TEST option, it sends the SDT command and retransmits the last recoverable message followed by any other available output. When IMS receives a TEST POSITIVE response to the SET AND TEST option, it sends the SDT command, dequeues the last recoverable message, and sends any other available output.

The contents of the STSN command as received from IMS and the resulting actions of the SLU P system and VTAM are summarized in the following table.

Table 161. Set-and-Test-Sequence-Numbers (STSN) summary

Remote system number action code	VTAM number action code	Controller sequence number ¹	VTAM sequence number ¹	Remote system action when STSN received	VTAM action when STSN received
01 set		Last inbound sync point message that IMS received		Set controller sequence number field to value indicated in STSN controller field.	Set VTAM value to value indicated in STSN controller field.
	01 set		Last outbound sync point message that IMS sends, and responded to by station.	Set host sequence number field to value indicated in VTAM field of STSN.	Set VTAM's last assigned value to value in VTAM field of STSN.
	11 set and test		Last recoverable message that IMS sends; no response from station received.	Set host sequence number field to value indicated in VTAM field of STSN.	Set VTAM's last assigned value to value in VTAM field of STSN.

Note:

1. The controller and VTAM sequence number fields can contain any value.

The valid SLU P system responses to the STSN command and the resulting IMS actions are summarized in the following tables.

Table 162. STSN response summary when the action code is the remote system number

Remote system action code value	Response action code must be ¹ :	IMS action when response received
01 set	01 test positive	Send SDT to station

Note:

1. Any other response causes IMS to terminate the session.

Table 163. STSN response summary when the action code is the VTAM number

VTAM number action code value	Response action code must be ¹ :	IMS action when response received
01 set	01 test positive	Send SDT to station
11 set and test	01 test positive	Dequeue last recoverable message and send SDT to station
11 set and test	11 test negative	Send SDT to station and retransmit last recoverable message.

Note:

1. Any other response causes IMS to terminate the session.

Suspending output from IMS

If the controller application program does not want or cannot receive any more output from IMS, the program can send the VTAM quiesce-at-end-of-chain (QEC) command to IMS. IMS returns a DR1 and the VTAM quiesce-complete (QC) command. IMS does not send any more output to the workstation until it receives the VTAM release-quiesce (RELQ) command.

A workstation can also suspend output from IMS by sending the VTAM LUSTATUS or SIGNAL commands to IMS.

Session termination

Session termination releases the workstation from its current logical connection to the VTAM application program. It makes the workstation available for session with other VTAM applications, or communications can be terminated altogether.

There are two types of session termination: orderly and immediate.

Definitions:


- In an *orderly termination*, the workstation is allowed to complete normal processing before the session is terminated.
- An *immediate termination* forces the workstation to terminate the session unconditionally.

Session termination can be invoked by any of the following:

- The IMS master terminal operator
- The VTAM network operator
- The workstation

The following figure summarizes the two types of session termination processing.

Orderly termination



Immediate termination




Figure 152. Termination processing

Each installation must determine specific procedures for session termination. When developing the procedures, be aware of the requirements for session termination processing.

Orderly termination

An orderly termination of the network is invoked by the IMS master terminal operator using the IMS /CHECKPOINT FREEZE, /CHECKPOINT PURGE, or /CHECKPOINT DUMPQ command with the QUIESCE parameter.

QUIESCE is a /CHECKPOINT parameter that initiates shutdown processing for a network. When QUIESCE is specified, IMS sends the VTAM shutdown (SHUTD) command to all workstations and waits until all workstations have completed normal processing and returned the VTAM shutdown-complete (SHUTC) command.

When all workstations have indicated that shutdown is complete, IMS performs checkpoint processing and then issues the VTAM CLSDST macro instruction. CLSDST causes VTAM to send the UNBIND command to all workstations. This command releases the workstations from session with IMS. The controller prohibits any further data transmission from the workstations.

During shutdown processing (the time between the VTAM shutdown-complete and CLEAR commands), the system's workstations should be prepared for any IMS output that results from shutdown processing.

During the processing of an orderly termination, the IMS master terminal operator can terminate the network unconditionally rather than wait for the orderly termination processing to complete. This can be done by invoking an immediate termination.

Immediate termination

The IMS master terminal operator, the VTAM network operator, or the workstation user can invoke an immediate termination of a SLU P workstation.

The IMS master terminal operator uses the /CHECKPOINT command with the FREEZE, PURGE, or DUMPQ parameter, but without the QUIESCE parameter, to invoke immediate termination of the network. IMS issues the VTAM CLSDST macro instruction. CLSDST causes VTAM to issue the UNBIND command to all workstations. This command releases the workstation from session with IMS. The controller prohibits any further data transmission from the workstation.

To terminate workstations or portions of a network selectively, the IMS master terminal operator can use the IMS /CLSDST or /STOP command. /CLSDST and /STOP commands cause IMS to issue the VTAM CLSDST macro instruction for the specified workstations. The /STOP command also prevents further sessions from being established until a /START command is issued for the workstation.

The z/OS VTAM network operator uses the VARY command to terminate a workstation immediately. z/OS VTAM network operator intervention using the VARY command might be required to terminate sessions in which a workstation error prevents an I/O operation from completing.

The workstation uses VTAM session control commands to terminate a session with IMS. Under normal processing circumstances, when the workstation decides to terminate its session with IMS, that workstation should send the VTAM request-shutdown command. IMS completes any input or output currently in progress for that workstation and then issues the VTAM CLSDST macro instruction.

If the workstation detects an error condition from which it cannot recover and then wants to terminate the session, it can send the VTAM terminate-session command. VTAM releases the workstation from the session and notifies IMS accordingly.

Shutting down an IMS network (SLU P)

Shutting down an IMS network may or may not include IMS.

The IMS /CHECKPOINT command is used to invoke termination of the network and a shutdown of IMS. The format of /CHECKPOINT used determines whether the network termination occurs immediately or waits for workstation processing to complete:

/CHECKPOINT FREEZE|DUMPQ|PURGE QUIESCE

Allows all workstations to complete normal processing before shutting down IMS.

/CHECKPOINT FREEZE|DUMPQ|PURGE

Causes immediate session termination for all workstations.

Related tasks:

“Immediate termination”

SLU P messages

The primary function of IMS support for the SLU P system is to receive and transmit data for a workstation and to ensure that the data is processed properly.

Input/output messages can consist of the following data types:

- IMS transactions
- IMS message switches
- IMS commands
- VTAM commands and indicators
- Message Format Service (MFS) control requests

IMS can process nongraphic characters when sending and receiving any of these messages.

Related concepts:

“IMS sensitivity to nongraphic message data” on page 450

Send/receive and bracket protocol

The change-direction indicator and bracket protocol, as defined by VTAM, are used to specify the beginning and end of synchronous transmissions and to control the flow of such transmissions.

Because of its queued input and output processing, IMS does not support multiple related inputs or multiple related outputs. While in an in-brackets state, the following can occur:

- One input message (coded BB/EB)
- One output message (coded BB/EB)
- One input message and an output reply (input coded BB/CD and output coded EB)

The input message and output reply are related only for:

- Fast Path transactions
- Terminal-response mode transactions
- Conversational transactions

For other transaction types, the input message and output reply might or might not be related.

Related concepts:

“Input bracketing protocol” on page 942

“Output bracketing protocol” on page 949

Chapter 53. SLU P message protocols

A single transmission must be used for VTAM commands and indicators and MFS control requests. Single or multiple transmissions can be used to send IMS transactions, commands, and message switches.

Multiple transmissions for one message are required if the workstation's transmission buffer is not large enough to hold all of the data to be sent. Multiple transmissions can also be used if the workstation user wants to segment the data. Multiple transmissions are logically related through the concept of chained messages. Each transmission of a multiple transmission message is identified by its position in the chain—that is, first-in-chain, middle-in-chain, last-in-chain as shown in the figure in “Output messages” on page 943.

In a SLU P system using chained input messages, the chaining indicator is used to specify the chain position of the transmission; it should not be set for an unchained message. The chaining indicator must be turned on for the first-in-chain transmission, turned off for each middle-in-chain transmission, and turned on again for the last-in-chain transmission.

When MFS is defined, input messages can be processed by the MFS. For input formatting, MFS does not distinguish between possible input components but assumes that all input comes from the 4701/4702 application program.

General format of input function management headers (Finance)

IMS transactions, message switches, and commands are in the standard IMS format (transaction code, logical terminal name, or command verb, followed by a blank and text) unless MFS or a Physical Terminal Input edit routine is used to edit the message into the standard format.

The first or only transmission of IMS transactions can have a variable-length FM header. If MFS is not defined for the workstation, the header is ignored. If MFS is defined, the header can contain either an MFS control request or a MID name. For MFS control requests, only the first 2 bytes of the FM header are required.

Byte 0 Header length including the length byte (binary). Bytes 0 and 1 constitute an MFS control request.

Byte 1 Message description (binary).

Byte 2 MID name length (binary).

Bytes 3-10

MID name (1 to 8 bytes) (EBCDIC).

IMS uses the header, but removes it before sending the message to the MPP or MFS. When IMS receives a message with a header, it interrogates the length of the message. If the message length is 2 bytes, IMS assumes that the message is an MFS control request. The MFS control request is indicated in the message description (byte 1).

Related reading: For more information on using the PAGEREQ function, see *IMS Version 11 Application Programming APIs*.

Input message descriptor byte (Finance)

If the message length is more than 2 bytes, IMS interrogates the message description to determine whether an MFS MID name is present. The format description specified by the MID name is used to format the message.

The format of the message descriptor byte is:

- Bit 0** Reserved
- Bit 1** Page advance requested (NEXTTPP)
- Bit 2** Message advance requested (NEXTMSG)
- Bit 3** Message advance protect requested (NEXTMSGP)
- Bit 4** Next logical page requested (NEXTLP)
- Bit 5** Reserved
- Bit 6** Format description that is indicated by the MID name length and MID name fields should be used to format this input message.
- Bit 7** Reserved

Bits 1, 2, 3, and 4 (MFS control requests) are mutually exclusive and are examined only if the message length is 2 bytes. Bit 6 must be set on if the input message is to be formatted by an MFS MID name. If bit 6 indicates an MFS MID name is present, bits 0 through 5 must be zero.

General format of input function management headers (SLU P)

IMS transactions, message switches, and commands are in the standard IMS format (transaction code, logical terminal name, or command verb, followed by a blank and text).

If a message is in nonstandard format, MFS or a Physical Terminal Input edit routine can be used to edit the message into the standard format. The first or only transmission of IMS transactions can have a variable-length FM header. If MFS is not defined for the terminal, the header is ignored. If MFS is defined, the header can contain either an MFS control request or optional MFS fields used to invoke MFS formatting. IMS uses the header in these cases, but removes it before passing the message on to the MPP or MFS. Incorrect header specifications or formats can cause the session to terminate. SLU P uses header type X'42'.

The format of the input FM header is:

- Byte 0** Header length including the length byte in binary (first byte)
- Byte 1** Header type (must be X'42')
- Byte 2** Message descriptor 1 (flag byte is binary)
- Byte 3** Message descriptor 2 (flag byte is binary)
- Byte 4** Input component identification (binary)

Optional MFS fields:

For DPM:

Bytes 5 and 6

Version ID (in binary) if bit 0 of byte 3 is on.

Bytes 7-15

MID name length, including length byte (1 byte), followed by the MID name (1 to 8 bytes). If bit 0 of byte 3 is off (version ID not included in the FM header), the MID name length and the MID name are in bytes 5 through 13.

For SCS:**Bytes 3-13**

MID name length (in binary), including length byte (1 byte), followed by the MID name (1 to 8 bytes).

When IMS receives a message with a header, it determines the length of the message. If the message length is 5 bytes, IMS assumes that the message is an MFS control request. The specific MFS control request is indicated in message descriptor 1 (byte 2). If the message length is more than 5 bytes, IMS examines the message description bytes to determine whether an MFS version ID and a MID name are present. When both are present, the version ID must immediately precede the MID name length. (Version ID does not exist for SCS1.) The format description specified by the MID name is used to format the accompanying input message, and the version ID is used to validate the format description level. If no FM header is provided, no version ID is provided, or a version ID of zero is provided, MFS bypasses the validity check on the format description level.

Related concepts:

 Version identification function for DPM formats (Application Programming APIs)

Input message descriptor bytes (SLU P)

Message descriptor 1 (byte 2 of the FM header) has the following format.

- Bit 0** Reserved
- Bit 1** Page advance requested (NEXTTPP)
- Bit 2** Message advance requested (NEXTMSG)
- Bit 3** Message advance protect requested (NEXTMSGP)
- Bit 4** Next logical page requested (NEXTLP)
- Bit 5** Reserved
- Bit 6** Format description indicated by the MID name length, and MID name fields should be used to format this input message
- Bit 7** Reserved

Bits 1, 2, 3, and 4 (MFS control requests) are mutually exclusive. Bit 6 must be set on if the input message is to be formatted by the MFS MID name provided. When bit 6 indicates a MID name is present, bits 0 through 5 must be zero.

Message descriptor 2 (byte 3 of the FM header) has the following format:

- Bit 0** Version identifier for MFS DPM. Bit 0 must be set on if bytes 5 and 6 of the FM header contain the version ID.
- Bit 1-7** Reserved

Input component identification (SLU P)

The input component identification is used for the input component selected. The component identification can be a value between 0 and 4, matching the ICOMPT specified on the NAME macro during IMS system definition or on an ETO user descriptor.

If a component identification is zero, or if no FM header is sent, the input is associated with component 1. Further, IMS uses the input component identification to reset component output protection. If a value of 0 (zero) is sent, output protection is reset on all components of the terminal. If a value 1 through 4 is sent, output protection is reset on only the specified component.

Related concepts:

“Component definition” on page 909

Related tasks:

Chapter 6, “Administering the Extended Terminal Option,” on page 87

Input bracketing protocol

A workstation in session with IMS must indicate begin-bracket (BB) or both begin-bracket and end-bracket (EB) on both the first transmission of a chained message (multiple transmission) and each transmission of unchained messages.

The end-bracket indicator places both IMS and the workstation between brackets and in a contention state. If only the begin-bracket indicator is sent, the workstation should send a change-direction (CD) indicator on each of the following:

- The last transmission of a chained message (multiple transmission)
- Each transmission of unchained messages

The CD indicator is necessary, because IMS cannot correlate multiple input messages. The CD indicator is implied if it is not specified.

If BB and CD are indicated, IMS indicates EB on its next output to the workstation. If IMS detects an error while receiving chained input on which only BB is specified, it returns an EB and an IMS error message.

Restriction: For SLU P terminals, a CD or EB indicator must accompany each input chain; otherwise, the session terminates.

Related concepts:

“Output messages” on page 943

Related reference:

“Error handling” on page 955

Activating MFS input formatting for Finance workstations

When MFS is used, input messages can be processed by the message and format descriptors.

When IMS receives an input message, IMS Basic Edit is performed unless a MID name accompanies the message. The MID name can be supplied by including it either in the input function management header or in the beginning of the message text with the required MFS escape characters (//). When the MID name is present,

MFS edits the message using the specified MID and its associated device input format (DIF). The MID name is provided by either the operator or the remote application program.

Output messages

Output messages from IMS can be one of several different types.

The different types of output messages from IMS include:

- Data replies to recoverable or irrecoverable input transactions
- Data replies to IMS commands
- Message switches
- VTAM indicators
- IMS system messages
- Broadcast messages
- A null (length=0) message containing end-bracket to return the workstation to between-brackets and to a contention state

All messages from IMS to a workstation are sent in a single transmission unless:

- The MPP, a message switch, a command, or MFS provides a multi-segment output message.
- The workstation's read buffer is too small to hold the single-segment output message provided by the MPP or MFS.
- A broadcast message is multi-segment.

Each segment of a message is sent in a single transmission whenever possible, that is, when only-in-chain is indicated. However, if a message segment exceeds the size of the receiving workstation read buffer, the segment is divided into as many transmissions as required until the complete segment is sent. Multi-segment output messages are handled like chained input messages: Each segment is identified appropriately as first-in-chain, middle-in-chain, or last-in-chain.

Exercise care when defining the sizes for IMS message queue data sets and workstation output buffers. Incorrect specification can result in multi-segment output chains from IMS. Output buffer sizes that are too small can result in multiple transmissions for large IMS segments. Message queue data set sizes that are too small can result in multiple IMS segments being created for large message processing program (MPP) inserts.

Related reading: For information on the MSGQUEUE and TERMINAL macros, see *IMS Version 11 System Definition*.

Definition: *MFS-paged output* is each physical page of a message destined for a display device and sent as if it were a complete message.

The following figure shows the relationship between transmissions sent to IMS, segments produced by the MPP, and segments transmitted by IMS.




Figure 153. Chained message interaction between a Finance Communication system and an IMS MPP

After successfully receiving a message, the SLU P system must interrogate the read-type field (SMSCRT), the read-flags field (SMSCRF), and, optionally, the read-flags field extension (SMSCRE) to determine the type and characteristics of the message received.

Related reading: For more information on these fields, see *IBM 4700 Finance Communication System: Controller Programming Library, Volume 3: Communication Program*.

If a non-MFS-formatted output segment is sent in multiple transmissions and the SCAN option is defined for that workstation, IMS does not allow a 4701/4702 device's select control, position-control, or write-control sequence to be split across transmissions. When transparent write-control sequences are split into two transmissions, the data length on the first transmission is modified, and the proper write control characters are inserted at the beginning of the second transmission.

Editing for device control sequences does not occur for non-MFS-formatted output if the NOSCANS option is specified during IMS system definition or on an ETO logon descriptor. If MFS formats the output segment, the device-control sequence (select and position) editing always occurs. Each transmission of the segment, except the last, ends in a three-character, null-horizontal-position sequence to suppress the automatic new-line function of the physical SLU P system output device.

MFS Distributed Presentation Management output (SLU P)

IMS sends the entire output message as a single chain of one or more related transmissions. For MFS DPM output using a device format defined with paging `OPTIONS=DPAGE` or `OPTIONS=PPAGE`, IMS sends each logical or presentation page in a single chain of one or more related transmissions as if it were a complete message.

For MFS DPM output using a device format defined with paging `OPTIONS=DPAGE` or `OPTIONS=PPAGE`, IMS sends an output function management header as part of the first or only transmission of each chain (logical or presentation page) of the message. The remainder of the transmission contains data fields (DFLDs) up to the defined record length. This function management header contains the DPM version ID, DPAGE or PPAGE name, and a MID name if specified.

When a forms literal is defined in the DPM format specification for paging `OPTIONS=DPAGE` or `OPTIONS=PPAGE` (`FORMS=` parameter on the DEV statement), IMS precedes the first logical or presentation page with an only-in-chain transmission consisting of a function management header containing only the forms literal. No version ID, DPAGE name, or PPAGE name is provided. Output components are protected at the end of each output chain; this prevents IMS from sending subsequent output, or logical or presentation pages (chains), until an appropriate input message or control request is received.

For MFS DPM output using a device format defined with paging `OPTIONS=MSG`, IMS sends the entire message as a single chain of one or more related transmissions. The function management header is sent as the first or only transmission and contains the DPM version ID, format name, and any MID name or forms literal defined in the MFS format specification. In this case, as for non-DPM output, component protection is provided only if the destination component is defined as PROGRAM2.

For DPM, the device type symbolic name is specified to MFS as DPM-Xn, where X can be A or B. Message formatting is specified on the TERMINAL macro or on an ETO logon descriptor for logical units defined as SLU P by using device type

symbolic names of the form DPM-An. DPM-Bn device type symbolic names refer to logical units defined for Intersystem Communication.

General format of output function management headers (Finance)

IMS appends a header to the first or only-in-chain transmission of all messages sent, except the IMS null message and VTAM commands and indicators.

This header describes the type of message that follows and, if appropriate, indicates the output component that is to receive the message. The header can contain optional MFS information.

The length of the output function management header varies from 3 to 29 bytes. The format is:

Byte 0 Header length including length byte (binary)

Byte 1 Message description (binary)

Byte 2 Output component identification (binary)

Bytes 3-28

MFS data (length byte: binary; names: EBCDIC)

Output message descriptor byte (Finance)

The bits of the message descriptor byte (byte 1), when set on, have the following meanings.

Bit 0 This message is an IMS system message or a broadcast message.

Bit 1 This message is formatted by MFS.

Bit 2 An MFS system control area (SCA) has requested device alarm for this message. (The controller application program can take any appropriate action, such as turning on a terminal's light.)

Bit 3 This message is a nonqueued message (no specific destination). Bit 0 is also set.

Bit 4 Reserved.

Bit 5 This message is Fast Path-recoverable output.

Bit 6 A MID name is present in the MFS data field.

Bit 7 A FORMS name is present in the MFS data field.

Output component ID byte (Finance)

The output component identification (byte 2) contains a value from X'01' to X'04'.

This value identifies the terminal component to which this message is directed. Output component identifiers are assigned during IMS system definition, on an ETO user descriptor, or by using the Signon (DFSSGNX0) and the Output Creation (DFSINSX0) exit routines. These identifiers are specified by the sender of the message (either IMS or the message processing program).

Related concepts:

“Output component selection” on page 909

MFS data bytes (Finance)

Bytes 3 through the end of the header can contain MFS data, and one or two fields can be present. Bits 6 and 7 of the message descriptor byte indicate whether these fields are included.

The first field, if present, is 2 to 9 bytes long and contains a 1-byte length indicator including length byte and a 1- to 8-byte MFS MID name. This MID name field is present if the message output description contains the name of the MID to be used for the next input message, and should be returned to IMS by the controller program. The second field, if present, is 2 to 17 bytes long and contains a 1-byte length indicator including length byte and a 1- to 16-byte FORMS name. FORMS name identifies the special form required for this message. Before printing the message, the application program should ensure that the form is in place and that page size and forms alignment are established.

General format of output function management headers (SLU P)

IMS appends a header to the first or only-in-chain transmission of all messages sent, except the IMS null message and VTAM commands and indicators.

This header describes the type of message that follows and, if appropriate, indicates the output component that is to receive the message. The header can also contain optional MFS information. SLU P uses header type X'42'.

The length of the output function management header varies from 5 to 42 bytes. The format is:

Byte 0 Message header length including length byte (binary)

Byte 1 Header type (X'42')

Byte 2 Message descriptor 1 (flag byte is binary)

Byte 3 Message descriptor 2 (flag byte is binary)

Byte 4 Output component identification (binary)

Bytes 5-41

MFS data (length byte: binary; names: EBCDIC)

Output message descriptor bytes (SLU P)

The bits of the message descriptor 1 (byte 2) have the following meaning when set on.

Bit 0 This message is an IMS system message or a broadcast message.

Bit 1 This message is formatted by MFS.

Bit 2 Reserved.

Bit 3 This message is a nonqueued message (no specific destination). Bit 0 is also set.

Bit 4 Reserved.

Bit 5 Fast Path-recoverable output.

Bit 6 A MID name is present in the MFS data field.

Bit 7 A forms name is present in the MFS data field.

The bits of message descriptor 2 (byte 3), when set on, have the following meaning:

Bit 0 This output is formatted by MFS DPM. The version ID might or might not be present, and one of the following is present in the MFS data field:

- DPM format name
- Logical page name (DPAGE statement label)
- Presentation name (PPAGE statement label)

Bit 1 Reserved.

Bit 2 At the end of this chain, the component indicated in byte 4 of the output function management header is protected.

Bits 3-7
Reserved.

If message descriptor 2, bit 2, is set on, IMS does not send another output chain until an appropriate input occurs. This input can be an input message, an MFS control request, or a READY-TO-RECEIVE (RTR) command.

Related concepts:

“Extended output component protection (SLU P)” on page 918

MFS data bytes (SLU P)

Bytes 5 through the end of the header can contain up to four MFS fields, depending on settings in message descriptor 1, bits 6 and 7, and message descriptor 2, bit 0.

If bit 0 is on, the first field, if present, is a 2-byte MFS DPM version identification. This field is present for all MFS DPM formatted output. The field can contain a version ID, or the value 0 (zero). In either case, one of the following fields (second, third, or fourth field) is present.

The second field, if present, is 2 to 9 bytes long and contains a 1-byte length indicator, including length byte followed by a 1- to 8-byte MFS MID name. This MID name field is present if the message output description contains the name of the MID to be used for the next input message, and should be returned to IMS by the controller program.

The third field, if present, is 2 to 9 bytes of data name, as defined by the paging option for the MFS device format.

The fourth field, if present, is 2 to 17 bytes long and contains a 1-byte length indicator, including length byte followed by a 1- to 16-byte user-specified MFS forms literal. The forms literal identifies the special setup or forms required for this message. The user must define the procedure to be followed at the terminal upon receipt of the forms literal.

The message format name (FMT statement label) is present if OPTIONS=MSG is specified; the logical page name (DPAGE statement label) is present if OPTIONS=DPAGE is specified; the presentation page name (PPAGE statement label) is present if OPTIONS=PPAGE is specified. This field appears in all output (message descriptor 2, bit 0) from DPM.

The content and format of the function management header for DPM formatted output can be influenced when defining the MFS device format through the

HDRCTL=FIXED or HDRCTL=VARIABLE option. If HDRCTL=VARIABLE is specified, each MFS header field is variable in size. If HDRCTL=FIXED is specified, the following MFS header fields are sent and padded to their maximum size.

- The MID name field in the output header is padded with trailing blanks for the maximum length of 8 bytes. Eight blanks are sent if no MID name is specified to format the next input through the user-supplied message output description.
- The message format name is padded to a maximum length of 6 bytes. The logical page name or the presentation page name is padded to a maximum length of 8 bytes.

Output bracketing protocol

For output messages, IMS specifies begin-bracket and end-bracket on single-segment output messages and on the first transmission of a multi-segment message.

If the input message specified begin-bracket and change-direction, the next output IMS sends specifies only end-bracket. If IMS detects an error when receiving chained input on which only begin-bracket is specified, it returns end-bracket and an IMS error message.

In the case of Fast Path, conversational, or response-mode transactions, the output message is the reply to the previous input message. In other cases, the output message might or might not be related to the previous input message.

IMS sends an only-in-chain output message (no FM header and data length=0), requesting exception DR2, and indicating FM header end-bracket when no other output is immediately available. This occurs in the following situations:

- An input message from a workstation defined to IMS as “negated terminal-response mode” specifies begin-bracket and change-direction.
- An input message from a workstation defined to IMS as “forced terminal-response mode” specifies begin-bracket and change-direction and is not an IMS transaction.
- An input message from a workstation defined to IMS as “transaction-dependent terminal-response mode” specifies begin-bracket and change-direction and is not a response-mode transaction.

After receiving exception-response sense codes 0811 and 0812, IMS also sends an only-in-chain output message (no FM header and data length=0), requesting exception DR2, and indicating FM header begin-bracket/end-bracket when no other output is immediately available.

Related reference:

“Error handling” on page 955

Activating MFS output formatting for SLU P

MFS output formatting occurs when an output message has an associated message output descriptor (MOD).

An output message has an associated MOD when any of the following occurs:

- The MPP supplies the name of a MOD (MOD name) with the output message.
- The input message was processed by a message input descriptor (MID) whose definition specified a MOD name for output formatting.
- The output message is a message switch from a device using MFS editing.

If no MOD is associated with the output message, standard IMS output editing occurs.

