

IMS
15.3.0

Application Programming
(2025-08-20 edition)



Note

Before you use this information and the product it supports, read the information in [“Notices” on page 883](#).

2025-08-20 edition.

This edition applies to IMS 15 (program number 5635-A06), IMS Database Value Unit Edition, V15.03.00 (program number 5655-DS5), IMS Transaction Manager Value Unit Edition, V15.03.00 (program