Input formatting for a workstation only occurs when a MID name and the message itself are sent to IMS. Specifying the next MID name is allowed in the message output description. Controller involvement is required to assure the proper MID is used, because IMS cannot directly control the receipt of input. Therefore, while IMS is formatting an output message, the workstation can be sending an input message. Input sent while IMS is processing output is queued by VTAM until IMS completes output processing. If the next MID name is specified, the input queued by VTAM appears as if it were the result of the current output. Unpredictable results occur if the internal MID name is used to format the message. IMS avoids this by sending the MID name as part of the input message.

IMS informs the controller that the name of the next MID is specified, and sends the MID name requested and the output message to the workstation. The controller then saves the MID name, displays the output message, reads the next operator input, adds the saved MID name to the transaction, and sends the transaction to IMS.

Response requests (Finance)

Output messages from IMS require specific responses. IMS commands, broadcast output, message switches, and non-Fast Path transactions defined to IMS as update-inquiry or recoverable-inquiry require a DR2, except replies to all but the last physical or logical page of non-operator, logical-paged MFS output.

Transactions defined to IMS as irrecoverable-inquiry or replies to operator logical-paged MFS output require an exception DR2. Except for the last physical page of a message whose MOD does not specify PAGE=YES, MFS-paged output requests an exception DR2 reply (regardless of how the transaction is defined), because an explicit MFS control or data input must follow each page before another page of output is allowed; otherwise, the message is dequeued.

Fast Path output messages are sent requesting an exception DR2 response regardless of the transaction type. The next input message from the workstation or an RTR command provides acknowledgment that the preceding output has been received and that IMS might dequeue the message.

Related reference:

“Input response requirements” on page 951

Response requests (SLU P)

Except for replies to all but the last physical or logical page of non-operator, logical-paged MFS output, output messages from IMS require specific responses. IMS commands, broadcast output, message switches, and non-Fast Path transactions defined to IMS as update-inquiry or recoverable-inquiry require a DR2.

Because an explicit MFS control or data input must follow each page before another page of output is allowed (or the message is dequeued), transactions defined to IMS as irrecoverable-inquiry and replies to operator logical-paged MFS output require an exception DR2. MFS-paged output, except the last physical page of a message whose MOD does not specify PAGE=YES, requests an exception DR2 reply regardless of how the transaction is defined.

Related reference:

Input response requirements

The following table summarizes input requirements for requesting responses for all types of data transmission between a workstation and IMS.

Table 164. Input response requirements by message type

Data type	Responses accepted	Responses accepted
	(ACK option)	(OPTACK option)
Update transaction	DR2 ¹ Exception DR1 or exception DR2 ³	DR1 or DR2 Exception DR1 or exception DR2 ^{1, 2}
Recoverable-inquiry transaction	DR2 ¹ Exception DR1 or exception DR2 ³	DR1 or DR2 Exception DR1 or exception DR2 ^{1, 2}
Irrecoverable-inquiry transaction	DR1 or DR2 Exception DR1 or DR2 ¹	DR1 or DR2 Exception DR1 or exception DR2 ^{1, 2}
Fast Path transaction ⁴	N/A	DR1 or DR2 Exception DR1 or exception DR2 ^{1, 2}
IMS message switch	DR2 ¹	DR1 or DR2 Exception DR1 or exception DR2 ^{1, 2}
IMS command	DR1 or DR2 Exception DR1 or DR2 ¹	DR1 or DR2 Exception DR1 or exception DR2 ^{1, 5}
VTAM command or indicator	DR1 ¹	DR1 ¹
SLU P system MFS control request	DR1 or DR2 Exception DR1 or DR2 ¹	DR1 or DR2 Exception DR1 or DR2 ¹

Notes:

1. Recommended or required response.
2. Fast Path users should always use OPTACK option with exception DRx.
3. With CD only for the commands /DIS, /RDIS, and /FOR.
4. ETO users should always use OPTACK option.
5. With Change Direction (CD) only.

Output response requirements

The following table summarizes the response to be requested by IMS for all types of output messages.

Table 165. Output response requested by message type

Output data type	Response requested
Update, Recovery	DR2
Inquiry, Recovery	DR2
Inquiry, Irrecovery	Exception DR1 and DR1 ³
Fast Path (Recovery)	Exception DR2 ¹ and DR2 ²
Last MFS Page (when PAGE=YES is not specified on the MOD)	Refer to above output data types for response request
Nonlast MFS Page (or all pages when PAGE=YES is specified on the MOD)	Exception DR2 and DR2 ²

Table 165. Output response requested by message type (continued)

Output data type	Response requested
IMS Command Replies: /FORMAT, /DISPLAY, /RDISPLAY Other Commands	DR2 Exception DR2 and DR2 ²
Test Mode Output	Exception DR1 and DR1 ³
Broadcast or Message Switch Output	DR2

Notes:

1. Next input message or RTR provides acknowledgment if exception DR2 is requested.
2. DR2 is requested when one of the following occurs:
 - The message is sent with BB (bidding with data).
 - Messages on the IMS message queues are waiting to be sent to the workstation.
 - The NFPACK option is defined on the system definition TERMINAL macro or on an ETO logon descriptor.
3. DR1 is requested when the message is sent with BB (bidding with data).

IMS transaction types

Transactions are the most common data type sent from a workstation to IMS. IMS supports two kinds of transactions—update and inquiry.

An update transaction can modify a database. An inquiry transaction can look at data in a database but cannot change or update it. Transactions are defined as update or inquiry during IMS system definition.

An additional attribute is defined for inquiry transactions—recoverable or irrecoverable. Recoverable-inquiry transactions are always recoverable no matter which element in the network fails. Irrecoverable-inquiry transactions are not recovered following an I/O error condition or IMS system restart.

All update transactions are recoverable.

All Fast Path transactions must be defined as recoverable, but can be either inquiry or update.

A workstation can be defined to handle one or both types of transactions. Other decisions and processing might be required of the controller or controller application program when both recoverable and irrecoverable transactions are handled. The decision to define a transaction as recoverable or irrecoverable requires an evaluation of the advantages and disadvantages that each transaction type offers to the individual operating environment.

The following topics describe the inquiry message types and the responses required for each.

Recoverable-inquiry transactions

To ensure that a recoverable transaction can be recovered, the workstation must perform certain required functions.

To ensure that a recoverable transaction can be recovered, the workstation must:

- Request a DR1 or DR2 on input to IMS, or, if the OPTACK option is defined, optionally request an exception DR1 or exception DR2 and specify begin-bracket and change-direction on input to IMS.
- Maintain the input message sequence number, and, optionally, a copy of the input message until the DR1 or DR2 is returned, or, if the OPTACK option is specified, until a reply message containing an end-bracket is returned.
- Return any DR2 requested by IMS, or, if Fast Path, ensure that another input message or RTR command is sent after having accepted responsibility for an IMS output message by either logging or writing out the data.

To ensure that a recoverable transaction can be recovered, IMS:

- Requests a DR2 response, or, if Fast Path output and no output messages are waiting to be sent, requests exception DR2 and waits for subsequent data or RTR as acknowledgment of the output.
- Maintains the message sequence number and a copy of the message until a DR2, input data, or RTR, if Fast Path, is returned.
- Returns any requested DR1 or DR2 after having accepted responsibility for the message by either logging or writing out the data.

If a failure occurs between the sending of a recoverable message and receipt of the acknowledgment or reply message, the sender cannot determine whether the message reached its destination. During the restart procedure, IMS uses the VTAM STSN command to inform the controller of the sequence numbers of the last inbound sync-point message IMS received and of the last outbound sync-point message IMS sent. Any messages that have not been received can then be retransmitted.

When recoverable messages are being sent, only one recoverable message can be outstanding at a time. This means that the sender should send one message and wait for the response or reply before sending another. IMS reads a message, places the message on the input queue, and returns the DR1 or DR2 response or reply before it accepts another message.

Irrecoverable-inquiry transactions

IMS treats an irrecoverable transaction in the same manner as a recoverable transaction, except that all processing required to achieve recoverability is eliminated. As a result, irrecoverable transactions require less processing time but can be lost in the event of a failure (for example, line failure, processor failure, queue failure) in the network.

An irrecoverable transaction need not request a DR1 or DR2. Because recoverability is not guaranteed, the receiver of the message is not required to acknowledge receipt. An irrecoverable transaction should, however, request an exception DR1 or DR2. The exception DR1 or DR2 requires fewer line transmissions than DR1 or DR2, because under normal circumstances no response is returned.

Verifying IMS receipt of irrecoverable messages

The SLU P system does not request that IMS acknowledge receipt of irrecoverable messages. However, if the system is transmitting multiple irrecoverable messages, indicating both begin-bracket and end-bracket, and wants to verify their receipt, it can send the VTAM CHASE command.

VTAM indicators must request a DR1. When the SLU P system receives the normal definite response 1, all messages preceding the CHASE command have been received by IMS.

IMS message switches

IMS treats a message switch that is sent or received just like a recoverable transaction.

For this reason, the following applies to message switches:

- A message switch must request a DR1 or DR2 if the workstation is defined with the ACK option. If the OPTACK option is specified, a message switch optionally requests an exception DR1 or DR2 and should specify begin-bracket and change-direction on input to IMS.
- A message switch is recovered by IMS.
- The sequence number of a message switch is used in message resynchronization.

IMS commands

The network system analyst decides whether a workstation is allowed to enter IMS commands and, if so, which commands are allowed.

IMS does not require that the controller request a specific response when sending an IMS command. DR1, DR2, and exceptions DR1 and DR2 are allowed; exception DR1 or DR2 is recommended.

IMS processes its commands and returns an IMS message when it finishes processing the command. Therefore, when a DR1 or DR2 is requested, IMS returns the DR1 or DR2, followed immediately by the message confirming command completion.

The additional line transmission required to send either the DR1 or DR2 and the command completion message can be eliminated if the controller requests exception DR1 or DR2.

If exception DR1 or DR2 is requested, the command completion message acknowledges receipt of the command and indicates that command processing is complete. If the ACK option is defined for the workstation, the VTAM sequence number of the command does not participate in message resynchronization as it does when only DR1 or DR2 is requested. If the OPTACK option is defined, all input participates in message resynchronization.

VTAM commands and indicators

When a VTAM command is sent to IMS, a DR1 must be requested.

When IMS sends a VTAM command, it requests a DR1. If the BIND or BID command is sent, the response must be either a DR1 or exception DR1. If the STSN command is sent, a DR1 is required.

The controller responds to all other VTAM commands and indicators for the workstation.

MFS control requests

When Message Format Service (MFS) is used, MFS control requests can be used to display paged messages and to control the display component screen.

The control request is specified in the input function management header.

DR1, DR2, or exception DR1 and DR2 can be requested. Using exception DR1 or DR2 is recommended, because the additional line transmission required to send DR1 or DR2 and the completion message can be eliminated if the controller requests exception DR1 or DR2.

Error handling

The following topics describe the procedures used by IMS and required of the controller or the controller application program to handle failures resulting from transmission or protocol errors.

IMS-detected errors

Whenever IMS detects an abnormal send or receive condition from VTAM, NCP, or the controller, it terminates the session. When an error on a received message is detected, IMS protects any display component and returns an exception DR1 or DR2 that includes 4 bytes of sense information.

Bytes 0, 1

System sense field

Bytes 2, 3

User sense field

System sense field (SSENSE)

Bytes 0 and 1 contain one of the following values.

X'0800'

The user sense field contains the user message from the user message table.

X'0819'

The user sense field contains the IMS message number 290 (in binary). No output is available. This value is set only for a non-MFS response to a previously received RTR command.

X'0826'

The user sense field contains the IMS message number.

User sense field (USENSE)

IMS uses the user sense field to pass on the number (converted to binary) of the appropriate error message.

For example, an invalid transaction code generates the message:

DFS064 DESTINATION CANNOT BE FOUND OR CREATED

If IMS receives an invalid transaction code, it returns an exception DR1 or DR2 response with X'0040' in the user sense field.

In the XRF complex, the alternate system sends SSENSE=X'0826', USENSE X'0F15', and a DFS3861I SYSTEM TAKEOVER OCCURRED message if an input transaction is lost across the XRF session takeover. You must then retransmit the last input record.

IMS binds each session with unconditional bracket termination. When IMS sends an exception DR1 or DR2 to input, it must remain in a send state (regardless of the current bracket state), and the workstation must go to a receive state. Both IMS and the workstation must go to a between-brackets state if the input message in

error specified end-bracket. IMS then places the workstation between-brackets and into a contention state by sending the entire IMS error message as a single chain with the appropriate bracket indicators.

An IMS error message can have multiple segments, with a maximum message length of 32,000 bytes. However, all current IMS error messages that are sent to the workstation are single-segment and less than 132 bytes. These messages adhere to the FM header and transmission rules for all output messages.

The generation of the user sense data is dependent on the standard IMS error message format. The session can be terminated if any user Output edit routine modifies this format.

Related concepts:

“Considerations for controller application programs for XRF systems” on page 901

Controller or station-detected errors

Whenever the controller detects an error on a message from IMS, or simply cannot accept the message at that time, an exception DR2 that includes 4 bytes of sense data is returned.

Bytes 0, 1

System sense field

Bytes 2, 3

User sense field

System sense field

Bytes 0 and 1 should contain either X'0802' (RECOVERABLE ERROR) or X'0811' (BREAK).

X'0802'

RECOVERABLE ERROR—If the message is recoverable, IMS returns it to the output queue and retransmits it following the next input from that workstation or whenever additional output for that destination is queued. Retransmission of chained messages begins from the first-in-chain transmission. Because any open bracket is closed upon receipt of the exception DR2, the retransmitted message contains end-bracket only or both begin-bracket and end-bracket indicators (depending upon whether input occurred or whether the bracket indicators were specified). Irrecoverable messages have been dequeued and cannot be recovered and retransmitted.

X'0811'

BREAK—Cancel output. The output message is terminated. This message is dequeued from IMS and is not retransmitted; it is not subject to recovery. Any remaining output messages that are queued for the workstation are sent. If additional output is not available after receiving this sense code, IMS sends a null output message.

X'0812'

TEMPORARY ERROR—Resource not available. IMS action same as for X'0802'.

X'0813'

TEMPORARY ERROR, BID OR BRACKET REJECTED—NOT READY TO RECEIVE follows. IMS action is the same as for X'0802'.

X'0814'

TEMPORARY ERROR, BID OR BRACKET REJECTED—READY TO RECEIVE follows. IMS action is the same as for X'0802'.

X'081C'

STATION/COMPONENT DOWN—If byte 3 indicates a valid component (X'01' through X'04'), IMS marks the indicated component as inoperable. If byte 3 indicates X'00' or an invalid component (other than X'01' through X'04'), IMS interprets this as a session termination request and terminates the session. If the message in progress is recoverable, IMS returns it to the output queue and retransmits it when the workstation or component is again operable. Irrecoverable messages cannot be recovered and retransmitted.

Any other value for bytes 0 and 1 are handled as a request to terminate the session.

If the BID option is defined, the controller application program should use either X'0813' or X'0814' in an exception DR1 to the BID command; otherwise, the session is terminated because of an invalid response.

The X'0813' sense code results in the message being returned to the IMS output queue if it is recoverable, or dequeued if it is irrecoverable. Any recoverable data returned to the queue causes the bid to be sent again following the next input from the workstation, or whenever additional output for that destination is queued.

The X'0814' sense code causes a recoverable message to be returned to the IMS output queue, or a irrecoverable message to be dequeued. IMS then waits for an RTR command. The workstation can now send input to IMS. Until an RTR is received, IMS does not initiate any output when in a between-brackets state. If IMS is left in an in-bracket send state (the last input carried BB/CD), IMS sends output with EB to terminate the current bracket.

Related concepts:

“Output messages” on page 943

User sense field

This field is used when the system sense field is set to X'081C'. Byte 2 is not inspected by IMS and, therefore, can contain any value.

Byte 3 contains the hexadecimal identification (X'nn') of the component to which the output message is directed. IMS marks that component as inoperable and continues to send output to other operable components of the workstation. If X'00' or an invalid component is specified, the session is terminated.

IMS binds each session with unconditional bracket termination. This means that both IMS and the workstation go to between-brackets and contention states for all supported sense codes except X'0813' and X'0814'. For these two sense codes, IMS goes to between-brackets and receive states; the workstation goes to between-brackets and send states.

VTAM logical unit status (LUSTATUS) command

If the SLU P system detects a component-down or workstation-down condition, and the SLU P system is not currently receiving a message from IMS, it can send the VTAM LUSTATUS command to notify IMS of the condition.

Four bytes of data, identical in format to the exception DR1 or DR2 for STATION/COMPONENT DOWN, must accompany the LUSTATUS. When IMS receives the LUSTATUS, it marks the indicated component inoperable or terminates the session.

The workstation can also notify IMS when the component is operable by sending the VTAM LUSTATUS. The system sense field must equal X'0001', and the user sense field byte 3 must equal a value from X'01' through X'04'. IMS resets the inoperable condition for the indicated component and sends any output queued for it.

All other values for LUSTATUS are considered requests to terminate the session.

VTAM ready-to-receive (RTR) command

The following list summarizes the previously defined IMS functions performed upon receipt of an RTR command.

- Any outstanding Fast Path output message is dequeued and IMS removes the workstation from terminal response mode.
- Any output-protected components are marked unprotected and available for output.
- Output that is suspended because of a X'0814' sense code on a previous exception response is now allowed.
- If IMS output is available to be sent, IMS returns a DR1 response, immediately followed by the output message.
- If no IMS output is available, or if the workstation is in a status where no output can be sent (output stopped or quiesced), IMS returns an exception DR1 followed by an error message indicating no output is available.
- If an RTR is received during the transmission of MFS-paged output, IMS returns an exception DR1 followed by an error message indicating invalid paging request.
- Any exception response returned by IMS causes any display component defined for the logical unit to be protected and unavailable for output.

VTAM CANCEL command

The VTAM CANCEL command is used by IMS, and should be used by the SLU P system to terminate a chained message in progress that has not completed.

IMS sends the CANCEL command to a workstation in these situations:

- The IMS master terminal operator issues the IMS /DEQUEUE command and IMS has not yet sent the last-in-chain transmission of a message.
- IMS receives an exception DR1 or DR2 to a transmission of a message and has not yet sent the last-in-chain transmission.

The controller or controller application program should use the CANCEL command in these situations:

- When an exception DR1 or DR2 is received from IMS prior to sending the last-in-chain transmission of a message.
- When multi-RU chained input to IMS must be canceled prior to sending the last-in-chain transmission of a message. In this case, if the message to be canceled contains only the begin-bracket indicator, the VTAM CANCEL command can include the end-bracket indicator, the change-direction indicator, or no additional indicators.

VTAM request-recovery command

If, during its normal operation, a workstation detects an error condition that does not terminate the session with IMS but does cause the workstation to resynchronize with IMS (for example, diskette failure in the controller or a program check), the application program can initiate message resynchronization.

Initiation of message resynchronization is done by sending the VTAM request-recovery (RQR) command. IMS responds with a CLEAR command followed by the set-and-test-sequence-numbers (STSN) command. Message resynchronization follows and is the same as that described under "Message recovery" on page 921.

Restriction: Session initiation and resynchronization caused by an SNA Request-Recovery (RQR) command is not allowed when the node is in response mode and the response reply message is not yet available for output; that is, the input response mode transaction is still queued or in the process of execution. Response-mode transactions are not recoverable or restartable prior to the application sync point; therefore, session input acknowledgment does not occur until input processing is complete.

Four bytes of data can accompany a SIGNAL command that is sent to IMS. The first two bytes are system-signal data. The last two bytes are user data that is ignored by IMS. The first two bytes are handled as follows, with all other values being ignored:

X'0000'

No specific action is taken. This data can, however, cause IMS to send any available output for operative, unprotected components if the workstation is idle due to a previous exception DR1 or DR2 sent to IMS. This input does not reset the output component protection or display protection.

X'0001'

This data is treated as an attention signal that causes IMS to stop sending and to wait for input at the end of the current output message. This type of signal can be used before sending either input or required DR1 or DR2 or exception DR1 or DR2 responses to IMS.

Part 12. Appendixes

Notices

This information was developed for products and services offered in the U.S.A.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America*

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who 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
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to

IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS," without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com/legal/copytrade.shtml)[®] 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 <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Programming interface information

This information documents Product-sensitive Programming Interface and Associated Guidance Information provided by IMS, as well as Diagnosis, Modification or Tuning Information provided by IMS.

Product-sensitive Programming Interfaces allow the customer installation to perform tasks such as diagnosing, modifying, monitoring, repairing, tailoring, or tuning of this software product. Use of such interfaces creates dependencies on the detailed design or implementation of the IBM software product. Product-sensitive Programming Interfaces should be used only for these specialized purposes. Because of their dependencies on detailed design and implementation, it is to be expected that programs written to such interfaces may need to be changed in order to run with new product releases or versions, or as a result of service.

Product-sensitive Programming Interface and Associated Guidance Information is identified where it occurs, either by an introductory statement to a section or topic, or by a Product-sensitive programming interface label. IBM requires that the preceding statement, and any statement in this information that refers to the preceding statement, be included in any whole or partial copy made of the information described by such a statement.

Diagnosis, Modification or Tuning information is provided to help you diagnose, modify, or tune IMS. Do not use this Diagnosis, Modification or Tuning information as a programming interface.

Diagnosis, Modification or Tuning Information is identified where it occurs, either by an introductory statement to a section or topic, or by the following marking: Diagnosis, Modification or Tuning Information.

Privacy policy considerations

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.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

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 <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

Bibliography

This bibliography lists all of the publications in the IMS Version 11 library, supplemental publications, publication collections, and accessibility titles cited in the IMS Version 11 library.

IMS Version 11 library

Title	Acronym	Order number
<i>IMS Version 11 Application Programming</i>	APG	SC19-2428
<i>IMS Version 11 Application Programming APIs</i>	APR	SC19-2429
<i>IMS Version 11 Commands, Volume 1: IMS Commands A-M</i>	CR1	SC19-2430
<i>IMS Version 11 Commands, Volume 2: IMS Commands N-V</i>	CR2	SC19-2431
<i>IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands</i>	CR3	SC19-2432
<i>IMS Version 11 Communications and Connections</i>	CCG	SC19-2433
<i>IMS Version 11 Database Administration</i>	DAG	SC19-2434
<i>IMS Version 11 Database Utilities</i>	DUR	SC19-2435
<i>IMS Version 11 Diagnosis</i>	DGR	GC19-2436
<i>IMS Version 11 Exit Routines</i>	ERR	SC19-2437
<i>IMS Version 11 Installation</i>	INS	GC19-2438
<i>IMS Version 11 Licensed Program Specifications</i>	LPS	GC19-2439
<i>IMS Version 11 Master Index and Glossary</i>	MIG	SC19-2440
<i>IMS Version 11 Messages and Codes, Volume 1: DFS Messages</i>	MC1	GC19-4232
<i>IMS Version 11 Messages and Codes, Volume 2: Non-DFS Messages</i>	MC2	GC19-4233
<i>IMS Version 11 Messages and Codes, Volume 3: IMS Abend Codes</i>	MC3	GC19-4234
<i>IMS Version 11 Messages and Codes, Volume 4: IMS Component Codes</i>	MC4	GC19-4235
<i>IMS Version 11 Operations and Automation</i>	OAG	SC19-2441
<i>IMS Version 11 Release Planning</i>	RPG	GC19-2442
<i>IMS Version 11 System Administration</i>	SAG	SC19-2443
<i>IMS Version 11 System Definition</i>	SDG	GC19-2444
<i>IMS Version 11 System Programming APIs</i>	SPR	SC19-2445
<i>IMS Version 11 System Utilities</i>	SUR	SC19-2446

Supplementary publications

Title	Order number
<i>IMS and SOA Executive Overview</i>	GC19-2516
<i>Program Directory for Information Management System Transaction and Database Servers V11.0</i>	GI10-8788
<i>IMS Version 11 Fact Sheet</i>	GC19-2451
<i>IRLM Messages and Codes</i>	GC19-2666

Publication collections

Title	Format	Order number
IMS Version 11 Product Kit	CD	SK5T-7375

Accessibility titles cited in the IMS Version 11 library

Title	Order number
<i>z/OS TSO/E Primer</i>	SA22-7787
<i>z/OS TSO/E User's Guide</i>	SA22-7794
<i>z/OS ISPF User's Guide Volume 1</i>	SC34-4822

Index

Special characters

- /ASSIGN command
 - cold starting ISC sessions 510
 - ETO limitations 120
- /CHANGE command
 - initiating an ISC session 510
- /DEQUEUE command 510
 - committing ISC output messages 526
- /DISPLAY
 - TMEMBER command 784
 - TRANSACTION command 754
- /DISPLAY command 433, 510
 - pending ISC output 510
 - QSTOP state 433
- /EXIT command
 - ISC conversation mode errors 524
 - terminating a conversation abnormally 403
- /MSASSIGN command
 - assigning physical link to logical link 679
- /MSVERIFY command 710, 711
 - dynamic validation 710
 - error responses 711
- /OPNDST command 111, 120, 428
 - relation to MODETBL keyword 428
- /SECURE OTMA NONE command 802
- /SIGN command for ETO STSN devices 131
- /SSR command 146
- /START command 431
 - making IMS ready 431
 - starting a SLU P network 927
- /START OTMA command 780
- /STOP command 404
- /STOP OTMA command 780

Numerics

- 119ABEND code 823
- 3270 copy command 454
- 3600 terminal
 - supported as Finance or SLU P 899
- 3650 system, supported as SLU P 899
- 4700 terminals
 - supported as Finance or SLU P 899
- 4730 Personal Banking Machine
 - connecting in the XRF environment 906
- 8100 system, supported as SLU P 899

A

- abend
 - CICS 650
 - MSC conversations 700
- access control environment element
 - cached user ID aging value 859
- accessibility
 - features xviii

- accessibility (*continued*)
 - keyboard shortcuts xviii
- ACEE
 - IMS Connect
 - specifying OTMA ACEE aging value 221
 - OTMA, refreshing the ACEE 802
 - refreshing OTMA ACEEs 802
- ACK
 - Send-only protocol 306
- ACK option
 - Finance Communication System 954
 - SLU P 954
- ACK timeout
 - send-then-commit transactions (CM1) 796
- acknowledgements
 - OTMA ACK timeout intervals 795
- administering
 - Fast Path 780
 - MSC 779
 - queue control facility 781
 - XRF 780
- administration
 - ETO 87
 - MFS 442
 - MSC (Multiple Systems Coupling) 695
- advanced program-to-program communication (APPC)/IMS
 - Error Extract Service 60
- affinity
 - definition 382
 - in an MSC-IMSplex configuration 723
- affinity routing 738
- aging value 859
- allocating a session, CICS 629
- alternate facility, CICS 634
- alternate PCB, secondary transaction 396
- alternate response PCB
 - modifying 910
 - use 910
- ALTPCB
 - IMS Connect and shared queues 197
 - OTMA
 - shared queues, retrieving output 785
 - OTMA destination descriptor 772
 - defining 774
 - masking destination names 775
 - shared queues and IMS Connect 197
- ALTPCBE 785
- API (application program interface)
 - explicit 373
 - implicit 373
- API (application program interface), ISC
 - IMS use of VTAM 471
- API (application program interface), ISC (intersystem communication)
 - asynchronous, CICS 609, 632

- API (application program interface), ISC (intersystem communication) (*continued*)
 - synchronous, CICS 609, 632
- APPC
 - outbound LU
 - specifying 57
- APPC (advanced program-to-program communication)
 - APPC/IMS 392
 - communications manager as the 44
 - destination code 393
 - destination structure 393
 - DFSAPPC 393
 - editing and formatting 437
 - IMS messages 392
 - input segment 393
 - LU 6.2 messages 393
 - queuing, SERIAL=YES option 393
 - segments, single or multiple 393
 - structure 393
 - switching 397
 - SYNCLVL=SYNCPCT 44
 - types 392
- APPC (advanced program-to-program communication)/IMS 373
 - API 373
 - explicit 48, 373
 - implicit 48, 373
- APPC/MVS Administration utility
 - example (ATBSDFMU) 55
- application programs
 - CPI-C driven 49
 - modified IMS 49
 - remote standard IMS 50
 - standard IMS 49
- asynchronous output delivery 73
- CPI Communications application
 - program 37
- CPI-C initialization 392
- default conversation
 - characteristics 73
- DFSAPPC 70
 - message switching 70
 - option keywords 71
- discardable messages 48
- Error Extract Service 60
- establishing APPC/IMS 54
 - APPC=Y or N 54
 - system definition 54
- introduction 47
- local service for remote APPC
 - transaction 50, 52
- LTERM interface 53
- LU 6.2
 - devices 373
 - relationship to APPC/IMS 47
 - terminal support 392
- LUADD option keywords 58
- LUs, managing multiple 69
- LUs, reassigning 70

- APPC (advanced program-to-program communication)/IMS (*continued*)
- MOD name 53
 - mode name 57
 - modified IMS application programs 51
 - MSC 62
 - application program failure and transaction recovery 66
 - intermediate IMS failure and transaction recovery 65
 - local IMS failure and transaction recovery 65
 - local transaction discardability 63
 - LU 6.2 recoverability flows 66
 - LU 6.2 sessions and transaction recovery 63
 - MSC links and transaction recovery 64
 - recovering transactions in APPC 62
 - remote IMS failure and transaction recovery 65
 - standard IMS applications 50
 - transaction point of failure 63
 - MSC (Multiple Systems Coupling) processing remote transactions in an IMSplex 728
 - MSC and modified IMS applications 51
 - network-qualified LU name 68
 - nondiscardable messages 48
 - PARMLIB 58
 - partner LU name 57
 - qualifying LU names 68
 - RACF 74
 - recoverability 433
 - recoverable versus irrecoverable transactions 62
 - remote service for remote APPC transaction 52
 - restrictions 49
 - security 74
 - SIADD 58
 - side information 392
 - mode_name 57, 392
 - partner_LU_name 57, 392
 - TP_name 57, 392
 - sym_dest_name 57
 - Sync_Level options (NONE, CONFIRM, SYNCPT) 49, 51
 - SYS1.APPCSI 58
 - SYS1.APPCTP 54
 - SYS1.PARMLIB(APPCPMxx) example 58
 - timeout service 59
 - TP name 58
 - TP Profile 52
 - ATBSDFMU utility 54
 - definition 52
 - DFSTPPE0 56
 - DFSTPROF 55
 - dialog example 54
 - TSO ICQASRM0 55
 - VSAM data set 54
 - TPN (transaction program name) 57, 71
- APPC (advanced program-to-program communication)/IMS (*continued*)
- transaction retry characteristics 68
 - deadlock 68
 - Fast Path retry conditions 68
 - lock reject 68
 - transaction security 74
 - ACCESS (EXECUTE) 74
 - UACC (NONE) 74
 - XRF 68
- APPC/IMS 37
- and protected transactions 829
- APPC/MVS Administration utility example (ATBSDFMU) 55
- application call processing 145
 - application program interface 471
 - application programming
 - OTMA
 - DL/I calls 806
 - application programs
 - callout function overview 851
 - CICS, with IMS 631
 - converting Finance to SLU P 902
 - correlating responses to asynchronous callout requests 852
 - definition 371
 - existing programs, with ISC 482
 - functions IMS provides 373
 - IMS
 - OTMA 754
 - ISC sample 669
 - remote 371
 - SLU P controller 902
 - XRF considerations, SLU P 901
 - application programs, z/OS
 - using the ODBA (Open Database Access) interface 747
 - using the Open Database Access (ODBA) interface 747
 - application threads 143
 - Application Transparent Transport Layer Security feature (AT-TLS)
 - IMS Connect 212
 - application-data section of OTMA
 - message prefix 883
 - APPLID= parameter, ISC session
 - definition 499
 - architected message 817
 - asynchronous callout requests
 - OTMA configuration overview 845
 - OTMA support overview 845
 - asynchronous messages
 - OTMA, rerouting output 841
 - OTMA, super member 843
 - asynchronous output
 - alternate client ID 339
 - ETO
 - valid destinations 126
 - implementing 332
 - managing 332
 - managing and controlling output messages 333
 - message flow 342
 - OTMA, managing output 840
 - OTMA, purging output 841
 - RESUME TPIPE, alternate client ID 339
- asynchronous output (*continued*)
- retrieving output for another client 339
- asynchronous output processing
- commit mode 339
 - socket type 339
 - sync level 339
 - timer settings 339
- asynchronous processing
- ATTACH EB 476
 - IMS-CICS session 617
 - ISC execution mode 475
 - SCHEDULER FM header 563
- ATCCONxx member (VTAM nodes) 430
- ATCSTRyy member (VTAM start lists) 430
- ATRABCK service 829
- ATRACMT service 829
- ATREINT service 829
- ATTACH FM header
 - bit contents 577, 578
 - format 577, 585
 - MFS 570
 - parameter description 578, 585
 - process initiation 562
 - with CICS 638
- ATTACH FM headers 476
 - EB indicator 476
- ATTIU parameter, FM header length 562
- automated operator programs 386
- autopaged output, synchronous 611
- ## B
- back-end subsystem
- CICS 616
 - definition 461
 - IMS 612, 614
 - routing transactions 495
- backout work unit, message resynchronization 507
- balancing group
 - definition 385
 - routing codes and 385
- balancing resource demand in MSC (Multiple Systems Coupling) 696
- bandwidth
 - MSC links 696
- basic edit 437
 - bypassing 449
 - editing options with ISC 448
 - IMS functions 447
 - IMS systems 437
 - input message segments 451
 - input messages 447
 - ISC messages 473
 - non-MFS programs 483
 - output message segments 450
 - output messages 448
 - SLU 1 transparent data 448
- BB (begin bracket) indicator
 - definition 481
 - LUSTATUS command 546
 - use (figure) 907
- begin bracket indicator
 - LUSTATUS command 546
 - use (figure) 907

- BID command 533, 913, 922
 - BID option
 - caused by session termination 957
 - design considerations 913, 922
 - effects
 - display screen protection 917
 - MFS paging 917
 - output messages 914
 - bind
 - IMS-CICS session 647
 - ISC session
 - parallel 507
 - single 507
 - negotiable 506
 - rejected, ISC 554
 - requesting asynchronous process 564
 - BIND race 507
 - BINPDSB1= parameter, BINTRNDS option 448
 - BIS (bracket initiation stopped) command
 - IMS-CICS session 630
 - session shutdown 558
 - BPE (Base Primitive Environment)
 - header data 256
 - header format 250
 - bracket and send/receive management
 - Finance Communication System
 - direction indicators 907
 - protocol 907
 - ISC, how determined 479
 - SLU P
 - direction indicators 907
 - bracket contention
 - invalid paging 920
 - resolving 507, 920
 - bracket initiation stopped (BIS) command
 - IMS-CICS session 630
 - session shutdown 558
 - bracket protocol
 - IMS 533, 937
 - input bracketing 942
 - input messages, ISC 533
 - output bracketing 949
 - output messages, ISC 536
 - bracket rejection
 - Finance Communication System 928, 957
 - ISC 555
 - SLU P 928, 956, 957
 - BREAK code X'0811'
 - actions taken 956
 - buffer
 - MSC (Multiple Systems Coupling)
 - considerations 695
 - MSC (Multiple Systems Coupling)
 - linking 677
 - buffer sizes
 - BUFSIZE parameter 702
 - IMS-to-IMS sessions 494
 - ISC 500
 - MSC physical links 702
 - buffers
 - BUFSIZE parameter 702
 - MSC links 702
 - buffer size and format in bandwidth mode 715
 - buffers (continued)
 - MSC links (continued)
 - buffer size and format in non-bandwidth mode 716
 - determining buffer size 714
 - format of link buffers 714
 - MSC links and bandwidth 696
 - MSC non-VTAM links 702
 - MSC VTAM links 702
 - sizes for MSC links 702
 - small buffer devices 452
 - BUFSIZE parameter for MSC LU 6.2 702
- ## C
- caching scheme, user ID 800
 - Callable Interface (C/I) 887
 - introduction to 887
 - callout
 - asynchronous
 - correlating a response 852
 - routing responses to IMS Connect client 853
 - routing responses to input terminal 852
 - IMS Connect support 223
 - OTMA
 - routing 849
 - SOAP Gateway 850
 - synchronous
 - acknowledgment messages (ACK and NAK) for IMS Connect 228
 - acknowledgment messages, IMS Connect IRM_F3_REROUT 228
 - acknowledgment messages, IMS Connect NAK 228
 - acknowledgment messages, IMS Connect SYNCNAK 229
 - coding user-written IMS Connect clients 223
 - error responses, returning to IMS Connect 230
 - IMS Connect correlation token structure 257, 262
 - IMS Connect message format 225
 - IMS Connect, returning responses to 307
 - responses, returning to IMS Connect 307
 - RESUME TPIPE call 226
 - RESUME TPIPE error scenarios 227
 - retrieving callout requests 226
 - returning callout responses 229
 - send-only protocol 229
 - timeout, OTMA ACK 795
 - terminals
 - avoiding hung terminals 851
 - callout function 847
 - asynchronous
 - correlating responses to callout requests 852
 - IMS Connect TCP/IP clients 849
 - IMS TM Resource Adapter
 - support overview 849
 - z/OS Sysplex Distributor 849
 - callout function (continued)
 - overview of application programs 851
 - callout messages
 - IMS Connect
 - RESUME TPIPE request 329
 - callout request 847
 - callout requests
 - OTMA
 - implementing asynchronous callout support 848
 - OTMA support 845
 - synchronous
 - OTMA configuration overview 845
 - OTMA support 846
 - callout support
 - asynchronous
 - programming IMS Connect clients 848
 - CANCEL command
 - paging errors 549, 556
 - protocol 541
 - SLU P session 958
 - CANCEL request
 - sender ERP 554
 - candidate printers 454
 - CBresynch 834
 - CBresynch command 862
 - CCTL
 - connections for CICS 745
 - CD (change direction) indicator 533
 - definition 481
 - Finance Communication System 905
 - LUSTATUS command 546
 - send/receive protocol 907, 937
 - SLU P 905
 - soliciting 529
 - chained message communication
 - sequence 943
 - chains, logical terminal 398
 - CHANGE (/CHANGE) command
 - initiating an ISC session 510
 - change direction (CD) indicator
 - Finance Communication System 905
 - send/receive protocol 907, 937
 - SLU P 905
 - channel-to-channel (CTC)
 - MSC (Multiple Systems Coupling)
 - physical link type 677
 - MSC physical link
 - defining 703
 - CHASE command 542
 - CHECK security level 800
 - CICS
 - _CXX_LSYSLIB environment variable 13, 16
 - accessing IMS databases 743
 - applications
 - IVP application for CICS, classic Java APIs for IMS 17
 - running with classic Java APIs for IMS 18
 - running, with IMS Universal drivers 13
 - CCTL connections 745
 - CICSPSB DLL 16

- CICS (*continued*)
 - classic Java APIs for IMS 15
 - configuration
 - DRA tasks 745
 - configuring for 16
 - DFHJVMMPR environment
 - member 13, 17
 - dfjvmpr.props 17
 - DFST2CIC DLL 13
 - DL/I support 744
 - DRA
 - coding PSBs 743
 - DRA configuration 745
 - IMS Universal drivers
 - running applications 13
 - IMS Universal drivers, type-2
 - configuring 12
 - IMS Universal JDBC driver, type-2
 - configuring 12
 - installing type-2 IMS Universal drivers 13
 - ISC
 - application-related concepts 631
 - functions supported 607
 - ISC node
 - defining 620
 - LIBPATH variable 17
 - LIBPATH_SUFFIX variable 13
 - Makefile 13, 16
 - running applications that use classic
 - Java APIs for IMS 18
 - sequential buffering 744
- CICS resource definition 621
- CICS-IMS communication
 - alternate facility 634
 - application coding for 631
 - asynchronous processing flow 617
 - ATTACH parameters 638
 - CICS transactions 628
 - coding function management
 - headers 638
 - coding system definition options 620
 - Command Level API 609
 - configuration 463
 - conversation mode 634
 - defining CICS transactions 628
 - device mapping function 463
 - alternate 634
 - principal 634
 - functions
 - description 609
 - overview 466
 - IMS commands 635
 - initiating sessions 629
 - integrity of session 646
 - LU 6.1 links
 - compatible nodes 622
 - description 621
 - Macro-Level Resource Definition 621
 - multiple links 626
 - Resource Definition Online 621
 - MFS support 644
 - mirror transaction 617
 - MSC links 465
 - passing data to IMS with ISC 467
- CICS-IMS communication (*continued*)
 - preparing CICS tables 621
 - principal facility 634
 - processing flows
 - RECEIVE 616
 - RETRIEVE 617
 - SEND INVITE 612
 - SEND LAST 614
 - SEND/RECEIVE 612
 - START/RETRIEVE 617
 - recovery and restart 645
 - SCHEDULER parameters 641
 - session
 - binding 647
 - initiation 629
 - integrity 646
 - processing outstanding traffic 647
 - reestablishing 646
 - resynchronizing 647
 - sync points 636
 - termination 630
 - sync points 636
 - transactions
 - abnormal termination 632, 649
 - attributes supported 609, 632
 - definition 628
- CICS-ISC installation options 620
- CIMS class 411
- Class 1 terminals
 - overview 416
 - status recovery modes 408
 - XRF support 416
- Class 2 terminals
 - overview 417
 - status recovery modes 408
 - XRF support 416
- Class 3 terminals
 - overview 418
 - status recovery modes 408
 - XRF support 416
- classic Java APIs for IMS
 - WebSphere Application Server for z/OS
 - configuring 27
- CLASSPATH ENVAR keyword 20
- client
 - commands issued 832
 - high-performance access to IMS 756
 - number that can connect to OTMA 756
 - OTMA, in XRF environment 780
- client communications
 - and IBM WebSphere 172
 - local option 172
- Client_Bid resynch 834
- client-bid request
 - message flow 825
 - message prefix contents 825
- CM0
 - OTMA, managing output 840
 - OTMA, purging output 841
 - OTMA, rerouting output 841
 - OTMA, super member 843
 - purge function 299
 - reroute function 301
- CM1 (send-then-commit) transactions
 - ACK timeout 796
- COBOL
 - XML conversion support
 - overview 178
- coding CICS
 - applications for ISC 631
 - system definition macros 620
 - tables 621
- coexistence
 - IMSplex and MSC 721
 - MSC and IMSplex 721
- cold start, recovering ISC sessions 510
- COMM macro statement 499
 - APPLID= keyword 426
 - RECANY= keyword 427
- command
 - /DISPLAY TRANSACTION 754
- CBresynch 834
 - issued by client 832
- REPresynch 834
- REQresynch 834
- SRVresynch 834
- TBresynch 834
- Command Level API, CICS 609
- command recognition character (CRC) 146
- commands 144
 - /CHANGE 164
 - /DBDUMP DATABASE 817
 - /DBRECOVERY AREA 817
 - /DISPLAY SUBSYS 164
 - /DISPLAY TMEMBER 784
 - /SECURE OTMA NONE 802
 - /SSR 146
 - /START AREA 817
 - /START REGION 817
 - /START SUBSYS 144
 - /STOP AREA 817
 - /STOP REGION 817
 - /STOP SUBSYS 144
- asynchronous reply 611
- CICS EXEC
 - asynchronous API 609
 - creating DFC protocols 609
 - functions available 609
 - program flow 612
- CICS-IMS session 635
- CRC (command recognition character) 146
- DEQUEUE TMEMBER 790
- DISPLAY STATUS 790
- DISPLAY TMEMBER 790
- DISPLAY TRACE 790
- IMS
 - parallel ISC sessions 474
 - start system 431
- issuing from an ISC session 474
- LUSTATUS protocol, ISC 548
- PRNs, not specified on IMS
 - commands 583
- recovery at ISC session failure 512
- response requirements 954
- securing with RACF 411
- START TMEMBER 790
- STOP TMEMBER 790
- TRACE TMEMBER 790
- transaction reply 612
- used to start system 927

- commands (*continued*)
 - VTAM
 - BID 913
 - BIS 558, 630
 - CANCEL 541, 958
 - CHASE 542
 - commands and indicators 470, 903
 - LUSTATUS command 546, 958
 - LUSTATUS with CICS 611, 634
 - ready to receive (RTR) 549
 - Request-Recovery 959
 - RSHUT 549
 - RTR 958
 - SIGNAL 558, 959
 - commit
 - LUSTATUS command 548
 - mode 817
 - processing 817
 - sample flows 819
 - work unit 507, 645
 - Commit Mode 291
 - CM1 timeout value 291
 - Commit mode 0 291
 - Commit mode 1 291
 - commit mode 0
 - purge function 299
 - reroute function 301
 - commit-then-send
 - OTMA, managing output 840
 - OTMA, purging output 841
 - OTMA, rerouting output 841
 - OTMA, super member 843
 - purge function 299
 - reroute function 301
 - commit-then-send exchanges
 - OTMA message flow 827
 - OTMA message prefix contents 827
 - commit-then-send flow 819, 827
 - Common Queue Server (CQS)
 - checkpoint data set 379
 - client
 - exit routines 380
 - client, definition 377
 - definitions 377
 - overview 380
 - shared queue environment, in a 379
 - communications
 - IMS
 - prerequisite knowledge xv
 - communications controller 371, 372
 - communications network 369
 - components
 - responsibilities 370, 385
 - Finance Communication System 900
 - SLU P system 900
 - communications, establishing 376
 - COMPINOP state 404
 - component definition
 - LTERM naming 478, 909
 - parameter definitions 537
 - selection
 - input component 478, 909
 - output component 481, 909
 - SLU P input component 910
 - SLU P system 909
 - component protection
 - extended output 918
 - state 405
 - configurations
 - Finance Communication System 900
 - ISC 463
 - SLU P system 900
 - VTAM 430
 - connecting multiple IMS systems 695
 - connection
 - description 471, 905
 - IMS session parameters
 - ISC 504, 506
 - restriction against altering 905
 - XRF complex, establishing
 - connections 906
 - connection factory
 - IMS Universal Database resource
 - adapter 5
 - connections
 - IMS
 - prerequisite knowledge xv
 - contention, bracket 507, 920
 - context 807
 - control block 445
 - MFS 445
 - control block mapping
 - EEVT 151
 - control blocks
 - ISC 475
 - control character 451
 - control function, MFS 437
 - controlling output 452
 - Conversation Abnormal Termination exit
 - routine 694
 - conversation mode
 - errors
 - ISC 523
 - explained 402
 - IMS-CICS 634
 - normal termination, ISC
 - extension 524
 - restriction 495
 - conversational processing 396
 - conversational processing in MSC
 - (Multiple Systems Coupling) 698, 699
 - abend 700
 - planning for 697
 - remote destination verification 698
 - conversational programs
 - IMS Connect support 292
 - conversational transactions 396
 - conversations
 - OTMA and IMS 779
 - protected 807
 - conversion
 - COBOL, XML example 234
 - XML to COBOL example 234
 - coordinating performance information for
 - MSC 718
 - COR token for synchronous callout 262
 - correlation
 - correlating responses to asynchronous
 - callout requests 852
 - correlator token 760
 - coupling facility
 - definition 377
 - CPI Communications 37
 - CPI Communications application
 - program 37
 - abnormal termination 38
 - APSB call 37
 - ATBCMTTP verb 38
 - Backout call 38
 - backout processing 38
 - Commit call 37
 - commit processing 37
 - DB2 for z/OS plan name, use of 37
 - ESS Attach Facility 37
 - in-flight unit of recovery,
 - definition 41
 - indoubt unit of recovery,
 - definition 41
 - normal termination 38
 - programming requirements 40
 - pseudonym files 40
 - recovery 39
 - resolve-in-doubt processing 39
 - return codes 38
 - RTT 37
 - session failure 38
 - SQL calls 37
 - SRRCMIT 37
 - system restart 39
 - two-phase commit process and 41
 - CPI-C (Common Programming Interface
 - for Communications) 373
 - APPC/IMS 392
 - initialization 392
 - side information 392
 - transactions 394
 - CQS (Common Queue Server) 377
 - checkpoint data set 379
 - client
 - exit routines 380
 - client, definition 377
 - definitions 377
 - overview 380
 - shared queue environment, in a 379
 - CRC (command recognition
 - character) 146
 - creating ESMT 149
 - creating work areas for ESAP 159
 - CRGGRM service 829
 - CRGSEIF service 829
 - CSM (complete status message)
 - format 261
 - CTC (channel-to-channel)
 - MSC (Multiple Systems Coupling)
 - physical link type 677
 - MSC physical link
 - defining 703
 - CTXBEGC service 829
 - CTXEINT service 830
 - CTXSWCH service 830
 - custom service, installing 28
 - customizing IMS 776
- D**
- data descriptor FM header
 - IMS use 644
 - input format 585
 - output format 585

- data descriptor FM header (*continued*)
 - use on input 572
 - use on output 572
- data flow 395
 - in an MSC (Multiple Systems Coupling) network 681
 - message switching 397
 - program-to-output terminal 395
 - program-to-program 396
 - terminal-to-program 395
- data flow control (DFC) protocols
 - bracket and half-duplex 533
 - bracketing messages
 - input 533
 - chaining 542
 - commands
 - BID 533
 - ERP PURGE 542
 - error handling 523
 - examples 653
 - exception response 523
 - half-session synchronization 525
 - input messages, backing out 525
 - paged message errors 549
 - PURGE after exception response 542
 - recoverability aided by LWA 525
 - response requirements 525
 - sense codes 557
 - symmetrical session shutdown 558
 - sync point and response 525
 - CICS-IMS 636
 - input 529
 - output 531
- data flow reset, handling errors 554
- data partitioning, MSC (multiple system coupling) 677
- data stream profile, ATTACH FM header 579
- data transmission
 - definite responses, ISC 469
 - error handling 955
 - error recovery procedure 645
 - exception responses, protocol
 - CICS-IMS 635, 645
 - Finance Communication System 903
 - ISC 469
 - messages
 - input 937
 - output 943
 - response requirements
 - IMS commands and indicators 470, 903
 - IMS message switches 529, 954
 - irrecoverable-inquiry
 - transactions 529, 953
 - LU 6.2 application program 432
 - MFS control requests 955
 - recoverable-inquiry
 - transactions 529, 953
 - summary 951
 - VTAM commands and indicators 470, 903
 - SLU P, message data types 937
 - data transparency 473
 - data types, IMS advanced function
 - network 937

- data-set-full condition 453
- database division 697
- database recovery adapter (DRA)
 - setup 747
- database recovery adapter (DRA) start-up table
 - using with ODBA (Open Database Access) 747
 - using with Open Database Access (ODBA) 747
- Database Resource Adapter (DRA)
 - connections
 - configuring 741
- DB2 Attach Facility 137
 - preparing the system 137
- DB2 for z/OS
 - CPI-C
 - plan name, use of 37
 - Java dependent regions
 - accessing DB2 for z/OS 137
- DB2 for z/OS stored procedures
 - classic Java APIs for IMS 20
 - IVP 21
 - Java class libraries for IMS DB, overview 19
 - running with classic Java APIs for IMS 24
- DB2_HOME ENVAR keyword 20
- db2sqljjdbc.properties 23
- DBCTL
 - initializing 745
 - OTMA unsupported 786
- DBFHAGU0 (Input Edit/Routing exit routine) 386
- DCCTL (Data Communication Control) 375
- DCCTL (Data Communications Control)
 - generation 376
 - IMS BTS 376
 - procedures 375
 - TM batch 376
- DD statements
 - DFSRESLB 22
- dead letter queue HWS\$DLQ 346
- DEADQ STATUS 121
- deblocking algorithm, ATTACH FM header 580
- deferred program switch 397
- defining
 - IMS 496
 - ISC node 496
- definite responses
 - Finance Communication System 903
 - response requirements (figure) 951
 - SLU P 903
- demand-paged messages
 - controlling
 - QMODEL FM headers 572
- dependent region connections 166
- dependent regions
 - commit-then-send transactions 794
 - OTMA usage 794
 - send-then-commit transactions 794
- DEQUEUE (/DEQUEUE) command
 - committing ISC output messages 526
- DEQUEUE command
 - FORCESS versus SYNCSESS 510

- DEQUEUE command (*continued*)
 - IMS-to-IMS ISC session 510
- dequeuing messages, implications 510
- descriptors
 - DFSOTMA descriptor 775
 - ETO
 - group logon, definition 82
 - generic ETO logon 82
 - logon
 - ETO 81
 - MFS device 83
 - MSC (Multiple Systems Coupling) 82
 - OTMA client descriptor 771
 - OTMA destination descriptor 772
 - defining 774
 - masking destination names 775
 - OTMA, specifying 771
 - specific logon, definition 82
 - user 82
- destination
 - determination message 776
- destination process name (DPN)
 - ATTACH FM header 581
 - message routing, ISC 484
 - SCHEDULER FM header 581
- destination queue name, ATTACH FM header 585
- Destination Resolution exit routine (DFSYDRU0) 776
- destination system (MSC) 686
- device class control 451
- device input format (DIF) 445
- device output format (DOF) 445
- DFC (data flow control) protocols
 - bracket and half-duplex 533
 - bracketing messages 533
 - input 533
 - output 536
 - chaining 542
 - commands 533
 - BID 533
 - BIS 558
 - CANCEL 541
 - CHASE 542
 - LUSTATUS 546
 - RSHUT 549
 - RTR 549
 - SBI 558
 - SIGNAL 558
 - ERP PURGE 542
 - error handling 523
 - conversation mode 524
 - response mode 523
 - selective receiver ERP 549
 - examples 653
 - exception response 523
 - half-session synchronization 525
 - input messages, backing out 525
 - paged message errors 549
 - PURGE after exception response 542
 - recoverability aided by LWA 525
 - response requirements 525
 - irrecoverable-inquiry
 - transactions 525
 - recoverable-inquiry
 - transactions 525
 - sense codes 557

DFC (data flow control) protocols
(*continued*)

- symmetrical session shutdown 558
- sync point and response 525
 - availability 525
 - CICS-IMS 636
 - exceptions for synchronous input 525
 - input 529
 - irrecoverable messages 529
 - MFS output messages, ISC 527
 - output 531
 - recoverable messages 527
 - requested on input 525
 - requested on output 526

dfjvmp:props file 17

DFS3650I (session status message) 427

DFSAREA0 module 748

DFSAERG0 module 748

DFSAERM0 module 748

DFSAPPC message switching 70

DFSCCMD0 776

DFSCDLI0 module 748

DFSCMC10

MSC

- physical links, defining 703

DFSCMUX0 exit routine 708

DFSCTRNO 776

DFSEMODL macro 149

DFSEWAL macro 152

DFSFEBJ0 (Front-End Switch exit routine)

- special support 461

DFSMSCE0 776

DFSMSxxx control block 700

DFSOTMA descriptor 775

DFSPBxxx

- OTMA parameters 767
 - GRNAME 768
 - OTMA= 769
 - OTMAASY 770
 - OTMAMD 770
 - OTMANM 768
 - OTMASE 769
 - OTMASP 770

DFSQSP0 776

DFSRESLB DD statement 22

DFSIML0 exit routine 455

DFSYDRU0 exit routine 776

DFSYIOE0 exit routine 776

DFSMSG DSECT 855

DFSYPX0 exit routine 776

DFSYRTUX exit routine 776

DIF (device input format) 445

direct-control subsystem 631

directed routing 693, 711

- MSC 693

password not passed across link 711

disabling enforcement

- resource name uniqueness 384
- resource type consistency 384

DISPLAY (/DISPLAY) command

- QSTOP state 433

display screen protection

- BID option 917, 918
- definition 917, 918
- Finance Communication System terminals 917

display screen protection (*continued*)

- MFS 917, 918
- NOBID option 917, 918

Distributed Presentation Management (DPM) 450

- FM headers 568
- option 450
- OPTIONS=DPAGE or PPAGE 568, 945
- output 945
- SLU P supports 914

Distributed Relational Database Architecture (DRDA)

- IMS Connect support 176

distributed transaction processing, ISC 461, 484

distributed two-phase commit 350

DL/I

- OTMA
 - calls used 806

documentation

- exit routines 707
- terminal profiles 425

DOF (device output format) 445

DPM (Distributed Presentation Management) 450

- FM headers 568
- option 450
- OPTIONS=DPAGE or PPAGE 568, 945
- output 945
- SLU P supports 914

DPN (destination process name)

- ATTACH FM header 581
- message routing, ISC 484
- SCHEDULER FM header 581

DR2 response

- exception
 - Fast Path 924, 925
 - irrecoverable-inquiry transactions 951
 - MFS output 950, 951
 - nonrecoverable output 911
 - when MOD does not specify PAGE=YES 934, 952
- requirements
 - MFS paged output messages, not Fast Path 910
 - normal IMS output messages 910
 - recoverable-inquiry transaction 953

DRA

- CICS
 - DRA configuration 745

DRA (database recovery adapter)

- setup 747

DRA (database recovery adapter) start-up table

- using with ODBA (Open Database Access) 747
- using with Open Database Access (ODBA) 747

DRA (Database Resource Adapter)

- connections
 - configuring 741

DRDA (Distributed Relational Database Architecture)

- IMS Connect support 176

DSCA operand, DEV statement 455

duplicate databases 696

dynamic allocation VTAM subpools 478

dynamic terminals 374

E

EB (end bracket) indicator

- definition 481
- LUSTATUS command 546
- use (figure) 907
- used with ATTACH 476

ECSA

- OTMA usage 794

edit

- basic edit during ISC 473
- ISC
 - default editor 473
- MFS
 - CICS 612
 - during ISC 473
- options
 - data communication exit routine 473
 - input and output 473
 - ISC input and output, list 462

editing facilities

- invoking FM headers 561
- ISC overview 473

editing messages

- basic edit and nongraphic messages 451
- bypassing Basic or MFS editing 449
- editing performed by IMS 447
- output segments 450
- transparency option 449

editing options

- COMPTn parameter of TERMINAL macro 501
- OUTBUF parameter of TERMINAL macro 500, 920
- SLU P, COMPTn parameter of TERMINAL macro 919
- types
 - DPM-An 919, 946
 - MFS-SCS1 919
 - SCS1 919

EEVPPEA 157

EEVPEWA 159

EEVT (external entry vector table) 158

EEVT mapping 158

EMH (expedited message handler)

- queue option
 - overview 385

EMH buffer 413

EMHL control region initialization

- parameter 413

encryption

- OTMA restriction 786

end bracket indicator

- definition 481
- LUSTATUS command 546
- use (figure) 907
- used with ATTACH 476

environments supported, IMS 754

ERP (error recovery procedure)

- CICS-IMS session 645
- extended 645
- FM header 562
- format 587
- function management header 645
- implemented by IMS 902
- selective receiver
- sense codes 550

ERP PURGE, after ISC exception response 542

Error Extract Service

- APPC/IMS 60

error handling 955

- BREAK code X'0811' 956
- output messages 956
- CANCEL command 541
- use 958
- when sent by IMS 958
- controller-detected errors
- system sense field 955
- use 956
- user sense field 957

error messages 569

FM header 562

IMS-detected errors 955

IMS-issued error messages in XRF complex 955

ISC

- CICS-IMS session 645
- paging errors, during data flow
- reset state 554
- paging errors, for nonpaged messages 556
- paging errors, for paged messages 549
- selective receiver ERP 549
- sender ERP 554
- sender ERP sense codes 555
- length of message 562
- LUSTATUS command 547, 958
- MFS-detected errors 569
- queuing messages 562
- RQR command 959
- SIGNAL command 959
- SLU P 955
- IMS-detected errors 955
- VTAM logical unit status
- command 634

error messages

- conditions causing
- BB-only specified on input 942
- BB-only specified on output 949
- length 562, 956
- MFS-detected errors 569

error recovery procedure (ERP)

- CICS-IMS session 645
- extended 645
- FM header 562
- format 587
- function management header 645
- implemented by IMS 902
- selective receiver
- sense codes 550

errors

- ISC session termination, causing 554

ESMT (external subsystem module table)

- creating 149
- loading external subsystem modules 158
- work area definitions 152

establishing connection 471

establishing connections for external subsystems 142

ETO

- device characteristics table
- building 108
- IDC0 Trace facility 129
- recovering ETO terminals with XRF 104
- terminals
- planning for growth 92

ETO (Extended Terminal Option) 374

- /SIGN command for ETO STSN
- devices 131
- 3275 devices 94
- 3600/Finance 130
- ABENDU0015 110
- advantages 87
- availability 87
- LTERMs 87
- algorithm 103, 116
- logon descriptor 103
- LTERM allocation 116
- associated printing techniques 118
- asynchronous output 126
- destinations, valid 126
- autologoff 123
- autologon 125
- autosignoff 122
- benefits of using 79
- commands that reset status and release control blocks 116
- commands that retain status 116
- common logon descriptors 101
- concepts
- summary 85
- conversations 131, 397
- customizing 84, 88
- dead-letter queue 126
- default CINIT/BIND user data
- formats 112
- defining
- autosignoff and autologoff
- timer 124
- parameters 121
- deleting control blocks 128
- after logoff 128
- after signoff 128
- delivering output to non-originating terminal 127
- descriptors 84
- added 110
- creating during system
- definition 101
- definition 81
- deleted 110
- generic logon, definition 82
- group logon, definition 82
- introduction to coding 100
- logon 101
- logon, definition 81
- MFS 107

ETO (Extended Terminal Option) (continued)

- descriptors (continued)
- MFS device, definition 83
- MSC (Multiple Systems Coupling) 109
- MSC (Multiple Systems Coupling), definition 82
- specific logon, definition 82
- updated 110
- user 104
- user, definition 82
- using 84
- VTAM TERMINAL macro 101
- device characteristics table 94, 107
- device type, defining 92
- DFS2085 113
- DFS3641W 110
- DFS3645 113
- DFS3649A 115
- DFS3650I 115
- DFS3672 113
- DFSINX0 106
- DFSSGNX0 106
- DFSUSER descriptor 106
- DLQT 126
- dynamic terminal
- definition 79
- dynamic terminals 374
- static terminals, using together 91
- dynamic user, definition 80
- exit routines 84
- Signon exit routine (DFSSGNX0) 118
- using 84
- exit routines, coding 110
- exit routines, list of 110
- guideline selection 102
- logon descriptors 102
- LOGOND parameter 102
- initialization 110
- descriptor validation 110
- DFSINTX0 110
- logging off 128
- logon
- /OPNDST command 111
- dynamic, limiting to specific terminal types 111
- INITOTHER 111
- INITSELF 111
- signon data 111
- USS LOGON 111
- logon descriptors
- creating during system
- definition 102
- NTO, 3600/Finance terminals 104
- LTERM
- creating and reusing control blocks 112
- LTERM with specific destination 115
- LU 2
- devices 93
- screen size and model information 94
- LU 6.1 (ISC) terminals 129
- MFS 98

ETO (Extended Terminal Option)
(*continued*)

- MFS device characteristics table 94, 107
- MFS device descriptors
 - MFSDCT utility (DFSUTB00) 108
- MFSDCT utility 94
- MODETBL on ETO logon
 - descriptor 428
- MSC (Multiple Systems Coupling) 100
 - descriptor 100
 - MSNAME macro 100
 - support 100
- multiple signons 114
- node user descriptor 106
- non-SNA 3270 devices
 - printers and displays 92
 - screen size and model information 94
- NTO devices 94
- output
 - assigning 125
 - inadvertent output data streams 127
- overview 79
- planning 87
 - LU2 95
 - operations 99
 - user IDs 99
 - user queue names 99
- printers
 - defining 119
 - direct printing 118
 - overview 117
 - printer node names 118
 - sharing printers 120
- queue (dynamic user message),
 - definition 79
- RACF 374
- recommendations 106
- requirements 88
- response mode 131
- restrictions 88
- screen definitions examples (non-SNA 3270) 97
 - display (model specified) 97, 98
 - display (screen size specified) 97
 - LU0 video 97
 - model 2 printer 97
- screen size control byte 94
- security 98, 412
 - static versus dynamic terminals 98
- shared printing 120
- signing off, definition 127
- signing on, definition 113
- signon
 - LTERM allocation 113
 - providing signon data 113
- Signon exit routine (DFSSGNX0) 118
- SLU P 130
- SNA commands 111
- special processing modes 116
- starting ETO 110
- static terminal, definition 79

ETO (Extended Terminal Option)
(*continued*)

- static terminals
 - dynamic terminals, using together 91
- storing descriptors 101
- structure
 - terminal, definition 80
 - user, definition 80
- structure, creation and deletion 83
- STSN terminals 130
 - overview 130
 - support for /SIGN command 131
- system definition 101
- terminal
 - definition 79
- terminal-LTERM relationship 115
- terminals
 - using static and dynamic terminals together 91
- terminology 79
- undeliverable data, dead-letter queue 126
 - /DISPLAY STATUS USER command 126
 - /DISPLAY USER DEADQ command 126
- user descriptors
 - creating during system definition 105
 - criteria for selection 105
 - DFSUSER 106
 - installation created 105
 - installation-created 104
 - node user descriptor 104, 106
- VTAM CINIT LUNAME 102
- VTAM considerations 92
 - logon CINIT session control blocks 92
- VTAM PSERVIC parameters 92
- EWAL 152
- exception response
 - DFC command, purging 523
 - ISC response mode errors 523
- protocol
 - response requirements (figure) 951
 - VTAM facilities 469
- exclusive mode 403
- EXEC parameter, SSM 141
- execution mode
 - recovery at ISC session failure 512
 - specification, ISC 475, 477
 - supported by ISC, list 462
- execution modes
 - CICS communication 609
 - ISC 609
- EXIT (/EXIT) command
 - ISC conversation mode errors 524
 - terminating a conversation abnormally 403
- exit interface block (XIB) 188
- exit interface block (XIB1) 191
- exit interface block ODBM data store entries 192
- exit interface blocks
 - IMS Connect 188

exit interface blocks (*continued*)

- XIBDS 189

exit routines

- affinity routing 738
- for ETO 110
- for MSC 707
- IMS Connect
 - overview 183
 - overview, function-specific exit routines 186
 - overview, user message exit routines 184
- IMSplex affinity routing 738
- IMSplex message routing 735
- MSC message routing 735
- routing 735
- TM and MSC Message Routing and Control user exit routine 735

exit routines, with OTMA 776

expedited message handler (EMH) 385

- queue option
 - overview 385

explanation of stopped status 167

express_context_interest service. 830

extended output component protection 918

Extended Recovery Facility 780

Extended Recovery Facility (XRF) 391

- Class 1 terminals 419
- Class 2 terminals 417
- Class 3 terminals 418
- class of service 418
- classes of terminals 416
- definition 414
- establishing communication
 - Finance Communication System 906
 - ISC 505
 - SLU P 906
 - system takeover
 - considerations 906
- master terminals 401
- MNPS 415
- recovery modes 408
- SLU P application program 901
- takeover considerations
 - Finance Communication System 906
 - SLU P 906
- terminal support 391, 418
- terminals in 416

Extended Terminal Option (ETO) 79, 374

- /SIGN command for ETO STSN devices 131
- 3275 devices 94
- 3600/Finance 130
- ABENDU0015 110
- advantages
 - availability 87
 - LTERMs 87
- algorithm
 - logon descriptor 103
 - LTERM allocation 116
- associated printing techniques 118
- asynchronous output 126
 - destinations, valid 126

Extended Terminal Option (ETO)

(continued)

- autologoff 123
- autologon 125
- autosignoff 122
- benefits of using 79
- coding descriptors
 - introduction 100
- commands that reset status and
 - release control blocks 116
- commands that retain status 116
- common logon descriptors 101
- concepts
 - summary 85
- conversations 131, 397
- creating descriptors 101
- customizing 84, 88
- dead-letter queue 126
- default CINIT/BIND user data
 - formats 112
- defining
 - autosignoff and autologoff
 - timer 124
 - parameters 121
- deleting control blocks
 - after logoff 128
 - after signoff 128
- delivering output to non-originating
 - terminal 127
- descriptors 84
 - added 110
 - definition 81
 - deleted 110
 - generic logon, definition 82
 - group logon, definition 82
 - logon 101
 - logon, definition 81
 - MFS 107
 - MFS device, definition 83
 - MSC (Multiple Systems
 - Coupling) 109
 - MSC (Multiple Systems Coupling),
 - definition 82
 - specific logon, definition 82
 - updated 110
 - user 104
 - user, definition 82
 - using 84
 - VTAM TERMINAL macro 101
- device characteristics table 94, 107
- device type, defining 92
- DFS2085 113
- DFS3641W 110
- DFS3645 113
- DFS3649A 115
- DFS3650I 115
- DFS3672 113
- DFSINSX0 106
- DFSSGNX0 106
- DFSUSER descriptor 106
- DLQT 126
- dynamic terminal
 - definition 79
- dynamic terminals 374
 - static terminals, using together 91
- dynamic user, definition 80
- exit routines 84

Extended Terminal Option (ETO)

(continued)

- Signon exit routine
 - (DFSSGNX0) 118
 - using 84
- exit routines, coding 110
- exit routines, list of 110
- guideline selection
 - logon descriptors 102
 - LOGOND parameter 102
- initialization
 - descriptor validation 110
 - DFSINTX0 110
- logging off 128
- logon
 - dynamic, limiting to specific
 - terminal types 111
- logon descriptors
 - creating during system
 - definition 102
 - NTO, 3600/Finance terminals 104
- LTERM
 - creating and reusing control
 - blocks 112
- LTERM with specific destination 115
- LU 2
 - devices 93
 - screen size and model
 - information 94
- LU 6.1 (ISC) terminals 129
- MFS 98
- MFS device characteristics table 94,
 - 107
 - screen size and model
 - information 94
- MFS device descriptors
 - MFSDCT utility (DFSUTB00) 108
- MFSDCT utility 94
- MODETBL on ETO logon
 - descriptor 428
- MSC (Multiple Systems Coupling)
 - descriptor 100
 - MSNAME macro 100
 - support 100
- multiple signons 114
- node user descriptor 106
- non-SNA 3270 devices
 - printers and displays 92
 - screen size and model
 - information 94
- NTO devices 94
- output
 - assigning 125
 - inadvertent output data
 - streams 127
- overview 79
- planning 87
 - LU2 95
- operations 99
- user IDs 99
- user queue names 99
- printers
 - defining 119
 - direct printing 118
 - overview 117
 - printer node names 118
 - sharing printers 120

Extended Terminal Option (ETO)

(continued)

- queue (dynamic user message),
 - definition 79
- RACF 374
- recommendations 106
- reduction of time of system
 - definition 374
- requirements 88
- response mode 131
- restrictions 88
- screen definitions examples (non-SNA
 - 3270) 97
- screen size control byte 94
- security 412
 - signon 98
- shared printing 120
- signing off, definition 127
- signing on, definition 113
- signon
 - providing signon data 113
- Signon exit routine (DFSSGNX0) 118
- SLU P 130
- SNA commands 111
- special processing modes 116
- starting ETO 110
- static terminal, definition 79
- static terminals
 - dynamic terminals, using
 - together 91
- storing descriptors 101
- structure
 - terminal, definition 80
 - user, definition 80
- structure, creation and deletion 83
- STSN terminals 130
 - overview 130
 - support for /SIGN command 131
- system definition 101
- terminal
 - definition 79
- terminal-LTERM relationship 115
- terminals
 - using static and dynamic terminals
 - together 91
- terminology 79
- undeliverable data, dead-letter
 - queue 126
- user descriptors
 - creating during system
 - definition 105
 - criteria for selection 105
 - DFSUSER 106
 - installation created 105
 - installation-created 104
 - node user descriptor 104, 106
- VTAM CINIT LUNAME 102
- VTAM considerations
 - logon CINIT session control
 - blocks 92
 - VTAM PSERVIC parameters 92
- external entry vector table (EEVT) 158
- external execution mode, ISC 475
- external subsystem attach facility
 - accessing multiple external
 - subsystems 164
 - application call processing 145

- external subsystem attach facility
 - (continued)
 - application threads 143
 - attach processing
 - overview 142
 - CHANGE command 164
 - commands
 - /START SUBSYS 165
 - /STOP SUBSYS 165
 - connection initiation 160
 - connections
 - dependent regions 161
 - CRC 146
 - creating ESMT 149
 - creating work areas for ESAP 159
 - deferring control region identify 160
 - dependent region connections 166
 - dependent regions
 - establishing connections 161
 - DFSEMODL macro 149
 - DFSEWAL macro 152
 - DISPLAY SUBSYS command 164
 - EEVPPEA 157
 - EEVPEWA 159
 - EEVT 158
 - EEVT mapping 158
 - ESMT
 - creating 149
 - description 140
 - establishing connections 142
 - EWAL 152
 - exit routine interface 140
 - explanation of stopped status 167
 - external subsystem command
 - support 146
 - functions supplied by external
 - subsystem 140
 - identify process 142
 - IMS services available to ESAP 147
 - INQ parameter 145
 - language interface module 163
 - loading ESAP 157
 - loading external subsystem
 - modules 158
 - notify message 142, 161
 - OASN 164
 - overview 139
 - PDSE 149
 - processing 157
 - recovery coordinator 139
 - recovery token 145, 164
 - resource coordination 145
 - RTT 140
 - signon process 143
 - specifying external subsystems to
 - IMS 141
 - SSR command 166
 - subsystem connections 142
 - subsystem termination 165
 - terminating connections 144
 - termination ECB 144, 166
 - termination requested by external
 - subsystem 166
 - thread 142
 - token 139
 - two phase commit process 145

- external subsystem attach facility
 - (continued)
 - unique language interface entry
 - points 163
 - user authorization processing 143
- External Subsystem Attachment Package (ESAP) 140
- external subsystem command
 - support 146
- external subsystem module table
 - PDSE 149
- external subsystem, definition 374
- external subsystems
 - attach facilities 137
 - callout request 847
 - connecting to 137
- extracting multiple system transaction
 - statistics 719

F

- facilities
 - available on IMS-CICS session 609, 632
 - component definition 909, 910
 - display screen protection 917, 918
 - Finance Communication System 909
 - ISC 473
 - Message Format Service 914
 - message recovery 915
 - SLU P 909
 - terminal-response mode 910
 - test mode for ISC 474
- failure
 - CICS-IMS session 646
 - ISC, recovering from in-bracket 510
 - system, during CICS-IMS session 646
 - while in-brackets 509
- Fast Path 374
 - EMH buffer size 413
 - Input Edit/Routing exit routine (DBFHAGU0) 386
 - input message 910
 - messages 374
 - MFS NEXTMSG or NEXTMSGP
 - control commands 910
 - MFS paged output messages 910
 - NFPACK option on Terminal
 - macro 924, 925
 - output messages 910
 - overview 374
 - parameters
 - FORCERESP 923
 - FPACK/NFPACK 923
 - OPTACK 923
 - recoverable-inquiry transactions 433, 952
 - recovery 407
 - restrictions for applications 386
 - routing code 385
 - RTR command 924, 925
 - TERMINAL macro
 - options 923
 - required parameters 923
 - terminal-response mode 911
 - terminals 413

- Fast Path EMH
 - shared queues environment
 - overview 385
- Fast Path, administering 780
- Finance communication 899
- Finance Communication System
 - application program
 - converting Finance to SLU P 902
 - functions available 900
 - MFS considerations 901
 - bracket and send/receive
 - protocols 907, 937
 - facilities 909
 - major parts 899
 - message resynchronization
 - controller application
 - program 929
 - sample configuration 900
 - system takeover considerations 906
 - terminals supported 899
 - VTAM
 - commands and indicators 903
 - facilities 902
 - workstations 900
 - XRF complex, establishing
 - connections 906
- first speaker
 - ISC contention winner 472
 - secondary logical units 472
- flow
 - commit-then-send 819, 827
 - of data
 - in an MSC (Multiple Systems Coupling) network 681
 - of resynchronization 835
 - send-then-commit 821
 - send-then-commit with Confirm 823
- FM (function management) header
 - error recovery procedure (ERP)
 - format 587
- FM (function management) headers
 - coding for use with CICS 638
- Finance Communication System 939, 946
- IMS, data descriptor 644
- input FM header length
 - ATTACH FM header 562
- input message 473
- input process names 584
- inserted by IMS 450
- ISC 559
 - ATTACH 562
 - ATTACH, MFS 570
 - data descriptor 572, 585
 - DPM messages, MFS 568
 - error recovery procedure 562, 645
 - initiating a process 562
 - Input QMODEL FM Headers 573
 - introduction 559
 - invoking ISC edit 561
 - MFS 567
 - QMODEL 572, 587
 - QMODEL, CICS 644
 - RAP 563
 - RAP, MFS 576
 - Reply QMODEL FM Headers 575
 - SCHEDULER 563, 641

FM (function management) headers
(*continued*)

- ISC (*continued*)
 - SCHEDULER, MFS 570
 - synchronous processing 562
 - SYMSMSG 564
 - system message 564, 644
- message descriptor byte format 940
 - input 940, 941
 - output 946, 947
- MFS 644
- output FM header 946
- output message 475
- output process names 584
- processing mode requested 475
- security 463
- SLU P 940, 947, 949
- type X'42' (SLU P) 942

FM headers 450

FMH 450

format, message

- control characters in message
 - prefix 451
- creating with SDF II 445
- transparency option 449

forms 453

FPAK/NFPAK option, Fast Path 923

front-end subsystem

- CICS
 - SEND LAST command 614
 - SEND/RECEIVE command 612
 - START/RETRIEVE command 617
- function of 461
- IMS 616, 619

Front-End Switch

- OTMA, unsupported 786

Front-End Switch exit routine (DFSFEJ0)

- special support 461

FULL security level 800

full-duplex message flow 760

function abort

- detecting loop, ISC 556
- indicated on LUSTATUS
 - command 547

function management (FM) headers

- coding for use with CICS 638
- error recovery procedure (ERP)
 - format 587

Finance Communication System 939, 946

IMS, data descriptor 644

input FM header length

- ATTACH FM headers 562

input message 473

input process names 584

inserted by IMS 450

ISC 559

- ATTACH 562
- ATTACH, MFS 570
- data descriptor 572, 585
- DPM messages, MFS 568
- error recovery procedure 562, 645
- format reference 577
- header format 577
- initiating a process 562
- Input QMODEL FM Headers 573
- introduction 559

function management (FM) headers
(*continued*)

- ISC (*continued*)
 - invoking ISC edit 561
 - MFS 567
 - QMODEL 572, 587
 - QMODEL, CICS 644
 - RAP 563
 - RAP, MFS 576
 - Reply QMODEL FM Headers 575
 - SCHEDULER 563, 641
 - SCHEDULER, MFS 570
 - supported by IMS, summary 559
 - synchronous processing 562
 - SYMSMSG 564
 - system message 564, 644
- message descriptor byte format 940
 - input 940, 941
 - output 946, 947
- MFS 644
- output message 475
- output process names 584
- processing mode requested 475
- security 463
- SLU P 940, 947, 949
- type X'42' (SLU P) 942

functions

- ISC and CICS 466, 609
- ISC versus MSC 460

G

generic resource groups 381

- benefits 381
- definition 382
- overview 381

generic resource member, definition 382

generic resource name

- definition 382

GRAPHIC= parameter

- NO option 450
- SEG statement 450

GRNAME parameter 768

H

half session

- definition 503, 507
- pairs 501, 507
- synchronization, ISC 525

half-duplex message flow 760

half-duplex protocol, IMS use 533

HANDLE CONDITION, CICS 629

headers

- input message 450
- ISC 450
- output message 450
- type X'42'
 - component identification 942
 - data bytes 948
 - message descriptor byte 941, 947

HIOP (high input/output pool)

- VTAM output buffers
 - RECASZ execution parameter 427

HIOP storage pool

- temporary shortage, reregistering
 - OTMA 782

hold queue, asynchronous

- retrieving asynchronous output for
 - alternate client IDs 339

horizontal partitioning in MSC (Multiple Systems Coupling) 677

HWS\$DLQ 346

HWSIMSCB macro 237

HWSJAVA0

- local option client
 - communication 172

HWSOMCTL DSECT 265

HWSOMUSR DSECT 283

HWSSMPL0 202

- and IRM 242
- message structures in simple
 - flow 262

HWSSMPL1

- and IRM 242
- message structures in simple
 - flow 262

HWSXIB macro 188

HWSXIB1 macro 191

HWSXIBDS macro 189

HWSXIBOD macro 192

I

identify process 142

immediate program switch 397

immediate session termination 517

IMS

- control region size and OTMA 794
- conversation
 - and commit-then-send mode 818
- device support with OTMA 756
- DFSnnnnn messages 346
- high-performance access 756
- IMS.ADFSMACT 855
- OTMAASY start-up parameter 809
- processing protected transactions 830
- recoverable transactions 304
- scheduler message block (SMB)
 - OTMA 758
- standard flow 819
- transactions
 - using a nonsynchronized
 - tpipe 832
 - using a synchronized tpipe 832
- z/OS Resource Recovery Services
 - exits supported 829

IMS connect

- BPE header format 250
- header format 249

IMS Connect

- alternate client ID 339
- ALTPCB output and shared
 - queues 197
- Application Transparent Transport
 - Layer Security feature
 - (AT-TLS) 212
- asynchronous callout requests from
 - TCP/IP clients 849
- asynchronous callout support
 - programming 848

IMS Connect (*continued*)

- asynchronous OTMA output
 - retrieving 328
- asynchronous output
 - auto message control 337
 - grouping 340
 - noauto message control 335
 - nooption message control 336
 - requesting from end-user application 333
 - RESUME TPIPE request 329
 - sharing 340
 - single message control 334
 - single with wait 335
- asynchronous output support 339
- availability, verifying 347
- calling out to external services from IMS 847
- callout
 - coding user-written clients for
 - synchronous callout 223
 - RESUME TPIPE call for
 - synchronous callout 226
 - RESUME TPIPE error scenarios 227
 - retrieving synchronous callout requests 226
 - returning synchronous callout responses 229
 - send-only protocol 229
 - synchronous, message format 225
- callout requests
 - retrieving from OTMA 328
- callout support 223
- callout support, synchronous
 - acknowledgment messages (ACK and NAK) 228
 - acknowledgment messages, IRM_F3_REROUT 228
 - acknowledgment messages, NAK 228
 - acknowledgment messages, SYNCNAK 229
 - returning error responses 230
- cancel timer 325
- changing RACF passwords 208
- client call flows 342
- client communications, restrictions for local 172
- client communications, TCP/IP 172
- client IDs
 - avoiding duplicates 314
 - canceled 314
- CM0
 - ignore DL/I PURG call 252
- CM0, purge function 299
- CM0, reroute function 301
- CM1 ACK/NAK timeout value 291
- commit mode 0, purge function 299
- commit mode 0, reroute function 301
- commit-then-send
 - ignore DL/I PURG call 252
- commit-then-send, purge function 299
- commit-then-send, reroute function 301
- conversational program support 292

IMS Connect (*continued*)

- conversational protocols 295
 - send-then-commit, sync
 - level=confirm, ACK response 297
 - send-then-commit, sync
 - level=confirm, NAK response 298
 - send-then-commit, sync
 - level=none, client terminated transaction 296
 - send-then-commit, sync
 - level=none, program terminated transaction 296
- data stores
 - status 189
- data stores, IMS DB
 - status 192
- dead letter queue 346
- defining
 - overview 180
- DFS messages
 - client response 344
- Distributed Relational Database Architecture (DRDA) 176
- exit interface block 188, 191
 - format, IMS DB connections 191
 - XIB1 format 191
- exit interface block (XIB) 188
 - format of entry 188
- exit interface block (XIB1) 191
- exit interface block data store entry
 - format of entry 190
- exit interface block ODBM data store (XIBOD) 192
- exit interface blocks 188
 - format of XIBOD 193
 - XIBDS 189
- exit routines
 - exit interface block (XIB) 188
 - exit interface block (XIB1) 191
 - input message format on return from exit 254
 - macro support 187
 - overview 183
 - overview, user message exit routines 184
- exit routines, function specific
 - overview 186
- exit routines, security 201
- HWS\$DLQ 346
- HWS\$JAVA0
 - user-defined messages 185
- HWSOMCTL DSECT 265
- HWS\$SMPL0
 - message structures in simple flow 262
- HWS\$SMPL1
 - message structures in simple flow 262
- HWS\$SOAP1 178
- HWSXIB macro 188
- HWSXIB1 macro 191
- HWSXIBDS macro 189
- HWSXIBOD macro 192
- ignore PURG call for CM0
 - multi-segment output 252

IMS Connect (*continued*)

- IMS TM Resource Adapter
 - callout support 849
 - duplicate client IDs 177
 - generated client IDs 177
 - message structures 253
 - support overview 177
 - two-phase commit support 357
- IMS Universal drivers 176
 - alias name 177
 - connection routing 177
 - one-phase commit 354
 - two-phase commit support 351
- IMSpIex support 195
- intervals for message timer 319
- introduction 171
- invoking
 - overview 180
- IRM
 - fixed portion 238
- IRM_F3_IPURG 252
- local option 171
- macros 187
- MAXSOC parameter
 - usage 311
- message formats
 - IRM fixed portion 238
 - output to message exit 257
- message structures 257, 260, 261, 262
 - examples 262
- message timeout intervals 318
- message timer
 - canceled 325
- mixed-case passwords 210
- ODBM
 - status 192
- OMHDRIPG 252
- OTMA conversational protocol 293
- OTMA conversational protocols
 - send-then-commit, sync
 - level=none 293, 294
- OTMA destination descriptor 772
- OTMA message header
 - HWSOMCTL DSECT 265
 - message control fields 265
 - notes 289
 - security data fields 281
 - state data fields 271
- OTMA message headers
 - format of user data section 283
- OTMA timeout for CM0
 - acknowledgments 326
- output message structure
 - to clients 252
- overview 171
- password management 208
- ping function 347
- protocols
 - RESUME TPIPE request 329
 - send only 305
 - send only for callout responses 307
- purge function 299
- purging multiple-message output 301
- purging output 299

- IMS Connect (*continued*)
 - purging undeliverable output
 - HWSSMPL0 and HWSSMPL1 300
 - when output is purged 300
 - RACF support 199
 - RACF, default user ID 200
 - RACF, enabling direct support 200
 - reroute function 301
 - rerouting commit-then-send
 - output 302
 - purging multiple-message output 304
 - specifying a destination 302
 - when output is routed 303
 - rerouting output 301
 - RESUME TPIPE 339
 - synchronous callout, coding for 226
 - RESUME TPIPE call
 - and timeout 338
 - RESUME TPIPE protocol
 - example flows 330
 - RESUME TPIPE request 328, 329
 - Secure Sockets Layer (SSL) 212
 - security 199
 - Application Transparent Transport Layer Security feature (AT-TLS) 212
 - asynchronous hold queue 202
 - error response to sample exit
 - RACROUTE calls 204
 - IMS Universal drivers 201
 - OTMA Resume TPIPE Security exit routine (DFSYRTUX) 202
 - overview 179
 - PassTicket, RACF 210
 - Resume tpipe 202
 - Secure Sockets Layer (SSL) 212
 - specifying OTMA ACEE aging value 221
 - trusted users 220
 - user message exit routines 185
 - security exit routine 201
 - security for 202
 - security, default RACF user ID 200
 - send-only protocol 305
 - synchronous callout, coding for 229
 - Send-only protocol
 - with acknowledgment 306
 - with serial delivery 307
 - send-only protocol and callout requests 307
 - send-then-commit timeout value 291
 - shared queues and ALTPCB
 - output 197
 - SOA composite business application support 292
 - SOAP Gateway 850
 - socket connections 308
 - non-persistent 309
 - persistent 308
 - transaction 309
 - sockets
 - maximum number 311
 - processing for transactions 310
 - reset threshold 313

- IMS Connect (*continued*)
 - sockets (*continued*)
 - setting percentages for warnings 313
 - UNIX System Services maximum number 312
 - warning threshold 313
 - SSL
 - default setup 218
 - initialization 214
 - sample JCL 218
 - z/OS key management 214
 - super member 340
 - Sysplex Distributor 179
 - TCP/IP
 - failures 316
 - KeepAlive intervals 315
 - TCP/IP communications 171
 - TCP/IP settings 365
 - timeout intervals, setting 318
 - timeout specifications 317
 - IMS DB clients 317
 - IMS TM clients 317
 - timeout value, CM1 ACK/NAK 291
 - timer
 - canceling 325
 - transaction expiration 327
 - setting time in IRM 328
 - transaction protocols 291
 - transactions
 - restrictions 291
 - two-phase commit 349, 360
 - IMS TM Resource Adapter support 357
 - IMS Universal drivers support 351
 - Unicode 363
 - UNIX System Services
 - socket limits 312
 - user ID, default for RACF 200
 - user message exit routines
 - overview 184
 - security 185
 - user message exits
 - output message structure 249
 - user-written client application
 - message structures 253
 - XIB (exit interface block)
 - format of entry 188
 - format, IMS DB connections 191
 - XIB1 format 191
 - XIB1
 - format 191
 - XIBDS
 - format of entry 190
 - XIBOD
 - format of entry 193
 - XML
 - COBOL conversion example 234
 - converting to COBOL 231
 - message structures 232
 - XML adapter 178
 - XML conversion support,
 - overview 178
 - XML converter 178
 - XML converters 232
 - z/OS Sysplex Distributor 179

- IMS DB resource adapter, installing 28
- IMS message switches, response requirements 954
- IMS Monitor
 - MSC considerations 718
 - Report Print Program 719
- IMS services available to ESAP 147
- IMS TM
 - IMS Connect
 - timeout specifications 317
- IMS TM Resource Adapter
 - calling out to external services from IMS 847
- IMS Connect
 - duplicate client IDs 177
 - generated client IDs 177
 - support overview 177
 - two-phase commit support 357
- IMS Connect callout support 849
- IMS Connect message structures 253
- one-phase commit 359
- two-phase commit
 - IMS Connect support 357
- IMS TM resources
 - resource type consistency
 - disabling enforcement 384
- IMS Universal Database resource adapter
 - one-phase commit 354
- type-2
 - installing in WebSphere
 - Application Server for z/OS 8
- type-4
 - EAR file, installing on WebSphere
 - Application Server 6
 - WebSphere Application Server for z/OS
 - application, installing 11
 - classpath, setting 9
 - data source, installing 10
- IMS Universal Database resource adapter, installing 5
- IMS Universal Database resource adapters
 - two-phase commit 350
- type-2 connectivity
 - WebSphere Application Server for z/OS configuration overview 6
- IMS Universal drivers
 - CICS
 - installing on CICS 13
 - configuring connections 3
- IMS Connect
 - timeout specifications 317
 - two-phase commit support 351
- IMS Connect security 201
- IMS Connect support 176
 - connection routing 177
- one-phase commit 354
- security
 - IMS Connect 201
- two-phase commit
 - IMS Connect support 351
- WebSphere Application Server
 - configuring 4
- IMS-CICS communication
 - alternate facility 634
 - application coding for 631

IMS-CICS communication (*continued*)

- asynchronous processing flow 617
- ATTACH parameters 638
- CICS transactions, definition 628
- coding
 - function management
 - headers 638
 - system definition options 620
- facility
 - alternate 634
- IMS commands 635
- IMS-CICS ISC 628
- initiating sessions 629
- integrity of session 646
- LU 6.1 links
 - compatible nodes 622
 - description 621
 - Macro-Level Resource Definition 621
 - multiple links 626
 - Resource Definition Online 621
- MFS support 644
- preparing CICS tables 621
- principal facility 634
- processing flows
 - RECEIVE 616
 - RETRIEVE 619
 - SEND INVITE 612
 - SEND LAST 614
 - SEND/RECEIVE 612
 - START/RETRIEVE 617
- recovery and restart 645
- SCHEDULER parameters 641
- session
 - binding 647
 - initiation 629
 - processing outstanding traffic 647
 - reestablishing 646
 - resynchronizing 647
 - sync points 636
 - termination 630
- sync points 636
- terminating a session 630
- transactions
 - attributes supported 609, 634
 - types supported 609, 634

IMS-to-IMS communication, LU 6.1 protocols 494, 496

IMS-to-IMS sessions, ISC protocol restrictions 495

IMS.FORMAT, output from MFS 438

IMS.PROCLIB 417, 701

IMSCTRL macro, MSC system definition 700

IMSIDs

- when an IMSplex and MSC network coexist 728

IMSMSV procedure 701

IMSplex

- affinity
 - in an MSC-IMSplex configuration 723
- APPC and OTMA messages
 - processing MSC remote transactions 728
- definition 377
- environment requirements 196

IMSplex (*continued*)

- IBM Management Console for IMS and DB2 for z/OS 195
- IMS Connect configuration file 196
- IMS Connect support 196
- IMSIDs
 - when an IMSplex and MSC network coexist 728
- installation 196
- link definitions in an IMSplex
 - deleting 727
- message routing
 - in an MSC-IMSplex configuration 721
- MSC (Multiple Systems Coupling)
 - coexistence 721
 - migration 723
 - sharing MSNAME definitions and SYSIDs 724
- MSC MSNAME definitions
 - duplication 726
 - sharing MSNAME definitions and SYSIDs 724
- MSNAME definitions
 - deleting 727
- OM access 195
- pseudoabend U0830
 - avoiding 734
- SCI (Structured Call Interface) 195
- SYSIDs
 - cloning MSC SYSIDs in an IMSplex 728
 - managing MSC SYSIDs in an IMSplex 728
 - sharing MSNAME definitions and SYSIDs 724
- terminal management 383

IMSplex affinity routing 738

IMSplex support 195

IMSplexes

- MSC
 - APPC and OTMA remote transactions 728
 - back-end processing of remote APPC or OTMA transactions 729

Information Management System (IMS) 373

initiation

- session 431

INOP state 404

input bracketing

- replies 942

input component

- identification for header type X'42' 942
- relationship to output component 477
- versus output component 909

input editing option

- ISC 473
- SLU P 919

input errors, IMS detection of ISC 553

input message

- backing out ISC 525
- bracketing 940, 942
- chained messages 542, 943

input message (*continued*)

- Fast Path 923
- Finance Communication System chaining indicator 939
- header
 - DPM 940
 - Finance format 939
 - MID name location 939, 941
 - optional MFS (Message Format Service) fields 940
 - SCS 941
 - SLU P format 940
 - use 939
- ISC bracketing 533
- message descriptor byte
 - header type X'42' 941
 - MID name indicator bit 940, 941
- MFS (Message Format Service) input formatting, activating 942
- MFS input formatting, activating 567
- multiple transmissions for one message 943
- origin of for MFS (Message Format Service) purposes 943
- requirements 939
- response requirements 951
- SLU P
 - chaining indicator 939
 - restriction 942
- sync point
 - availability 525
 - ISC 525

input message structure 364

- from clients 251

input message, response requirements 529

input system (MSC) 686

Input/Output Edit exit routine (DFSYIOE0) 776

inquiry flag processing 145

integrity

- message
 - IMS-CICS session 646
 - in ISC 462
 - NOCHECK PROTECT 617

interchange unit code, ATTACH FM header 579

intermediate IMS

- MSC
 - security checking 711

intermediate system (MSC) 686

internal execution mode, ISC 476

Intersystem Communication (ISC) 459

- application program
 - accessing 482
 - example 669
 - not using MFS 483
 - routing example 486, 488, 489, 490, 491, 493
- ATTACH FM header
 - bit contents 577
 - format 577
- ATTDBA 580
- ATTDSP 579
- binding an ISC session 506
- CICS
 - abends, transaction 649

Intersystem Communication (ISC)
(continued)

CICS (continued)

- asynchronous transactions,
 - defining 629
- ATTACC 640
- ATTDBA 640
- ATTDSP 640
- ATTDQN 640
- ATTDSP 640
- ATTIU 640
- ATTPRN 639
- ATTRDPN 639
- ATTRPRN 639
- data stream profile (ATTDSP) 640
- deblocking algorithm (ATTDBA) 640
- in-doubt processing 628
- interchange unit code (ATTIU) 640
- network name 622
- number of sessions 623
- outbound destination process name field (ATTDPN) 638
- outbound destination process name field (SCDDPN) 641
- primary resource name field (ATTPRN) 639
- primary resource name field (SCDRPN) 642
- return destination process name (ATTRDPN) 639
- return destination process name (SCDRDPN) 642
- return primary resource name (ATTRPRN) 639
- return primary resource name (SCDRPRN) 642
- sample program DFSISC00 669
- sample program, ACB generation 674
- sample program, installing 669
- sample program, JCL to compile and bind 672
- sample program, MFS formats 673
- sample program, PSB generation 674
- sample program, system definition statements 672
- SCDDP 642
- SCDDPN 641
- SCDDQN 642
- SCDPN 642
- SCDRDPN 642
- SCDRPRN 642
- session initiation 630
- session names 623
- session parameters 623
- terminal device-dependent data 467
- transaction abends 649
- CICS-IMS 607
 - application-related concepts 631
- control block storage, parallel sessions 475

Intersystem Communication (ISC)
(continued)

- conversation mode termination extension 524
- data descriptor FM header
 - use on input 572
- data flow control
 - example 653, 655, 658, 660, 663
- DFC protocols 533
- DFC sync point 525
 - exception 525
 - irrecoverable messages 529
 - MFS messages 527
 - on output 526
 - recoverable messages 527
- editing facilities
 - invoking FM headers 561
 - overview 473
- editing options 448
- error handling
 - MTO notified 550
 - sender ERP 554
- error recovery procedure
 - examples 665
 - examples, receiver-detected errors 666
 - examples, sender-detected errors 665
- error recovery, purge after exception response 542
- examples of ATTACH and SCHEDULER parameters with MFS 601
- examples of ATTACH parameters with SYSMMSG 599
- execution mode
 - external specification 475
- facilities 473
- FM headers
 - ATTDBA 580
 - ATTDSP 579
 - CICS session 638
 - format reference 577
 - introduction 559
 - MFS 567
 - QGET 587
 - QGETN 589
 - QPURGE 590
 - QSTATUS 590
 - QXFR 591
 - RAP 594
 - routing messages 484
 - SCHEDULER 594
 - supported by IMS 559
 - SYSERROR 595
 - SYSSTAT 595
- half session, primary versus secondary 506
- IMS commands, issuing 474
- IMS support, CICS 607
- IMS-to-IMS session
 - design considerations 494
 - MSC and ISC coexistence 494
- interchange unit code, ATTACH FM header 579
- introduction 459

Intersystem Communication (ISC)
(continued)

- ISC edit ATTACH parameter examples 597
- ISC node
 - defining 620
- ISC protocol
 - communications design 481
- ISC sessions
 - unsupported commands 620
- message integrity, CICS 646
- message resynchronization sessions 507
- message routing
 - FM headers 484
- message switch
 - ATTACH FM header 577
- MSC coexistence 494
- multichain input message switches restriction 579
- network operation 503, 504
- network, bringing up 503
- network, starting 503
- node definition
 - macros used 496
- output editing option 473
- output protocols 481
- parallel sessions issuing IMS commands 474
- protocols
 - defined 503
 - restrictions on IMS-to-IMS sessions 495
- QGET 587
- QGETN 589
- QPURGE 590
- QSTATUS 590
- QXFR 591
- RAP 594
- resynchronization procedure 509
- routing examples 485
- routing parameters 485
- SC (session control) protocols 504
- SC protocols
 - session initiation 504
- SCHEDULER 594
- selective receiver ERP
 - X'0864xxxx': function abort 551
 - X'0865xxxx': function abort 551
 - X'0866xxxx': function abort 552
- sense codes
 - received 557
- statically defining 496
- sync-point indicators 529
- SYSERROR 595
- SYSSTAT 595
- system definition
 - macro parameters 498
- test mode 474
- transaction types supported, CICS session 609, 634
- versus CICS 634
- VTAM facilities 469
- XRF
 - USERVAR, session initiation 505
- XRF complex 505

- invalid destinations in MSC (Multiple Systems Coupling) 694
- IOPCB 758
- IRM (IMS request message)
 - fixed portion format 238
- IRM (IMS Request Message) 238, 242, 363
- IRM_TIMER
 - usage 318
 - values 318
- irrecoverable messages, ISC sync points 529
- irrecoverable output 819
- irrecoverable transactions 495
- irrecoverable-inquiry transactions
 - LU 6.2 application program 432
 - response requirements 432, 953
- ISC
 - CICS
 - asynchronous processing, overview 607
 - functions supported 607
 - synchronous processing, overview 607
 - functions supported
 - CICS sessions 607
- ISC (Intersystem Communication) 459
 - application program 669
 - accessing 482
 - example 669
 - not using MFS 483
 - routing example 486, 488, 489, 490, 491, 493
 - using MFS (Message Format Service) 482
 - ATTACH FM header
 - bit contents 577
 - format 577
 - ATTDBA 580
 - ATTDSP 579
 - basic functions 459
 - binding an ISC session 506
 - CICS
 - abends, transaction 649
 - asynchronous transactions, defining 629
 - ATTACC 640
 - ATTDBA 640
 - ATTDP 640
 - ATTDPN 638
 - ATTDQN 640
 - ATTDSP 640
 - ATTIU 640
 - ATTPRN 639
 - ATTRDPN 639
 - ATTRPRN 639
 - data stream profile (ATTDSP) 640
 - deblocking algorithm (ATTDBA) 640
 - in-doubt processing 628
 - interchange unit code (ATTIU) 640
 - network name 622
 - number of sessions 623
 - outbound destination process name field (ATTDPN) 638

- ISC (Intersystem Communication)
 - (continued)
 - CICS (continued)
 - outbound destination process name field (SCDDPN) 641
 - primary resource name field (ATTPRN) 639
 - primary resource name field (SCDPN) 642
 - return destination process name (ATTRDPN) 639
 - return destination process name (SCDRDPN) 642
 - return primary resource name (ATTRPRN) 639
 - return primary resource name (SCDRPN) 642
 - sample program DFSISC00 669
 - sample program, ACB generation 674
 - sample program, installing 669
 - sample program, JCL to compile and bind 672
 - sample program, MFS formats 673
 - sample program, PSB generation 674
 - sample program, system definition statements 672
 - SCDDP 642
 - SCDDPN 641
 - SCDDQN 642
 - SCDPN 642
 - SCDRDPN 642
 - SCDRPN 642
 - session initiation 630
 - session names 623
 - session parameters 623
 - terminal device-dependent data 467
 - transaction abends 649
 - CICS and IMS functions 466
 - CICS-IMS 607
 - application-related concepts 631
 - Class 2 terminal 463
 - configurations 463
 - control block storage, parallel sessions 475
 - conversation mode termination extension 524
 - data descriptor FM header
 - use on input 572
 - data flow control 459
 - example 653, 655, 658, 660, 663
 - DFC protocols 533
 - availability 525
 - backing out input messages 525
 - BID command 533
 - bracket and half-duplex 533
 - bracketing input messages 533
 - bracketing output messages 536
 - CANCEL command 541
 - chaining 542
 - CHASE command 542
 - conversation mode errors 524
 - exception response 523
 - half-session synchronization 525

- ISC (Intersystem Communication)
 - (continued)
 - DFC protocols (continued)
 - LUSTATUS 546
 - paged message errors 549
 - recoverability supplemented by LWA 525
 - response mode errors 523
 - response requirements 525
 - RSHUT command 549
 - RTR command 549
 - selective receiver ERP 549
 - sync point, input messages 525
 - DFC sync point 525
 - exception 525
 - irrecoverable messages 529
 - MFS messages 527
 - on output 526
 - recoverable messages 527
 - distributed processing 461
 - distributed services, support for 462
 - editing facilities
 - invoking FM headers 561
 - overview 473
 - editing options 448
 - error handling
 - MTO notified 550
 - sender ERP 554
 - error recovery procedure
 - examples 665
 - examples, receiver-detected errors 666
 - examples, sender-detected errors 665
 - error recovery, purge after exception response 542
 - examples of ATTACH and SCHEDULER parameters with MFS 601
 - examples of ATTACH parameters with SYSMSG 599
 - execution mode 462, 475
 - external specification 475
 - internal definition 476
 - internal, with CICS 609
 - processing mode 477
 - synchronous versus asynchronous 476
 - facilities 473
 - FM headers 484, 559, 638
 - ATTACH 562
 - ATTACH, MFS 570
 - ATTDBA 580
 - ATTDSP 579
 - CICS session 638
 - data descriptor 572
 - DPM messages, MFS 568
 - error recovery procedure 562
 - format reference 577
 - initiating a process 562
 - introduction 559
 - invoking ISC edit 561
 - MFS 567
 - QGET 587
 - QGETN 589
 - QMODEL 572
 - QPURGE 590

ISC (Intersystem Communication)
(continued)

- FM headers (continued)
 - QSTATUS 590
 - QXFR 591
 - RAP 594
 - RAP (reset attached process) 563
 - RAP (reset attached process), MFS 576
 - reply QMODEL 575
 - routing messages 484
 - SCHEDULER 563, 594
 - SCHEDULER, MFS 570
 - supported by IMS 559
 - synchronous processing 562
 - SYSError 595
 - SYSMSG 564
 - SYSSTAT 595
- half session, primary versus secondary 506
- IMS commands, issuing 474
- IMS facilities 461
- IMS support for 461
- IMS support, CICS 607
- IMS-to-CICS configuration 463
- IMS-to-IMS configuration 464
- IMS-to-IMS session 494
 - buffer sizes 494
 - defining parallel session 494
 - defining single session 494
 - dequeuing messages 510
 - design considerations 494, 496
 - limit of parallel sessions 494
 - MSC and ISC coexistence 494
 - protocol restrictions 495
 - remote control 494
 - restriction on conversation mode 495
 - routing to back-end IMS 495
- interchange unit code, ATTACH FM header 579
- introduction 459
- ISC edit ATTACH parameter examples 597
- ISC node
 - defining 620
- ISC protocol
 - communications design 481
- ISC sessions
 - unsupported commands 620
- message integrity 462
- message integrity, CICS 646
- message resynchronization 507
 - sessions 507
 - summary 531
- message routing 484
 - FM headers 484
 - parameters 484, 485
 - through MSC links 493
- message switch
 - ATTACH FM header 577
- MSC coexistence 494
- multichain input message switches
 - restriction 579
- network operation 503, 504
- network, bringing up 503
- network, starting 503

ISC (Intersystem Communication)
(continued)

- node definition 496
 - COMM macro 499
 - example 496
 - macros used 496
 - NAME macro 499
 - SUBPOOL macro 499
 - summary 501
 - TERMINAL macro 500
- output editing option 473
- output protocols 481
- parallel sessions issuing IMS commands 474
- passing CICS data to IMS 467
- protocols 459
 - defined 503
 - restrictions on IMS-to-IMS sessions 495
 - summary 459
- QGET 587
- QGETN 589
- QPURGE 590
- QSTATUS 590
- QXFR 591
- RAP 594
- resynchronization procedure 509
 - cold start, to recover sessions 510
 - commands, recovery at failure 512
 - controlling backout of work unit 510
 - execution mode, recovery at failure 512
 - maintaining sequence numbers 509
 - output available at restart 510
 - recovering from in-bracket failure 510
 - restart process 512
 - session failure, IMS failure 512
 - session failure, IMS still running 512
- routing 459
- routing examples 485
- routing parameters 485
- SC (session control) protocols 504
- SC protocols
 - binding sessions 506
 - cold start recovery 510
 - commands, recovery at failure 512
 - controlling pending work units 510
 - designing restart procedures 509
 - execution mode, recovery at failure 512
 - maintaining sequence numbers 509
 - output available at restart 510
 - recovering from in-bracket failure 510
 - resolving bind race 507
 - restart process 512
 - running the session 516
 - session failure, IMS failure 512

ISC (Intersystem Communication)
(continued)

- SC protocols (continued)
 - session failure, IMS still running 512
 - session initiation 504
 - session states 515
 - session termination 517, 554
 - XRF, session initiation 505
- SCHEDULER 594
- security 463
- selective receiver ERP
 - X'0864xxxx': function abort 551
 - X'0865xxxx': function abort 551
 - X'0866xxxx': function abort 552
- sense codes 557
 - received 557
 - sent 557
- session control 459
- statically defining 496
- sync-point indicators 529
- SYSError 595
- SYSSTAT 595
- system definition
 - macro parameters 498
- test mode 474
- transaction types supported, CICS
 - session 609, 634
 - versus CICS 634
- VTAM API 471
 - CICS 609
 - IMS 471
- VTAM facilities 469
- XRF
 - MNPS, session initiation 505
 - USERVAR, session initiation 505
 - XRF complex 463, 505
- ISC data flow control examples 653
- ISC edit
 - default editor 473
 - non-MFS programs 483
- ISC-CICS installation options 620

J

Java

- accessing DB2 for z/OS data from IMS 137
- classic Java APIs for IMS
 - configuring connections 15
- connecting to IMS from external Java environments 1
- DB2 for z/OS
 - accessing from Java dependent regions 137
- dependent regions
 - accessing DB2 for z/OS 137
- Java application programs
 - calling out to external services from IMS 847
 - callout support overview 849
- JAVA_HOME ENVAR keyword 20
- JAVAENV data set
 - sample 21

- JBPs
 - accessing DB2 for z/OS data from IMS 137
- JCL
 - to protect transactions 829
- JDBC
 - type-2
 - IMS Universal drivers and CICS 12
 - IMS Universal JDBC driver and CICS 12
- JMPs
 - accessing DB2 for z/OS data from IMS 137
- JOBLIB
 - TM and MSC Message Routing and Control user exit routine 735

K

- keyboard shortcuts xviii

L

- language interface 163
- language interface entry points 163
- language interface module 163
- legal notices
 - notices 963
 - trademarks 963
- LIBPATH ENVAR keyword 21
- libraries, online change 445
- link
 - MSC (Multiple Systems Coupling) 677
 - priorities, setting 705
- link paths
 - MSC (Multiple Systems Coupling) 679
- Link-Receive Routing exit routine in MSC 737
- LINKLIST
 - TM and MSC Message Routing and Control user exit routine 735
- links
 - bandwidth of MSC links 696
 - deleting MSC link definitions in an IMSplex 727
 - MSC
 - optimum link types 713
 - MSC (Multiple Systems Coupling)
 - definitions in an IMSplex 724
 - logical 679
 - physical 677
 - MSC bandwidth mode, capacity 714
 - MSC logical link benchmarks 713
 - MSC logical link capacity 714
 - MSC logical link high value statistics 717
 - MSC logical link statistics 712
 - MSC logical link statistics and buffer sizes 714
 - MSC logical link statistics, resetting 713
 - MSC, controlling bandwidth 696

- list structure
 - definition 379
 - overflow, definition 379
 - primary, definition 379
- loading ESAP 157
- loading external subsystem modules 158
- local option client communication 172
- local system
 - MSC (Multiple Systems Coupling)
 - defined 681
- lock mode 403
- LOCK state 433
- Log Analysis report 720
 - ID column for MSC entries 720
 - MSC transactions 720
- Log Merge utility
 - log output, control of 720
 - MSC logs 719
- Log Transaction Analysis utility, MSC
 - statistics 719
- log write-ahead (LWA) 462
- logging
 - MSC 718
- logging off 377
- logging off, definition 377
- logging on, definition 377
- logical destinations
 - MSC (Multiple Systems Coupling) 684
- logical link paths
 - MSC (Multiple Systems Coupling) 679
- logical links
 - defining 704
 - MSC (Multiple Systems Coupling) 679
- logical network design 397
- logical page advance function (NEXTLP), MFS 916
- logical page requests function (PAGE=nn), MFS 916
- logical terminal (LTERM)
 - chains 398
 - component definition 579
 - convention for naming 478
 - eliminating the need for 910
 - definition 391
 - ETO 374
 - master terminal 401
 - method for naming 909
 - MSC and 737
 - network design 397
 - queues 399
 - relationship to physical terminals 397
 - remote 691
 - separating input and output devices 399
 - subpools, users and components 478
- logical unit
 - definition 370
 - programmable, definition 371
- logical unit (LU)
 - multiple, managing 69
 - qualifying LU names 68
 - reassigning 70
- logical unit of work 645

- logical unit status command 546
- logon descriptor
 - ETO
 - definition 81
 - generic 82
 - group, definition 82
 - specific, definition 82
- logon mode
 - default logon mode table 471
 - session initiation 905
- LOGON MODE table 428
- LTERM (logical terminal)
 - chains 398
 - component definition 579
 - convention for naming 478
 - eliminating the need for 910
 - definition 391
 - ETO 374
 - master terminal 401
 - method for naming 909
 - network design 397
 - queues 399
 - relationship to physical terminals 397
 - remote 691, 737
 - separating input and output devices 399
 - subpools, users and components 478
- LU
 - APPC outbound LU specification 57
 - outbound
 - specifying 57
- LU (logical unit)
 - multiple, managing 69
 - qualifying LU names 68
 - reassigning 70
- LU 6.2
 - descriptors 392
- LUSTATUS command 546
 - CICS 611, 634
 - commit 548
 - conversation mode
 - requesting normal termination 524
 - sending errors 524
 - function abort 547
 - NO-OP 548
 - paging errors 549
 - protocol 546
 - queue empty 547
 - response mode errors 523
 - SLU P system 958
 - terminating test mode 474
- LWA (log write-ahead) 525

M

- macro keywords 426
 - APPLID on COMM macro 426
 - COMM on BUFPOOLS macro 427
 - MODETBL on ETO logon
 - descriptor 428
 - MODETBL on TERMINAL
 - macro 428
 - PASSWD on COMM macro 426
 - RECANY on COMM macro 427

- macros
 - IMS Connect 187
 - XCF, and OTMA 785
- maintenance 443
- making CICS ready 629
- making IMS ready 503, 927
- master terminal 428
 - device choice 401
 - logon requirements 428
 - MSC routing 737
 - reserving an LTERM 391
 - XRF complex 401
- master terminal operator (MTO) 428
- master-slave MSC (Multiple Systems Coupling) relationship 677
- MAXFILEPROC parameter, UNIX System Services 312
- memory-to-memory (MTM)
 - MSC (Multiple Systems Coupling) physical link type 677
 - MSC physical link defining 703
- message
 - architected form 817
 - bypassing MFS editing 449
 - control characters 451
 - destination determination 776
 - editing of output segments 450
 - editing performed by IMS 447
 - examples of how to select 781
 - extending 756
 - flow
 - commit-then-send 819, 827
 - deallocate 817
 - definition 758
 - in full-duplex environment 764
 - resynchronization 837
 - send-then-commit 821
 - send-then-commit with Confirm 823
 - use of queues in tpipe 763
 - flow in an OTMA environment 758
 - formatting and editing 445
 - full-duplex flow 760
 - half-duplex flow 760
 - IMS handling 378
 - in-flight
 - device LUs 430
 - program LUs 430
 - output segment editing 450
 - prefix
 - contents 827
 - OTMA application-data section 883
 - OTMA security data 880
 - OTMA security-data section 880
 - OTMA user data section 883
 - state-data section 867
 - syntax 855
 - requeuer 781
 - resynchronization
 - sample 839
 - sample
 - acknowledge receipt of CBresynch 839
 - acknowledge receipt of SRVresynch 839
- message (*continued*)
 - sample (*continued*)
 - client-bid request with resynchronization 839
 - REPresynch command 840
 - REQresynch command 840
 - SRVresynch command 839
 - successful resynchronization 840
 - sample OTMA message 884
 - scheduling
 - definition 394
 - Fast Path 385
 - selective recovery 781
 - sensitivity to nongraphic message data 450
 - sequence numbers 816
 - sequential order 823
 - switch 807
 - in shared queues environment 811
 - translation 363
 - Z2 field 451
- message advance function (NEXTMSG) 916
- message advance protect function (NEXTMSGP) 916
- message buffering with a Fast Path-capable system 386
- message contention 507, 920
- Message Control/Error exit routine (DFSCMUX0) 708
- message descriptor byte
 - input FM header 940, 941
 - output message 946, 947
- Message Format Service
 - OTMA, unsupported 785
- Message Format Service (MFS) 437
 - administration 442
 - advantages 443
 - application programs, accessing with ISC 482
 - Bid options 917
 - bypassing MFS 449
 - bypassing with ISC application programs 482
 - components 438
 - components, overview 445
 - control functions
 - Finance Communication System 915
 - NEXTLP 916
 - NEXTTP 916
 - PAGE=nn 916
 - control requests, response requirements 955
 - data bytes output message 948
 - default format 595
 - defining 914
 - delete characters 450
 - description 375
 - DPM 945
 - DPM option 450
 - edit of ISC messages 473
 - editing 449
 - escape characters, not supported in ISC 567
 - facilities available 914
- Message Format Service (MFS) (*continued*)
 - FM header
 - activating output formatting 567
 - FM headers
 - activating input formatting 567
 - editing 567
 - facilities available 567
 - types 567
 - formatting 928
 - activating input 928
 - input formatting 567
 - input messages 437
 - input segments 450
 - introduction 437
 - invalid paging request 918, 920
 - Language utility 438, 445
 - libraries, online change 445
 - message editor 445
 - message recovery 915
 - MID/MOD chaining 915
 - MSC (Multiple Systems Coupling) 700
 - online error detection 569
 - online performance 443
 - output formatting 915, 949
 - MOD name 949
 - typical application program procedure 949
 - output formatting, ISC 567
 - output messages
 - editing segments 450
 - how MFS defines 437
 - overview 437
 - page delete function, not supported in ISC 568
 - paging, CICS 609, 633
 - pool manager 445
 - SDF II 445
 - Service utility 445
 - SLU P
 - benefits 915
 - station-by-station, availability 915
 - sync point, ISC messages 527
 - terminal keyboard lock and unlock 450
 - message handling 378
 - message headers 450
 - message input descriptor (MID) 445
 - message integrity 462
 - message output descriptor (MOD) 445, 949
 - message queue, definition 385
 - message recovery
 - Finance Communication System 921
 - message resynchronization 531, 922
 - MFS 915
 - restriction 922
 - SLU P system 921
 - message resynchronization 507
 - associated system definition options 509
 - CICS-IMS session 647
 - description 509, 647
 - design considerations 509, 929
 - direction and bracket indicators 937

message resynchronization (*continued*)

- effects on normal data
 - transmission 922, 929
- Fast Path 925
- how and when initiated 507, 922
- last inbound/outbound 922
- options for message sequence
 - numbers 929
- performing 647
- polarity of half-session pairs 509
- purpose 507
- requirements 507, 922
- send/receive and bracket
 - protocol 937
- sequence numbers 509, 930
- STSN
 - flow, primary-to-secondary 518
 - flow, secondary-to-primary 519
 - format 520, 930
 - use of 513

message routing

- examples 485
- in an MSC (Multiple Systems Coupling) network 682
- in ISC
 - across an MSC link 493
 - use of routing parameters 485

message structures

- for IMS TM Resource Adapter 252
- for user-written IMS Connect client
 - application programs 252

message switch

- ATTACH FM header 577
- examples 486
- ISC 484, 485

message switch, ISC 473

message switches, response
 requirements 529

message switching, DFSAPPC 70

messages

- affinity routing 738
- flow
 - using transaction pipes 763
 - using XCF 763
- IMSplex 735
- IMSplex affinity routing 738
- MSC 735
- OTMA
 - sending 815
- protected conversation, OTMA
 - restrictions 785
- routing
 - in an MSC-IMSplex
 - configuration 721
 - OTMA destination descriptor 772
 - OTMA destination descriptor,
 - defining 774
 - OTMA destination descriptor,
 - masking destination names 775
 - TM and MSC Message Routing
 - and Control user exit
 - routine 735
 - terminal/input routing 735
 - TM and MSC Message Routing and
 - Control user exit routine 735

MFS (Message Format Service) 437

- administration 442

MFS (Message Format Service) (*continued*)

- advantages 443
- application programs, accessing with
 - ISC 482
- Bid options 917
- bypassing MFS 449
- bypassing with ISC application
 - programs 482
- components 438
- components, overview 445
- control functions 915
 - Finance Communication
 - System 915
 - NEXTLP 916
 - NEXTMSG 916
 - NEXTMSGP 916
 - NEXTPP 916
 - PAGE=nn 916
 - SLU P 916
- control requests, response
 - requirements 955
- data bytes output message 948
- default format 595
- defining 914
- delete characters 450
- description 375
- DPM 945
- DPM option 450
- edit of ISC messages 473
- editing 449
- escape characters, not supported in
 - ISC 567
- facilities available 914
- FM headers 567
 - activating input formatting 567
 - activating output formatting 567
- ATTACH 570
- data descriptor 572
- DPM messages 568
- editing 567
- facilities available 567
- MOD name 567
- QGET 574
- QGETN 573
- QMODEL 572, 644
- QPURGE 575
- QSTATUS 575
- QXFR 575
- RAP (reset attached process),
 - MFS 576
- reply QMODEL 575
- SCHEDULER 570
- specifying MID name 567
- types 567

formatting 928

- activating input 928
- activating output 943

input formatting 567

- input messages 437
- input segments 450

introduction 437

invalid paging request 918, 920

Language utility 438, 445

libraries, online change 445

message editor 445

message recovery 915

MFS (Message Format Service) (*continued*)

- MID/MOD chaining 915
- MSC (Multiple Systems Coupling) 700
- online error detection 569
- online performance 443
- output formatting 915, 949
 - MOD name 949
 - typical application program
 - procedure 949
- output formatting, ISC 567
- output messages
 - editing segments 450
 - how MFS defines 437
- overview 437
- page delete function, not supported in
 - ISC 568
- paging, CICS 609, 633
- pool manager 445
- SDF II 445
- Service utility 445
- SLU P 915
 - benefits 915
 - station-by-station, availability 915
- sync point, ISC messages 527
- terminal keyboard lock and
 - unlock 450

MFS bypass option effect 438

MFS device characteristics table
 entries 389

- how used 94
- use with non-SNA 3270 devices 96

MFS device descriptor, definition 83

MFS Distributed Presentation
 Management (DPM) 473

MFS Service utility 441

MFSTEST procedure 445

MID (message input descriptor) 445

MID/MOD chaining 915

migration
 IMSplex from MSC 723

- MSC to IMSplex 723

mirror transaction, CICS 617

mixed-case passwords
 IMS Connect 210

MNPS (multinode persistent sessions)
 XRF 415

MOD (message output descriptor)
 name specification 567, 949

- purpose 445

modes, terminal
 conversation 402

- ETO and exclusive mode 403
- exclusive 403
- lock 403
- LU 6.2 402
- response 402
- SNA QUIESCE 405
- test mode 403

MODETBL
 specifying a default LOGON MODE
 identifier 428

MODETBL= keyword
 overriding the defaults 471

- use during ISC logon 471

- modified application program
 - LU 6.2 descriptor 60
 - remote execution, MSC 51
- modules
 - DFSERA0 748
 - DFSARG0 748
 - DFSARM0 748
 - DFSCDLI0 748
- monitoring and tuning, MSC 718
- monitoring performance 786
- MSASSIGN (/MSASSIGN) command
 - assigning physical link to logical link 679
- MSC
 - message routing 721
- MSC (Multiple Systems Coupling) 62, 677, 779
 - /MSVERIFY command 710
 - administration
 - APPC 695
 - affinity
 - in an MSC-IMSplex configuration 723
 - APPC and OTMA messages
 - remote processing 728
 - APPC/IMS
 - application program failure and transaction recovery 66
 - intermediate IMS failure and transaction recovery 65
 - local IMS failure and transaction recovery 65
 - local transaction discardability 63
 - LU 6.2 recoverability flows 66
 - LU 6.2 sessions and transaction recovery 63
 - MSC links and transaction recovery 64
 - remote IMS failure and transaction recovery 65
 - transaction point of failure 63
 - bandwidth mode, capacity 714
 - bandwidth, link 696
 - benchmarking logical link
 - performance 713
 - buffer sizes 702
 - buffer sizes, links 696
 - buffers
 - format in bandwidth mode 715
 - format in non-bandwidth mode 716
 - size in bandwidth mode 715
 - size in non-bandwidth mode 716
 - size, optimizing with link statistics 714
 - coexistence with ISC 494
 - comparison to ISC 460
 - concepts 677
 - data flow 681
 - data partitioning 677
 - defining priorities 705
 - definition 677
 - definition of transaction codes 700
 - descriptor 82
 - design considerations 695
 - destination system 686

- MSC (Multiple Systems Coupling)
 - (continued)
 - DFSCMC10
 - physical links, defining 703
 - directed routing, program-to-program switch 694
 - functions compared to ISC 460
 - Generic Resources, VTAM 735
 - IMSIDs
 - in an IMSplex 728
 - IMSplex
 - sharing MSNAME definitions and SYSIDs 724
 - IMSplex with shared queues
 - coexistence 721
 - IMSplexes
 - APPC and OTMA remote transactions 728
 - APPC transactions, sending to an MSC system outside of an IMSplex 734
 - back-end processing of remote APPC or OTMA transactions 729
 - enabling back-end processing of remote APPC and OTMA transactions 733
 - message routing in an IMSplex 722
 - message routing when an IMSplex and MSC coexist 722
 - OTMA transactions, sending to an MSC system outside of an IMSplex 734
 - input system 686
 - intermediate system 686
 - introduction 677
 - ISC facility 460
 - link bandwidth 696
 - link buffer sizes 696
 - link definitions
 - in an IMSplex 724
 - link definitions in an IMSplex
 - deleting 727
 - link statistics
 - assessing optimum link buffer size 714
 - link type, optimum 713
 - links
 - logical 679
 - physical 677
 - local system
 - defined 681
 - logical destinations 684
 - logical link capacity 714
 - logical link paths 679
 - logical link statistics 712
 - logical link statistics and buffer sizes 714
 - logical link statistics, high values 717
 - logical links
 - benchmarking 713
 - buffer sizes and statistics 714
 - capacity and statistics 714
 - deleting from an IMSplex 727
 - high value statistics 717
 - resetting statistics 713

- MSC (Multiple Systems Coupling)
 - (continued)
 - logical links (continued)
 - statistics and buffer sizes 714
 - statistics and link capacity 714
 - statistics for benchmarking 713
 - statistics, high values 717
 - statistics, resetting 713
 - logical links, SDLC link 679
 - LU 6.2 application transactions buffer size 702
 - message routing 682
 - in an MSC-IMSplex configuration 721
 - migration to IMSplex 723
 - minimizing resource consumption 695
 - monitoring and tuning 718
 - MSC conversation failure 700
 - MSNAME definitions
 - sharing in an IMSplex 724
 - MSNAME definitions in an IMSplex
 - deleting 727
 - MSNAME duplication in an IMSplex
 - with shared queues 726
 - operating procedures 712
 - optimum link type 713
 - overview 375
 - overview of network and routing 681
 - performance 712, 718
 - physical links
 - CTC, defining 703
 - defining 701, 702
 - deleting from an IMSplex 727
 - MTM, defining 703
 - types 701
 - VTAM, defining 704
 - pseudoabend U0830
 - avoiding 734
 - Queuing Summary Report 719
 - recoverable versus irrecoverable transactions 62
 - recovering transactions in APPC 62
 - remote system
 - defined 681
 - resetting logical link statistics 713
 - routing path 683
 - security checking
 - intermediate IMS 711
 - security considerations 711
 - serial transaction processing 707
 - shared queues
 - MSNAME duplication 726
 - sharing MSNAME definitions and SYSIDs 724
 - SPA 699
 - standard application programs 50
 - statistics
 - assessing optimum link buffer size 714
 - statistics, for logical links 712
 - support for APPC/IMS 61
 - SYSID (system identifier) 688
 - SYSID tables
 - deleting MSNAME definitions in an IMSplex 727

MSC (Multiple Systems Coupling)
(*continued*)

- SYSIDs
 - cloning in an IMSplex 728
 - managing in an IMSplex 728
 - sharing in an IMSplex 724
- system definition 700
- system definition verification 710
- system identifier (SYSID) 688
- tuning links 712
- utility for verifying names 709
- VTAM Generic Resources 735

MSC Program Routing exit routine 737

MSDB

- OTMA, read only access 786

MSLINK definitions

- deleting from an IMSplex 727

MSLINK macro

- defining 704

MSNAME definitions

- deleting from an IMSplex 727
- in a shared queues group
 - duplication 726
- sharing among systems in an IMSplex 724

MSNAME macro

- logical link paths 679
- logical path name 705
- MSC (Multiple Systems Coupling)
 - Directed Routing 693
 - reports 719
 - for MSC 719
- SYSID keyword 705

MSPLINK definitions

- deleting from an IMSplex 727

MSPLINK macro

- defining CTC physical links 703
- defining MTM physical links 703
- defining physical links 701
 - buffer sizes 702
- defining VTAM physical links 704

MTM (memory-to-memory)

- MSC (Multiple Systems Coupling)
 - physical link type 677
- MSC physical link
 - defining 703

MTO (master terminal operator) 428

- /OPNDST command 428
- assignment of logical link 677
- ISC errors 550
- notification of session rejection 516
- restarting the terminal 449

MULT1 parameter 482, 537

MULT2 parameter 482, 537

multichain input message switches

- restriction 579

multinode persistent sessions (MNPS)

- XRF 415

multiple external subsystems 164

multiple signons 427

- description 114
- naming conventions for 114
- sysplex environment, in 114

multiple systems 695

Multiple Systems Coupling (MSC) 677

- /MSVERIFY command 710
- bandwidth mode, capacity 714

Multiple Systems Coupling (MSC)
(*continued*)

- bandwidth, link 696
- benchmarking logical link
 - performance 713
- buffer sizes 702
- buffer sizes, links 696
- buffers
 - format in bandwidth mode 715
 - format in non-bandwidth mode 716
 - size in bandwidth mode 715
 - size in non-bandwidth mode 716
 - size, optimizing with link statistics 714
- definition of transaction codes 700
- DFSCLM10
 - physical links, defining 703
- Generic Resources, VTAM 735
- IMSplexes
 - APPC and OTMA remote transactions 728
 - APPC transactions, sending to an MSC system outside of an IMSplex 734
 - back-end processing of remote APPC or OTMA transactions 729
 - enabling back-end processing of remote APPC and OTMA transactions 733
 - OTMA transactions, sending to an MSC system outside of an IMSplex 734
- link bandwidth 696
- link buffer sizes 696
- link statistics
 - assessing optimum link buffer size 714
- link type, optimum 713
- logical link statistics 712
- logical link statistics and buffer sizes 714
- logical link statistics and link capacity 714
- logical link statistics, high values 717
- logical links
 - benchmarking 713
 - buffer sizes and statistics 714
 - capacity and statistics 714
 - high value statistics 717
 - resetting statistics 713
 - statistics and buffer sizes 714
 - statistics and link capacity 714
 - statistics for benchmarking 713
 - statistics, high value 717
 - statistics, high values 717
 - statistics, resetting 713
- optimum link type 713
- overview of network and routing 681
- performance, logical links 712
- physical links
 - CTC, defining 703
 - defining 701, 702
 - MTM, defining 703
 - types 701

Multiple Systems Coupling (MSC)
(*continued*)

- physical links (*continued*)
 - VTAM, defining 704
- resetting logical link statistics 713
- security considerations 711
- SPA 699
- statistics
 - assessing optimum link buffer size 714
- statistics, for logical links 712
- system definition 700
- system definition verification 710
- tuning links 712
- VTAM Generic Resources 735

Multiple Systems Verification utility (DFSUMSV0) 693

Multiple Systems Verification utility, using 709

N

NAME macro 499

- defining an ISC session 499
- for remote logical terminals 679

MSNAME macro 709

- used to relate components 479

name uniqueness, resource

- disabling enforcement 384

naming conventions 442

NCP (Network Control Program) 372

NCP (Network Control Programs) 371

negotiable bind 506

network 369

- administration 389
- activities 389
- APPC/IMS, operating with 376
- communications 369
- defining 425
- definition 425, 426, 709
- design considerations 389
- documenting requirements 389
- logical terminal network design 397
- multiple systems, effect of 709
- nonswitched communications network 400
- operating
 - establishing communication 376
 - Finance Communication System 899, 927
 - IMS-CICS 629, 630
 - ISC 503, 504
- operations, preparing for 425
- optional components 369
- planning 389
- restarting after remote takeover 434
- shutting down 504, 935
- starting
 - Finance Communication System 431, 927
- terminating
 - Finance Communication System 934, 936
 - VTAM and NCP parameters 425

network architecture models 753

network role 371, 373

network-qualified LU name 68

- network, communications 369
- networking 695
- NO-OP indicated on LUSTATUS command 548
- node definition, ISC 496, 502
- node name 391
- node user descriptor 106
- non-conversational program 811
- non-persistent socket connections 309
- NONE security level 799
- nongraphic message data 450
- nonswitched communications network 400
- NOSCAN option 945
- notify message 142, 161
- NRESTART command 454
- NTO terminals 401
- null output message
 - purpose 943
 - when not sent 914
 - when sent 949

O

- OASN (origin application sequence number) 164
- ODBA (Open Database Access)
 - accessing IMS databases 747
 - application programs, binding 748
 - APSB security 750
 - connections
 - configuring 741
 - defining security 748
 - interface 747
 - modules
 - which ones to place in JOBLIB 748
 - which ones to place in STEPLIB 748
 - overview 747
 - RAS security 749
 - security
 - APSB 750
 - RAS 749
 - setup 747
- one-phase commit
 - IMS Connect
 - IMS Universal drivers 354
 - IMS TM Resource Adapter 359
 - IMS Universal Database resource adapter 354
 - IMS Universal drivers 354
- online change
 - DISPLAY MODIFY command 433
 - libraries 445
 - Online Change utility 445
- online performance 443
- Open Database
 - IMS Connect
 - timeout specifications 317
- Open Database Access (ODBA)
 - accessing IMS databases 747
 - application programs, binding 748
 - APSB security 750
 - connections
 - configuring 741
 - defining security 748

- Open Database Access (ODBA)
 - (continued)
 - interface 747
 - modules
 - which ones to place in JOBLIB 748
 - which ones to place in STEPLIB 748
 - overview 747
 - RAS security 749
 - security
 - APSB 750
 - RAS 749
 - setup 747
- open systems interconnection 753
- Open Transaction Manager Access
 - Callable Interface (C/I) 887
 - introduction to 887
 - restrictions 890
 - security for 890
 - headers 265
 - sample messages 884
- Open Transaction Manager Access (IMS Open Transaction Manager Access)
 - Callable Interface (C/I)
 - initializing 889
- Open Transaction Manager Access (OTMA) 753
 - ACEE, refreshing 802
 - administering 767
 - asynchronous hold queues
 - security 804
 - benefits 755
 - callout
 - implementing asynchronous callout support 848
 - callout configuration overview 845
 - callout requests 845
 - callout, asynchronous
 - correlating responses 848
 - capabilities 754
 - client
 - initiating protected transactions 829
 - client descriptors
 - DFSOTMA 775
 - CM0 output
 - acknowledgment time out 842
 - timeout for acknowledgment 842
 - commit-then-send output
 - acknowledgment time out 842
 - timeout for acknowledgment 842
 - configuration parameters 767
 - configuring 767
 - coupling facility
 - cross system 753
 - XCF 753
 - cross system coupling facility 753
 - default client values, setting 775
 - descriptors, specifying 771
 - DFSOTMA descriptor 775
 - enabling 767
 - GRNAME 768
 - IMS administration 779
 - IMS environments supported 754
 - installing 767
 - introduction 753

- Open Transaction Manager Access (OTMA) (continued)
 - M descriptors
 - DFSOTMA 775
 - member descriptors
 - DFSOTMA 775
 - member name override, enabling 770
 - message
 - prefix length 785
 - message prefix
 - message-control fields, explanations 859
 - security data fields, explanations 881
 - user data fields, explanations 883
- MSC (Multiple Systems Coupling)
 - processing remote transactions in an IMSplex 728
- OTMA C/I
 - Timing out OTMA C/I sessions 890
- OTMA client descriptor 771
- OTMA client, introduction 813
- OTMA destination descriptor 772
 - defining 774
 - masking destination names 775
- OTMA= 769
- OTMAASY 770
- OTMAMD 770
- OTMANM 768
- OTMASE 769
- OTMASP 770
- parameters
 - GRNAME 768
 - OTMA= 769
 - OTMAASY 770
 - OTMAMD 770
 - OTMANM 768
 - OTMASE 769
 - OTMASP 770
- PROCLIB member DFSPBxxx 767
- program-to-program switches
 - asynchronous output, specifying 770
- protocol
 - overview 757
- response flags
 - delayed acknowledgement support 859
- restrictions 785
- security 799
 - asynchronous hold queues 804
 - general considerations 805
 - levels of checking, defining 769
 - specifying OTMA client security 801
 - specifying OTMA member security 801
 - system definition 801
- security, RACF 799
- server definition 754
- starting
 - defining when OTMA starts 769
- Startup procedures 767
- synchronous callout support 846
- system definition 767
 - security 801

- Open Transaction Manager Access (OTMA) *(continued)*
 - time out CM0 acknowledgment 842
 - transaction expiration
 - in seconds 798
 - overview 797
 - STCK format 798
 - WebSphere MQ
 - synchronized tpipes, defining 770
 - XCF
 - group name, defining 768
 - member name, defining 768
- operator logical paging
 - paged messages 573
 - paging errors 569
 - QXFR FM header 591
 - SLU P, MFS option 916
 - sync points 527
 - VTAM indicators 536
 - when in effect 532
- OPNDST (/OPNDST) command 428
 - relation to MODETBL keyword 428
- OPNDST command
 - logging on 111
 - results of using 120
- OPTACK option
 - Fast Path 923
 - Finance Communication System 954
 - SLU P 920, 951
- options, CICS
 - system definition 620
 - table preparation 621
- OPTIONS= DPAGE or PPAGE
 - MFS DPM 568, 945
 - output message 591
- orderly session termination
 - CICS 630
 - ISC 517
 - VTAM 934
- OTMA
 - /DISPLAY TRANSACTION
 - output format 893
 - application-data section of message
 - prefix 883
 - architected transaction attributes 893
 - asynchronous output
 - IMS Connect 328
 - buffer pool
 - HIOP 794
 - LUMP 794
 - callout
 - SOAP Gateway 850
 - callout requests
 - IMS Connect 328
 - client
 - definition 813
 - client-bid message flow 825
 - clients
 - routing 757
 - CM0 output, managing 840
 - CM0 output, purging 841
 - CM0 output, super member 843
 - CM1 ACK timeout 796
 - CM1 transactions
 - timeout 796
 - commands
 - architected output 757

- OTMA *(continued)*
 - commands *(continued)*
 - IMS 817
 - OTMA terminology 817
 - commit
 - summary of processing 818
 - commit-then-send message flow 827
 - commit-then-send output, managing 840
 - commit-then-send output, purging 841
 - commit-then-send output, rerouting 841
 - commit-then-send output, super member 843
 - conversations, administering
 - IMS 779
 - DBCTL unsupported 786
 - DL/I calls
 - CHNG 806
 - INQY 806
 - PURG 806
 - SETO 806
 - express PCB and program switch 808
 - Front-End Switch, unsupported 786
 - I/O PCB, and program switch 808
 - IMS Connect
 - specifying OTMA ACEE aging value 221
 - timeout specifications 317
 - IMS conversations 779
 - LTERM 761
 - message
 - IMS resynchronization support 831
 - queue data set size 793
 - recoverable 831
 - message encryption, restriction 786
 - message flood detection 789
 - Message Format Service, unsupported 785
 - message prefix
 - cancel resume tpipes request
 - format 880
 - chain flag (TMAMCCHN) 865
 - command type (TMAMCTYP) 861
 - commit-confirmation flag (TMAMCCCI) 861
 - message type (TMAMCMGT) 859
 - no messages on tpipes hold queue
 - format 880
 - overview 815
 - prefix flag (TMAMCPFL) 865
 - processing flag (TMAMCPFG) 864
 - reason code (TMAMCRSC) 866
 - response flag (TMAMCRSI) 860
 - send-sequence number (TMAMCSSN) 866
 - sense code (TMAMCSNC) 866
 - sense code (TMAMCSNS) 866
 - TMAMAGNG (userid aging value) 866
 - TMAMCCCI (commit-confirmation flag) 861

- OTMA *(continued)*
 - message prefix *(continued)*
 - TMAMCCHN (chain flag) 865
 - TMAMCMGT (message type) 859
 - TMAMCPFG (processing flag) 864
 - TMAMCPFL (prefix flag) 865
 - TMAMCRSC (reason code) 866
 - TMAMCRSI (response flag) 860
 - TMAMCSNC (sense code) 866
 - TMAMCSNS (sense code) 866
 - TMAMCSSN (send-sequence number) 866
 - TMAMCTNM (tpipes name) 865
 - TMAMCTYP (command type) 861
 - tpipes name (TMAMCTNM) 865
 - userid aging value (TMAMAGNG) 866
 - message prefix, application-data section 883
 - message prefix, security-data section 880
 - message prefix, user-data section 883
 - messages
 - client-bid flow 825
 - commit-then-send flow 827
 - displaying current workload 787
 - flow, client-bid message 825
 - flow, commit-then-send message 827
 - flow, server-available message 826
 - prefix, contents 825, 826, 827
 - sample flows 824
 - server-available flow 826
 - MSDB, read only access 786
 - performance
 - resource monitoring 786
 - prefix
 - contents 825, 826, 827
 - program switch
 - usage scenarios 807
 - protected conversation messages, restrictions 785
 - protocol commands
 - cancel resume tpipes request
 - format 880
 - no messages on tpipes hold queue
 - format 880
 - overview 786
 - server state 876
 - purging output
 - NAK responses 841
 - RACF security classes 803
 - recoverable messages 831
 - rerouting output
 - NAK responses 842
 - send-only output 841
 - resource management
 - transaction instance block (TIB) 788
 - resource monitoring 786
 - resource recovery services 785
 - resources
 - monitoring 786
 - restart processing 780

OTMA (*continued*)

- resynchronization
 - assumptions 831
 - overview 830
- sample messages 884
- security
 - asynchronous hold queue 804
 - DFSRTUX 805
 - message prefix security specifications 803
 - OTMA Resume TPIPE security exit routine 805
 - specifying 800
- security-data section of message
 - prefix 880
- send-only
 - rerouting output 841
- server state protocol command 876
- server-available message flow 826
- shared queues
 - ALTPCB output, retrieving 785
- shutdown
 - overview 785
- SLU 2 Transaction flow
 - standard 760
- state data
 - REQresynch command format 870
- state-data
 - resume output for tpipe format 878
- super member 843
- terminal control commands
 - unsupported 786
- termination
 - overview 785
- timeout
 - CM1 transactions 796
- timeout interval
 - specifying for OTMA client 796
- timeout, ACK 795
- TMAMAGNG 866
- TMAMCCCI 861
- TMAMCCHN 865
- TMAMCMGT 859
- TMAMCPFG 864
- TMAMCPFL 865
- TMAMCRSC 866
- TMAMCRSI 860
- TMAMCRSQ 866
- TMAMCSEQ 866
- TMAMCSNC 866
- TMAMCSNS 866
- TMAMCSSN 866
- TMAMCTNM 865
- TMAMCTYP 861
- TMAMRTID 867
- transaction attributes segment 893
- user-data section of message
 - prefix 883
- XCF macros, reassembly 785

OTMA (Open Transaction Manager Access) 753

- /START OTMA 769
- ACEE, refreshing 802
- administering 767

OTMA (Open Transaction Manager Access) (*continued*)

- asynchronous hold queues
 - security 804
- benefits 755
- callout
 - implementing asynchronous callout support 848
- callout configuration overview 845
- callout requests 845
- callout, asynchronous
 - correlating responses 848
- capabilities 754
- client
 - in XRF environment 780
 - initiating protected transactions 829
- client descriptors
 - DFSOTMA 775
- CM0 output
 - acknowledgment time out 842
 - timeout for acknowledgment 842
- commit-then-send output
 - acknowledgment time out 842
 - timeout for acknowledgment 842
- configuration parameters 767
- configuring 767
- coupling facility
 - cross system 753
 - XCF 753
- cross system coupling facility 753
- default client values, setting 775
- descriptors, specifying 771
- DFSOTMA descriptor 775
- enabling 767
- GRNAME 768
- IMS administration 779
- IMS application programs 754
- IMS environments supported 754
- installing 767
- introduction 753
- IOPCB 758
- M descriptors
 - DFSOTMA 775
- member descriptors
 - DFSOTMA 775
- member name override, enabling 770
- message
 - prefix length 785
- message prefix
 - message-control fields, explanations 859
 - security data fields, explanations 881
 - user data fields, explanations 883
- message switch 807
- message-control information section of message prefix 855
- MSC (Multiple Systems Coupling)
 - processing remote transactions in an IMSplex 728
- OTMA C/I
 - Timing out OTMA C/I sessions 890
- OTMA client descriptor 771
- OTMA client, introduction 813
- OTMA destination descriptor 772

OTMA (Open Transaction Manager Access) (*continued*)

- defining 774
- masking destination names 775
- OTMA= 769
- OTMAASY 770
- OTMAMD 770
- OTMANM 768
- OTMASE 769
- OTMASP 770
- parameters
 - GRNAME 768
 - OTMA= 769
 - OTMAASY 770
 - OTMAMD 770
 - OTMANM 768
 - OTMASE 769
 - OTMASP 770
- Prerouting exit routine (DFSYPX0) 758
- PROCLIB member DFSPBxxx 767
- program switch 807
- program-to-program switches
 - asynchronous output, specifying 770
- protected messages 811
- protocol
 - overview 757
- response flags
 - delayed acknowledgement support 859
- restrictions 785
- resynchronization protocol 834
- security 799
 - general considerations 805
 - levels of checking, defining 769
 - specifying OTMA client security 801
 - specifying OTMA member security 801
- security, RACF 799
- server definition 754
- starting
 - defining when OTMA starts 769
- Startup procedures 767
- synchronous callout requests 846
- system definition 767
 - security 801
- time out CM0 acknowledgment 842
- transaction expiration
 - in seconds 798
 - overview 797
 - STCK format 798
- WebSphere MQ
 - synchronized tpipes, defining 770
- XCF
 - group name, defining 768
 - member name, defining 768

OTMA administration

- target members 790
- tmembers 790

OTMA C/I

- Timing out OTMA C/I sessions 890

OTMA client

- naming conventions for 814

OTMA client descriptors 771

OTMA destination descriptors 772

- OTMA Destination Resolution exit routine (DFSYPX0)
 - IMSplices 729
 - MSC 729
 - multiple IMS systems 729
- OTMA Resume Tpipe Security exit routine (DFSYRTUX) 776
- OTMA= parameter 769
- OTMAASY option 809
- OTMAASY parameter 770
- OTMAMD parameter 767, 770
- OTMANM parameter 767, 768
- OTMASE parameter 768, 769
- OTMASP parameter 768, 770
- output
 - available at ISC restart 510
 - controlling 452
 - MFS DPM 945
- output algorithms
 - RU chain 580
 - VLVB 580
- output bracketing 949
- output component
 - defining 481
 - relationship to input component 477
 - selection
 - identification number 478, 909
 - modifying 478, 909
 - system messages 564, 910
- output component ID byte, output message 946
- output component protection, extended 918
- output devices, control characters by type 451
- output editing option, SLU P 919
- output errors, MFS online detection 569
- output function management headers, ISC 638
- output message
 - between brackets
 - design considerations 914
 - figure 913
 - how handled 913
 - message switching 954
 - description 943
- Fast Path
 - Finance systems 923
 - SLU P systems 924
- Finance Communication System
 - multiple transmission 943
 - read type field (SMSCRT) 943
 - read-flags field (SMSCRF/SMSCRE) 943
- formatting, activating MFS 567, 949
- ISC bracketing 536
- MFS DPM 568
- null
 - purpose 943
 - when sent 949
- output bracketing 949
- output FM header
 - Finance format 946
 - ID byte 946
 - message descriptor byte (Finance) 946

- output message (*continued*)
 - output FM header (*continued*)
 - message descriptor byte (SLU P) 947
 - MFS data bytes (Finance) 947
 - MFS data bytes (SLU P) 948
 - response requests 950
 - SLU P format 947
 - segmenting 943
 - SLU P system
 - multiple transmission 943
 - read type field (SMSCRT) 943
 - sync point requested, ISC 526
 - temporarily stopping 934
 - types 943
 - when committed 526
- output message structure 364
- output protocols, determining 481
- output response requested by message type 951
- overload avoidance, MSC 696

P

- page advance function (NEXTTP), MFS 916
- page delete function, MFS 568
- page protection state 405
- paging errors
 - ISC 549
 - online detection, MFS 570
- parallel sessions
 - bind requirements 507
 - ISC, defining IMS-to-IMS sessions 494
 - users permitted 478
- parameter
 - OTMAMD 767
 - OTMANM 767
 - OTMASE 768
 - OTMASP 768
- parameters
 - OTMA configuration parameters, summary 767
- partner systems in MSC (Multiple Systems Coupling) linking 679
- PassTicket
 - IMS Connect 210
 - replay protection 212
- passwords
 - IMS Connect mixed-case passwords 210
 - IMS Connect, changing RACF passwords 208
- PCB (program communication block)
 - alternate PCB 479
 - for an LTERM not in I/O PCB 910
 - I/O PCB 479
- PDSE
 - external subsystem module table 149
- performance 786
- Persistent Session Tracking
 - Termination of 421
- persistent socket connections 308
- physical links
 - defining 701
 - buffer sizes 702

- physical links (*continued*)
 - defining (*continued*)
 - CTC 703
 - MTM 703
 - VTAM 704
 - MSC (Multiple Systems Coupling) 677
 - types 701
- physical terminals
 - defining 89
 - device class control 451
 - separating input and output devices 399
- ping
 - IMS Connect availability 347
- PL/I
 - XML conversion support
 - overview 178
- pool manager 445
 - MFS description 445
 - MFSTEST 445
- prefix
 - of OTMA message 815
 - rules 815
 - syntax 855
- Prerouting exit routine (DFSYPX0)
 - basic message flow 758
 - customizing IMS for OTMA 776
 - OTMA usage restrictions 785
- primary error recovery procedure 902
- primary resource name (PRN) 583
- printed output
 - /ASSIGN command 452
 - 3270 printer components 454
 - 3270R and ETO 454
 - ASSIGN (/ASSIGN) command 452
 - candidate printers 454
 - controlling 452, 455
 - operational considerations 455
 - printer component 453
 - shared printers 455
 - spooled output control 453
 - VTAMLIST definitions 454
- printers
 - candidate 454
 - sharing 455
- private context 807
- PRN (primary resource name)
 - ATTACH FM header 583
 - message routing, ISC 484
 - SCHEDULER FM header 583
- processing affinity
 - in an MSC-IMSpdex configuration 723
- processor utilization, MSC 696
- PROFILE security level 799
- program communication block (PCB) 479
- program specification block (PSB)
 - CICS DRA
 - coding guidelines 743
- program switch 807
 - CM0 messages 807
 - CM1 messages 807
 - race condition 809
 - usage scenarios
 - for protected transactions 811

- program switch (*continued*)
 - usage scenarios (*continued*)
 - single-stream 808
 - to multiple programs 809
 - with express PCB 808
 - with OTMAASY option 809
 - without ISRT to I/O PCB 808
- program-to-program (P2P) 807
- program-to-program switch 698, 783
 - conversational transactions 694
 - destination name 690
 - nonconversational transactions 694
 - remote destination verification 698
- protected
 - conversations 807
 - transactions 829
- protected resource, definition 40
- protection
 - display screen 917
 - extended output component 918
- protocol
 - output 481
 - restrictions on IMS-to-IMS sessions 495
- protocols
 - IMS Connect conversational 295
 - send-then-commit, sync
 - level=confirm, ACK response 297
 - send-then-commit, sync
 - level=confirm, NAK response 298
 - send-then-commit, sync
 - level=none, client terminated transaction 296
 - send-then-commit, sync
 - level=none, program terminated transaction 296
 - IMS Connect Send Only 305
 - OTMA
 - server state protocol command 786
- PSB (program specification block)
 - CICS DRA
 - coding guidelines 743
- pseudoabend U0830
 - avoiding 734
- PSTOP state 433
- publications
 - changes xix
 - organization xix
 - titles xix
- PURGE
 - error recovery procedure 542
 - begins 543
 - DFC protocol 542
 - DFC support layer 543
 - ends 543
 - multichain 542
 - single 542
 - spanning chains 543
 - state after selective receiver purge 545
 - state after sender purge 544
 - sender ERP 554
 - purge function, IMS Connect 299
 - PURGE state 433

Q

- QCF 782
- QEC command 934
- QERROR state 404
- QGET FM header
 - description 574
 - format 587
- QGETN FM header
 - description 573
 - format 589
- QLOCK state 404
- QMODEL FM headers
 - CICS 644
 - formats 587
 - QGET 574
 - QGETN 573
 - QPURGE 575
 - QSTATUS 575
 - QXFR 575
 - reply or output 575
 - request or input 573
 - supported by IMS 572, 587
- QPURGE FM header
 - format 590
 - received by IMS 575
- QSTATUS FM header
 - format 590
 - sent by IMS 575
- QSTOP state 433
- qualifying an LU name 68
- queue control facility 781
 - identifying message categories 782
 - support for non-shared queues 781
 - support for shared queues 781
- queue empty indication 547
- queue rotation 538
- queued subsystem 631
- queues
 - balancing group 385
 - logical terminals 399
 - shared 377
- quiesce-at-end-of-chain command, VTAM 934
- QXFR FM header
 - as output 575
 - format 591

R

- race condition
 - avoiding 809
 - defined 809
- race, BIND 507
- RACF
 - ACEE
 - refreshing OTMA ACEEs 802
 - command security 411
 - IMS Connect mixed-case passwords 210
 - IMS Connect support for PassTicket 210
 - IMS Connect support for PassTicket 210
 - IMS Connect, changing passwords 208
 - IMS Connect, default user ID 200

RACF (*continued*)

- IMS Connect, enabling direct support 200
- mixed-case passwords, IMS Connect 210
- OTMA security classes 803
- PassTicket and IMS Connect 210
- passwords, changing IMS Connect 208
- RACF (Resource Access Control Facility)
 - APPC transaction security 74
 - ETO security 374
 - FACILITY class definition 768
 - OTMA restrictions 786
 - OTMA security 799
 - security for MSC (Multiple Systems Coupling) 711
 - unauthorized terminal use 410
- RACROUTE call
 - IMS Connect
 - error responses on calls from sample exits 204
- RAP FM header
 - description 563
 - example 563
 - format 594
 - MFS 576
- Rapid Network Reconnect and IMS Shutdown 422
 - Changing Levels of Support 420
 - defining level of persistent support for VTAM 436
 - defining level RNR support for VTAM 436
 - Defining VTAM for 435
 - Persistent Session Tracking 421
 - Signon Security 423
 - Specifying Levels of Support 420
 - Terminal Reconnect Protocols 422
 - using with XRF or VGR 422
- Rapid Network Reconnect (RNR) planning 420
- RCVYCONV= 407
- RCVYFP= 407
- RCVYRESP= 407
- RCVYSTSN= 407
- RDPN (Return Destination Process Name)
 - ATTACH FM header 584
 - message routing, ISC 484
 - SCHEDULER FM header 584
- read flags field, output messages 943
- read type field, output messages 943
- ready-to-receive (RTR) command 549
- receive-any buffers 427
- recoverable
 - resources 829
 - transactions 831
- recoverable input, acknowledging 920
- recoverable message sent, sequence numbers 929
- recoverable messages, ISC sync points 527
- recoverable versus irrecoverable messages 950, 952
- recoverable versus irrecoverable transactions 62

- recoverable-inquiry transactions, response requirements
 - IMS 950, 953
 - ISC 527, 529
- recovering send-then-commit 819
- recovering transactions in APPC 62
- recovery
 - Fast Path 407
 - resource status 405
- recovery coordinator 139
- recovery environment
 - distributed resources, definition 42
 - local resources, definition 42
- recovery of OTMA messages, selective 781
- recovery procedures 417
- recovery token 145, 164
- RELQ command 934
- remote control of IMS 494
- remote destination name 690
- remote destination verification
 - MSC (Multiple Systems Coupling)
 - conversations 698
 - system integrity 694
- remote LTERMs 691
- Remote Site Recovery (RSR)
 - protected conversations and 45
- remote system
 - MSC (Multiple Systems Coupling)
 - defined 681
- remote takeover
 - determining message loss after 434
 - restarting the network after 434
 - VTAM definition for Transport Manager Subsystem 435
- remote terminal operator (RTO) 428
- reply resynch 834
- reports 719
 - IMS Monitor 719
 - MSC 719
- REPresynch 834
- REPresynch command 870
- REQresynch 834
- REQresynch command 870
- request resynch 834
- reroute function, IMS Connect 301
- reset attached process FM header 563
- resource
 - MSC (Multiple Systems Coupling)
 - considerations 695
 - status classification 405
 - status recovery 405
 - status recovery mode 406
- Resource Access Control Facility (RACF) 711
- resource coordination 145
- resource manager
 - IMS 43
- Resource Manager
 - IMS TM resources, benefits of RM 383
- Resource Manager (RM)
 - IMS TM resources
 - managing 383
- resource manager, definition 40
- resource name uniqueness
 - disabling enforcement 384
- resource recovery services
 - OTMA restrictions 785
 - protected conversation messages, OTMA restrictions 785
- resource structure
 - benefits 383
 - IMS TM resources
 - managing 383
- resource translation table 140
- resource type consistency
 - disabling enforcement 384
- resources
 - TM
 - sharing 384
- response mode
 - dynamic establishment 478
- errors
 - ISC 523
- ISC
 - input 636
- terminal
 - definition 910
 - design considerations 609, 912
 - effects of specifying
 - transaction-dependent 911
 - figure 911
 - introduction 402
 - methods of termination 912
 - restrictions 912
 - sizes of message queue data sets 943
 - sizes of workstation output buffers 943
 - station-by-station 911
 - when activated 911
- response requests, output message 951
- response requirements
 - IMS commands and indicators 934, 950
 - inquiry transactions
 - irrecoverable 525, 953
 - recoverable 525, 953
 - irrecoverable-inquiry transactions 432
 - ISC messages 527
 - LU 6.2 application program 432
 - MFS control requests 955
 - output ISC messages 526
 - verifying IMS receipt 526, 954
 - VTAM indicators and commands 470, 934
- response time 444, 720
- restart
 - coding CICS applications 650
 - output available, ISC 510
- restart and recovery 645
- restart transaction, CICS 618
- RESUME TPIPE
 - alternate client ID 339
 - IRM_RT_ALTCID 339
 - retrieving output for another client 339
- RESUME TPIPE protocol
 - example flows 330
- RESUME TPIPE request
 - IMS Connect 329
- RESUME TPIPE request (*continued*)
 - retrieving asynchronous output
 - IMS Connect 328
- resynchronization
 - deferred 836
 - flow 835
 - OTMA protocol 834
 - OTMA, overview 830
 - sample message 839
- resynchronization, message 507
- Return Destination Process Name 584
- Return Primary Resource Name 584
- RM affinity 406
- RMM (request mod message)
 - format 260
- RNR 712
- RNR (Rapid Network Reconnect)
 - planning 420
- routing
 - messages
 - in an MSC-IMSpIex
 - configuration 721
 - routing code, definition 385
 - routing exit routines
 - link routing 737
 - MSC (Multiple Systems Coupling)
 - conversations 698
 - program routing 737
 - Routing exit routines
 - Terminal Routing exit routine 737
 - routing messages
 - destination name 690
 - ISC examples 485
 - parameters 485
 - SYSIDs 690
 - routing path
 - MSC (Multiple Systems Coupling) 683
- RPRN (Return Primary Resource Name)
 - ATTACH FM header 584
 - message routing, ISC 484
 - SCHEDULER FM header 584
- RQ* messages, LUSTATUS
 - command 546
- RQD* 469, 527, 531
- RQE* 469, 527, 531
- RQR command, SLU P 904
- RR_BACKED_OUT 38
- RR_OK 38
- RR_PROGRAM_STATE_CHECK 38
- RRS 829
- RRS (z/OS Resource Recovery Services) 40
 - description 40
 - resource recovery with 40
- RRSAF 137
- RSHUT command 549
- RSM (request status message)
 - format 261
- RSR (Remote Site Recovery)
 - protected conversations and 45
 - restarting the network 434
- RTO (remote terminal operator) 429
- RTR (ready-to-receive) command 958
 - Fast Path 958
 - IMS functions 958
 - protocol 549

RTR (ready-to-receive) command
(*continued*)
 resetting component status to
 unprotected 918
 summary 549, 958
RTT (resource translation table) 140
RU (request unit) chaining in ISC 542

S

sample
 message
 client-bid 884
 response 884
 transaction 884
 OTMA message 884
sample programs
 using ISC between IMS and
 CICS 669
SBI (stop bracket initiation)
 command 558
 CICS 630
 IMS-CICS session 630
 session shutdown 558
 use in data flow control 459
SC (session control) protocols
 BIND parameters 513
 binding sessions 506
 negotiable versus nonnegotiable
 BIND 506
 parallel session 507
 resolving a race 507
 single session 507
 synchronizing sessions 507
 message resynchronization 507
 commands used 508
 designing procedures 509
 when required 507
 session initiation 516
 completing 516
 ISC VTAM 504
 session states 515
 session termination 517
 abnormal 517
 normal 517
STSN flow 518
 primary-to-secondary half
 session 518
 secondary-to-primary half
 session 519
STSN format 520
XRF complex, establishing an ISC
 session 505
XRF complex, establishing
 connection 505
XRF with MNPS, establishing
 connection 505
XRF with USERVAR, establishing
 connection 505
SCA (system control area), ISC
 support 482
SCHEDULER FM header
 ATTACH 564
 chained message support 564
 example 597
 format 594
 IMS-CICS session 641
SCHEDULER FM header (*continued*)
 introduction 563
 MFS 570
 parameter description 581, 584
SCHEDULER FM headers
 request for asynchronous
 execution 476
scheduler message block (SMB)
 OTMA 758
scheduling
 algorithm 394
 Fast Path messages 385
scratchpad area (SPA) 396
screen protection, Finance
 Communication System
 BID option 917
screens
 protection 405
 unprotected screen option 449
SDF II
 MFS formats 445
SDT (start data traffic) command 516
 completing session initiation 516
 IMS 933
secondary logical unit
 design considerations 472
 first speaker in ISC 472
secondary logical unit type P 910
Secure Sockets Layer (SSL)
 communication with IMS
 Connect 212
 IMS Connect 212
 TCP/IP 212
security
 APPC transactions 74
 Application Transparent Transport
 Layer Security feature
 (AT-TLS) 212
 IMS Connect 199
 Application Transparent Transport
 Layer Security feature
 (AT-TLS) 212
 changing passwords 208
 default RACF user ID 200
 exit routine 201
 mixed-case passwords 210
 overview 179
 PassTicket 210
 RACF, enabling direct
 support 200
 Secure Sockets Layer (SSL) 212
 specifying OTMA ACEE aging
 value 221
 trusted users 220
intermediate IMS
 MSC 711
ISC 463
MSC
 intermediate IMS 711
MSC (Multiple Systems
 Coupling) 711
OTMA 799
 specifying individual member
 security 801
OTMA ACEE aging value
 IMS Connect 221
OTMA clients, overview 754

security (*continued*)
 OTMA RACF security levels 799
 OTMA security-data section of
 message prefix 880
 PassTicket
 IMS Connect 210
 passwords
 changing, IMS Connect 208
 IMS Connect mixed-case
 passwords 210
 Secure Sockets Layer (SSL)
 IMS Connect 212
 TCP/IP
 Application Transparent Transport
 Layer Security feature
 (AT-TLS) 212
 Secure Sockets Layer (SSL) 212
 trusted users
 IMS Connect 220
Security
 IMS Connect
 errors on RACROUTE calls from
 sample exits 204
security considerations
 CICS 636
security levels
 CHECK 800
 FULL 800
 NONE 799
 PROFILE 799
security options 410
 APPC/IMS, RACF 412
 APPC/IMS, SAF 412
 command authorization 411
 DFSCCMD0 Command
 Authorization exit routine 411
 user ID 411
 ETO 412
 user ID 412
 password 412
 user ID 412
 RACF 410
 security profile 411
 DFSCTRN0 Transaction
 Authorization exit routine 411
 user ID 411
 signon verification 410
 transaction authorization 411
 RACF 411
 transaction command 412
 transaction security 74
 RACF 74
 UACC (NONE) 74
security support 202
 PassTicket replay protection 212
security-data section of OTMA message
 prefix 880
selective receiver ERP
 description 549
 sense codes 550
send and receive protocol
 bracketing
 input 942
 output 949
 IMS 937
SEND INVITE EXEC command,
 CICS 612

- SEND LAST EXEC command, CICS 614
- send-only protocol 305
- Send-only protocol
 - with acknowledgment 306
 - with serial delivery 307
- send-only transactions
 - rerouting output 301
- send-then-commit
 - ACK/NAK timeout value 291
 - flow 821
 - with Confirm flow 823
- send-then-commit transactions (CM1)
 - ACK timeout 796
- send/receive and bracket protocol 937
- SEND/RECEIVE EXEC command, CICS 612
- sender ERP, sense codes 555
- sending IMS commands from CICS 635
- sense code 557
 - definition 557
 - error 557
 - received during ISC 557
 - selective receiver ERP 549
 - sender ERP 555, 556
 - sent during ISC 557
- separating input and output devices 399
- sequence numbers
 - definition 815
 - description 513
 - maintaining 509
 - management 520, 930
 - recoverable 816
 - send-sequence numbers 816
 - storage 520
 - use 513, 930
- sequential buffering
 - CICS 744
- serial transactions
 - in an MSC network 707
- server resynch 834
- server state protocol command
 - overview 786
- server-available exchange
 - OTMA message flow 826
 - OTMA message prefix contents 826
- servers
 - OTMA definition 754
- services
 - ATLABCK 829
 - ATRACMT 829
 - ATREINT 829
 - CRGGRM 829
 - CRGSEIF 829
 - CTXBEGC 829
 - CTXEINT 830
 - CTXSWCH 830
- services available to ESAP 147
- session 371
 - binding 506
 - definition 371
 - establishing 376
 - ISC VTAM, establishing 504
 - parallel 507
 - remote control 494
 - requirements to bind ISC 506
 - single 507
- session initiation 431
 - session initiation (*continued*)
 - bind parameters
 - Finance Communication System 928
 - IMS-CICS 629
 - ISC VTAM 504
 - possible session states 515
 - requested by
 - CICS 629
 - master terminal operator,
 - IMS 504, 927
 - network operator, z/OS
 - VTAM 504, 927
 - system definition 504, 927
 - workstation 927, 942
 - step-by-step explanation 928
 - transmission sequence 928
 - ways to request 431
 - session local flag, ATTACH FM
 - header 585
 - session parameters
 - establishing connection 504, 905
 - use 504, 905
 - session termination 517
 - abnormal 517, 630
 - BID option specified, effects 957
 - CICS 630
 - conditional versus unconditional 472
 - definition 517, 934
 - immediate 936
 - definition 630, 934
 - master terminal operator,
 - IMS 517, 934
 - network operator, z/OS
 - VTAM 517, 934
 - station 517
 - workstation 934
 - normal 517, 630
 - orderly
 - definition 935
 - initiated 558, 935
 - requested by CICS 630
 - sequence 559, 935
 - summary (figure) 934
 - symmetrical (SBI/BIS) in ISC 558
- set-and-test-sequence-numbers (STSN)
 - command 518
- setting inquiry parameters 145
- shared queues 377
 - ALTPCB output and IMS
 - Connect 197
 - APPC and OTMA messages
 - processing MSC remote
 - transactions 728
 - benefits 378
 - commit-then-send messages 783
 - definition 377
 - EMH queue option
 - overview 385
 - enabling OTMA 778
 - environment
 - components of, illustration 380
 - operating in, overview 377
 - required components of 379
 - IMS Connect and ALTPCB
 - output 197
- shared queues (*continued*)
 - link definitions in an IMSplex
 - deleting 727
 - message routing
 - in an MSC-IMSplex
 - configuration 721
 - MSC (Multiple Systems Coupling)
 - coexistence 721
 - sharing MSNAME definitions and
 - SYSIDs 724
 - MSC MSNAME definitions
 - duplication 726
 - sharing MSNAME statements and
 - SYSIDs 724
 - MSNAME definitions
 - deleting 727
 - OTMA
 - ALTPCB output, retrieving 785
 - output message queue count 782
 - program-to-program message
 - switch 811
 - program-to-program switch 783
 - pseudoabend U0830
 - avoiding 734
 - send-then-commit messages 783
 - SYSIDs
 - cloning MSC SYSIDs in an
 - IMSplex 728
 - managing MSC SYSIDs in an
 - IMSplex 728
 - sharing MSNAME definitions and
 - SYSIDs 724
 - when an IMSplex and MSC
 - network coexist 728
 - tpipe status 784
 - unsolicited messages 783
 - z/OS system log, role of 379
- sharing printers
 - between systems 455
- shutdown
 - OTMA
 - overview 785
- shutting down an IMS network 504, 936
- SIGNAL command
 - protocol 558
 - VTAM 959
- signing off 377
- signing off, definition 377
- signing on, definition 377
- signon and static terminals 427
- Signon exit routine (DFSSGNX0)
 - ETO
 - associated printing 118
- signon process for external subsystem
 - connections 143
- single session, ISC bind requirements
 - (IMS-to-IMS) 494
- single-stream program switch 808
- SINGLE1 parameter 482, 537
- SINGLE2 parameter 482, 537
- slave-master MSC (Multiple Systems Coupling) relationship 677
- SLU 2 Transaction flow
 - with OTMA 760
- SLU P 899
 - application program
 - converting from Finance 902

SLU P (*continued*)

- application program (*continued*)
 - functions available 900
 - MFS considerations 901
 - XRF considerations 901
- bracket and send/receive protocols 907, 937
- controller-detected errors 955
- facilities 909
- IMS-detected errors 955
- input component definition 910
- message resynchronization
 - controller application program 929
- MFS DPM option 945
- network
 - major parts 899
- sample configuration 900
- SCAN/NOSCAN 902
- session termination
 - immediate 936
- terminals supported 899
- VTAM
 - commands 954
 - commands and indicators 903
 - facilities 902
 - indicators 954
- workstations 900
- XRF 923
- XRF complex, establishing
 - connections 906
- SNA commands 470
- SNA QUIESCE 405
- socket connections
 - IMS Connect 308
 - setting socket type for IMS TM clients 309
 - non-persistent 309
 - persistent 308, 345
 - transaction 309
 - types 308
- sockets
 - IMS Connect
 - processing for transactions 310
 - IMS Connect maximum and UNIX System Services 312
 - maximum number for IMS Connect 311
 - reset threshold 313
 - setting percentages for warnings 313
 - UNIX System Services socket
 - limits 312
 - warning threshold 313
- SPA (scratchpad area) 396
 - characteristics 396
 - MSC (Multiple Systems Coupling)
 - conversations 697
 - size 694, 697
 - transaction code field 397, 403
- specifying external subsystems to IMS 141
- splitting databases, MSC (Multiple Systems Coupling) 697
- SPOOL command
 - output control 453
- spooled output control 453
- SRVresynch 834

SRVresynch command 869

SSL

- IMS Connect
 - sample JCL 218
 - SSL default setup 218
 - SSL initialization 214
 - z/OS key management 214
- SSL (Secure Sockets Layer)
 - IMS Connect 212
 - TCP/IP 212
- START (/START) command
 - starting a SLU P network 927
- start data traffic (SDT) command 516
- start lists, VTAM 430
- START/RETRIEVE EXEC commands, CICS 617
- starting a SLU P network 431
- starting an IMS network 431
 - making IMS ready
 - /START command 503, 927
 - results 503, 927
 - prerequisites 431, 927
- starting workstations 431
- state-data
 - format
 - client-bid commands 867
 - for resume output for hold queue for tpipe 879
 - REPresynch commands 870
 - server-available commands 867
 - SRVresynch commands 869
 - TBresynch commands 871
 - transaction-related
 - information 871
 - section of message prefix 867
- station, user 899
 - definition of as logical unit 899
 - device selection 899
- Statistical Analysis Utility and MSC 718
- statistics
 - bandwidth mode, capacity 714
 - logical links, benchmarking 713
 - logical links, buffer sizes 714
 - logical links, high value statistics 717
 - logical links, link capacity 714
 - logical links, MSC 712
 - logical links, resetting statistics 713
 - MSC link types, optimum 713
 - MSC links
 - buffer size and format in
 - bandwidth mode 715
 - buffer size and format in non-bandwidth mode 716
 - buffers, determining optimum size 714
 - MSC logical links 712
 - MSC logical links, benchmarking 713
 - MSC logical links, buffer sizes 714
 - MSC logical links, high value statistics 717
 - MSC logical links, link capacity 714
 - MSC logical links, resetting statistics 713
- status
 - nonrecoverable 405
 - recoverable 405

status (*continued*)

- resource
 - classification 405
- significant 405
 - command 405
 - end-user 405
- status recovery mode 406
 - GLOBAL 406
 - LOCAL 406
 - NONE 407
 - RCVYCONV= 407
 - RCVYFP= 407
 - RCVYRESP= 407
 - RCVYSTSN= 407
 - resource types 407
- STEPLIB
 - TM and MSC Message Routing and Control user exit routine 735
- STOP (/STOP) command 404
- stop bracket initiation (SBI) command
 - IMS-CICS session 630
- stop bracket initiation command 558
- STOP state 404, 434
- stopped status 167
- stopping
 - workstations 517
- stored procedures 19
- structure
 - creation and deletion 83
 - list
 - definition 379
 - overflow, definition 379
 - primary, definition 379
 - pair, definition 379
- structure recovery data set (SRDS)
 - overview 379
- STSN command 518
 - action code 930
 - action code format 521, 930
 - action required 931
 - controller sequence number, verification 930
 - Finance Communication Systems 930
 - flow 518
 - message resynchronization 930, 933
 - response 932
 - requirements 932, 933
 - response requirements 930
 - summary 933
 - sync points 513
 - VTAM sequence number, verification 930
- STSN function, requirements for ETO 130
- STSN support for ETO devices 131
- SUBPOOL macro 478, 499
- subpools users
 - dynamic allocation 478
 - VTAM 478
- subsystem
 - direct-control 631
 - external, definition 374
 - queued 631
- subsystem connections 142
- Subsystem Startup Service
 - using 160
- subsystem termination 165

- super member
 - IMS Connect support 340
 - managing CM0 output 843
 - OTMA 843
- suspending output from IMS 934
- switch_context service. 830
- sync point
 - definition 395
 - input messages, ISC 525
 - ISC
 - CICS 636
 - input 529
 - output 531
 - requirements 527
- Synch Level 291
 - CONFIRM 292
 - NONE 291
 - SYNCH 292
- synchronization point
 - definition 395
- synchronization, ISC half sessions 525
- synchronous callout requests
 - OTMA configuration overview 845
 - OTMA support 846
 - OTMA support overview 845
- synchronous processing
 - ATTACH FM headers, ISC 562
 - CICS 612
 - definition 476
- syntax diagram
 - how to read xvi
- syntax, message prefix 855
- SYS1.VTAMLST 426
 - ATCCONxx member 430
 - ATCSTRyy member 430
 - defining IMS as an application
 - node 426
 - VTAM nodes 426
- SYSERROR FM header
 - format 595
- SYSID (system identifier) 688
- SYSID keyword
 - logical link paths 679
- SYSID tables
 - deleting MSNAME definitions in an IMSplex 727
- SYSIDs
 - cloning in an IMSplex 728
 - managing in an IMSplex 728
 - sharing among systems in an IMSplex 724
- SYSMSG FM header
 - format 595
- SYSMSG FM headers
 - ATTACH FM header 564
 - sending 565
 - types 565
- Sysplex Distributor
 - IMS Connect 843
 - IMS Connect support for 179
 - OTMA support for 843
- sysplex environment 377
 - definition 377
 - shared queues in 377
- SYSSTAT FM header
 - format 595
- SYSSTAT FM header, format 595

- system
 - definition
 - IMS ISC sample 496, 501
 - messages, length 565
 - system control area (SCA) 482
 - system definition
 - macros 701
 - MSC macros 700
 - MSC
 - general considerations 700
 - implications 709
 - partner 701
 - verifying 709
 - MSC (Multiple Systems Coupling)
 - exit routines 707
 - local 700
 - macros 700
 - setting link priorities 705
- system identification for logical link
 - paths 679
- system identifier (SYSID) 688
- system log
 - MSC 718
- system log, z/OS
 - shared queues and 379
- system resources 786

T

- TBresynch 834
- TBresynch command 871
- TCP/IP
 - calling out to external services from
 - IMS 847
 - IMS Connect 171
 - IMS Connect message structures 237
 - IMS Connect support for RACF
 - PassTicket 210
 - IMS Connect trusted users 220
 - IMS Connect-related TCP/IP
 - settings 365
 - IMS Connect, client
 - communications 172
 - KeepAlive intervals for IMS
 - Connect 315
 - message formats 237
 - purge function for output
 - messages 299
 - reroute function for output
 - messages 301
 - Secure Sockets Layer (SSL) 212
 - security
 - specifying OTMA ACEE aging
 - value 221
 - socket connections 308
 - non-persistent 309
 - persistent 308
 - transaction 309
- terminal
 - physical, defining 89
- terminal control commands
 - OTMA unsupported 786
- TERMINAL macro 500
 - defining an ISC session 500
 - MODETBL= keyword 428
- terminal modes 402
 - conversation mode 402
- terminal modes (*continued*)
 - exclusive 403
 - lock mode 403
 - response mode 402
 - SNA QUIESCE 405
- Terminal Reconnect Protocols 422
- Terminal Routing exit routine 737
- terminals
 - attached through VTAM 389
 - Class 1 with XRF 391
 - Class 2 with XRF 391
 - COMPINOP state 404
 - component protection state 405
 - connections 391
 - conversation mode 402
 - definition 370
 - device class control 451
 - documenting requirements 389
 - ETO and exclusive mode 403
 - exclusive mode 403
 - INOP state 404
 - lock mode 403
 - logical
 - chains 398
 - master terminal 401
 - queues 399
 - relationship to physical
 - terminals 397
 - LU 6.2 terminals and Fast Path 413
 - modes and states 402, 405
 - nonswitched communications
 - network 400
 - NTO 401
 - operating modes 402
 - page protection state 405
 - profiles 425, 427
 - QERROR state 404
 - QLOCK state 404
 - response mode 402
 - screen protection state 405
 - separating input and output
 - devices 399
 - small buffer devices 452
 - SNA QUIESCE 405
 - states 404
 - STOP state 404
 - support for, IMS 389
 - sysplex, in a
 - recovery status 409
 - test mode 403
- terminating communications, MSC
 - conversations 699
- terminating connections with external
 - subsystems 144
- terminating ISC Extension
 - conversations 524
- termination
 - OTMA
 - overview 785
- termination ECB 144, 166
- termination requested by external
 - subsystem 166
- termination, session 517
- test mode 403, 474
- thread, external subsystem 142

- time settings
 - IMS Connect
 - IMS DB clients 317
 - IMS TM clients 317
 - overview 317
 - IMS Connect timeout intervals 318
- timeout
 - IMS Connect
 - IMS DB clients 317
 - IMS TM clients 317
 - input messages 318
 - overview 317
 - OTMA ACK timeout 795
- timer settings 339
- TM and MSC Message Routing and
 - Control user exit routine
 - affinity routing 738
 - IMSplex affinity routing 738
 - JOBLIB 735
 - LINKLIST 735
 - message routing 735
 - STEPLIB 735
- TM resource
 - resource name uniqueness
 - disabling enforcement 384
- TM resources
 - sharing 384
 - sharing, disabling 384
- TMAMAGNG 866
- TMAMALTB 785
- TMAMCCCI 861
- TMAMCCHN 865
- TMAMCMGT 859
- TMAMCPFG 864
- TMAMCPFL 865
- TMAMCRSC 866
- TMAMCRSI 860
- TMAMCRSQ 866
- TMAMCSEQ 866
- TMAMCSNC 866
- TMAMCSNS 866
- TMAMCSSN 866
- TMAMCTNM 865
- TMAMCTYP 861
- TMAMRTID 867
- tmember operand 781
- TMSUFFIX ENVAR keyword 21
- tpipe 761
- tpipe operand 781
- tpipe_Bid resynch 834
- tpipes
 - idle, automatic removal 793
 - removal of idle tpipes 793
 - resource impact 791
 - retrieving asynchronous output for
 - alternate client IDs 339
 - retrieving output
 - IMS Connect 328
- trademarks 963
- traffic between two IMS systems 477, 478
- TRANSACT macro
 - EDIT=ULC 450
 - PRTY= keyword 705
 - translation to uppercase 447
- transaction
 - code, definition 385
- transaction (*continued*)
 - codes, unique 394
 - MSC statistics 719
 - multiple systems 677
 - states 433
- Transaction Analysis utility 718
- transaction code (remote destination) 690
- Transaction Manager
 - introduction 369
- Transaction Manager services 373
- transaction pipe
 - and message flow 763
 - definition 758, 761
 - differences from LTERMs 763
 - differences from UNIX pipes 763
 - flow in full-duplex environment 764
 - in an OTMA client/server environment 762
 - naming conventions for 814
 - non-synchronized 762
 - number a client can create 759
 - removal of idle transaction pipes 762
 - synchronized 762
 - use of queues and message flow 763
 - using 761
- transaction pipes
 - idle, automatic removal 793
 - removal of idle tpipes 793
 - resource impact 791
- transaction socket connections 309
- transaction sync point relationships 525, 636
- transaction types
 - commands 954
 - definitions 432, 952
 - inquiry
 - definition 951
 - recoverable or irrecoverable 952, 953
 - ISC session, during 474
 - message switches
 - IMS 954
 - ISC 579
 - ISC examples 486, 489
 - supported by ISC, list 462
 - test mode, in ISC 474
 - update 952
- transactions 432
 - abends 649
 - commit-then-send 817, 819
 - conversational and OTMA 785
 - expiration
 - IMS Connect support 327
 - OTMA support overview 797
 - OTMA, specifying in seconds 798
 - OTMA, specifying in STCK format 798
 - Fast Path and OTMA 785
 - flow for standard 819
 - IMS, using a nonsynchronized tpipe 832
 - IMS, using a synchronized tpipe 832
 - OTMA
 - specify expiration in seconds 798
 - specify expiration in STCK format 798
- transactions (*continued*)
 - OTMA (*continued*)
 - transaction expiration
 - overview 797
 - OTMA grouping 757
 - protecting 829
 - recoverable 831
 - send-then-commit 817, 821
 - unrecoverable 832
 - transparency option 449
 - Transport Layer Security (TLS) 212
 - Transport Manager Subsystem (TMS)
 - VTAM definition 435
 - trusted users
 - IMS Connect 220
 - tuning
 - buffers 718
 - MSC (Multiple Systems Coupling) environment 695
 - MSC tuning and monitoring 718
 - two phase commit process 145
 - two-phase commit
 - application component 350
 - application server 350
 - commit phase 357
 - communication resource manager (CRM) 350
 - context token 357
 - distributed client flow 358
 - distributed two-phase commit 350
 - enterprise information system 350
 - general description 349
 - global transaction 357
 - IMS Connect
 - one-phase commit for IMS Universal drivers 354
 - IMS Connect support 349
 - IMS Connect support for IMS TM Resource Adapter 357
 - IMS Connect support for IMS Universal drivers 351
 - IMS Connector for Java 350
 - IMS TM Resource Adapter 350
 - IMS Connect support 357
 - one-phase commit optimization 359
 - IMS Universal Database resource adapter
 - one-phase commit optimization 354
 - IMS Universal Database resource adapters 350
 - IMS Universal drivers
 - commit phase 351
 - context token 351
 - distributed client flow 351
 - global transaction 351
 - IMS Connect support 351
 - one-phase commit optimization 354
 - prepare phase 351
 - local option 360
 - prepare phase 357
 - resource adapter 350
 - resource manager 350
 - server distributed syncpoint manager (SDSRM) 350

- two-phase commit (*continued*)
 - transaction manager 350
 - X/Open XA protocol 350
 - z/OS Resource Recovery Services (RRS) 350
- two-phase commit process, definition 41
- TYPE macro
 - defining terminals 390

U

- Unaccessed ETO User Control
 - Blocks 121
- UNBIND command, stopping session initiation 516
- unconditional bracket termination, IMS
 - error handling 955, 956
- Unicode
 - IMS Connect 363
- unit of recovery
 - definition 41
 - in-flight, definition 41
 - indoubt, definition 41
- UNITYPE keyword on TYPE macro 416
- Universal drivers
 - type-2 connectivity
 - CICS, configuring 12
- UNIX pipes 763
- UNIX System Services
 - MAXFILEPROC parameter 312
 - socket limits 312
- unprotected screen option 449
- unrecoverable transactions 832
- user
 - sysplex, in a
 - recovery status 409
- USER 474
- user abend 119ABEND 823
- user authorization processing 143
- user descriptor 82, 106
- user ID
 - IMSpIex, in an
 - recovery status 409
- User ID caching scheme 800
- user message exits
 - description and structures 251
 - support 237
- user workstation 607
 - bracket protocol 533
- user workstation.
 - API (application program interface)
 - CICS asynchronous 632
 - CICS synchronous 632
- user-data section of OTMA message
 - prefix 883
- user-written client application
 - IMS Connect message structures 253
- users
 - operating modes 402
- USSTAB option defined in VTAM 429
- USTOPPED state 434

V

- verification
 - IMS terminal support 389

- verification (*continued*)
 - remote destinations 694
 - transaction definitions across systems 709
- vertical partitioning 677
- Virtual Telecommunications Access Method (VTAM) 371
- VLVB records 580
- VTAM
 - defining the network 425
 - MSC (Multiple Systems Coupling)
 - physical link type 677
 - MSC and VTAM Generic Resources 735
 - MSC physical link
 - defining 704
 - RNR (Rapid Network Reconnect)
 - planning 420
 - SLU P
 - commands 954
 - indicators 954
- VTAM (Virtual Telecommunications Access Method)
 - ACBNAME parameter 704
 - APPL statement 704
 - attached terminals 389
 - BB indicator 481, 546
 - BID command 533, 913
 - binary synchronous communications (BSC) 389
 - BIS command 558
 - LU 6.1 half sessions 558
 - use with CICS 630
 - bracket protocol 907
 - CD (change direction) 907
 - Finance Communication System 907
 - SLU P 907
 - bracket protocols 475
 - buffering 435
 - CANCEL command 541, 958
 - CD indicator 481, 905
 - CHASE command 542
 - COMM macro 704
 - commands and indicators 470, 903
 - devices with MFS support 389
 - EB indicator 481, 547
 - establishing sessions 905
 - Finance Communication System 905
 - SLU P 905
 - facilities 470
 - commands and indicators 470, 903
 - data transmission 470
 - Finance Communication System 902
 - IMS 903
 - ISC 470
 - SLU P 902
 - used by ISC 609, 632
 - generation 426
 - IMS as host subsystem 426
 - LOGON MODE identifiers 428
 - NCP buffer pool values 427
 - storage requirements 427
 - VTAM buffer pool values 427

- VTAM (Virtual Telecommunications Access Method) (*continued*)
 - generation (*continued*)
 - VTAM configurations 430
 - VTAM nodes 426
 - half-duplex protocol, ISC 533
 - IMS, relationship to 372
 - macros 472
 - SEND 472
 - TERMSESS 472
 - missing interrupt handler (MIH) 435
 - MSC (Multiple Systems Coupling)
 - linking 677
 - network role 371
 - node as chosen name 391
 - output buffers 427
 - pacing 435
 - parallel sessions 704
 - RQR command 904, 959
 - SBI command 558
 - LU 6.1 half sessions 558
 - use with CICS 630
 - SDLC using VTAM 704
 - SDT command 516
 - SESSION parameter 704
 - SIGNAL command 558, 959
 - synchronous data link control (SDLC) 389
 - Transport Manager Subsystem, definition for 435
 - UNBIND command 516
- VTAM Generic Resources
 - Multiple Systems Coupling (MSC) 735
- VTAM network administration 369
- VTAM support 756

W

- WebSphere Application Server
 - configuring
 - IMS Universal Database resource adapter, installing 5
 - connection factory
 - defining 5
 - properties 5
 - IMS Universal Database resource adapter
 - application, installing 6
 - IMS Universal drivers
 - configuration 4
- WebSphere Application Server for z/OS
 - classic Java APIs for IMS
 - IVP 29
 - IVP application, installing 31
 - IVP data source, installing 30
 - IVP, testing 31
 - overview 26
 - classpath, setting 32
 - configuring
 - access to IMS 27
 - custom service, installing 28
 - IMS DB resource adapter, installing 28
 - IMS Universal Database resource adapter 7

- WebSphere Application Server for z/OS
 - (continued)
 - configuring (continued)
 - type-2 IMS Universal Database
 - resource adapter, installing 8
 - type-2 IMS Universal drivers 7
 - IMS Universal Database resource adapter
 - application, installing 11
 - classpath, setting 9
 - data source, installing 10
 - running your application
 - application, installing 33
 - classpath, setting 32
 - data source, installing 32
 - prerequisites 32
 - server.policy file 28
 - type-2 connectivityIMS Universal Database resource adapters
 - overview 6
 - type-2 IMS Universal drivers
 - configuring 3
- WebSphere MQ
 - OTMA
 - synchronized tpipes, defining 770
- work areas, creating for ESAP 159
- work unit
 - backed out 507, 645
 - CICS 645
 - committed 507, 645
 - definition 507
 - example 508
 - pending, unilateral decisions 510
 - status at session initiation 515
- workload distribution, MSC 696
- workstation, user
 - API 471, 607
 - CICS asynchronous 609
 - CICS synchronous 609
 - bracket protocol
 - input messages, ISC 533
 - output messages, ISC 536
 - definition as logical unit 478
 - logical unit definition 372
 - terminology 478
- workstations 899
 - Finance Communication System 900
 - SLU P 900

X

- XCF (z/OS cross-system coupling facility)
 - basic OTMA message flow 763
- XML
 - conversion to COBOL 234
 - IMS Connect
 - converting XML to COBOL 231
 - message structures 232
 - XML conversion example 234
 - IMS Connect XML conversion support
 - overview 178
- XRF
 - SLU P 923
- XRF (Extended Recovery Facility) 391
 - APPC/IMS 59, 68
 - Class 1 terminals 419
 - Class 2 terminals 417

- XRF (Extended Recovery Facility)
 - (continued)
 - Class 3 terminals 418
 - class of service 418
 - classes of terminals 416
 - definition 414
 - establishing communication
 - Finance Communication System 906
 - ISC 505
 - SLU P 906
 - system takeover
 - considerations 906
 - master terminals 401
 - MNPS 415
 - processing with OTMA 780
 - recovery modes 408
 - SLU P application program 901
 - takeover considerations
 - Finance Communication System 906
 - SLU P 906
 - terminal support 391, 418
 - terminals in 416

Z

- z/OS
 - IMS Connect support for Sysplex Distributor 179
 - Sysplex Distributor, IMS Connect support 179
 - system log 379
- z/OS application programs
 - accessing IMS databases using ODBA (Open Database Access) 747
 - accessing IMS databases using Open Database Access (ODBA) 747
- z/OS cross-system coupling facility (XCF)
 - macros and OTMA 785
- z/OS program 754
- z/OS Resource Recovery Services
 - and protected transactions 829
 - exits supported by IMS 829
- z/OS Resource Recovery Services (RRS) 40
- z/OS Sysplex Distributor
 - IMS Connect support for 179
- Z2 field in message prefix 451



Product Number: 5635-A02

Printed in USA

SC19-2433-05



Spine information:

IMS Version 11

Communications and Connections

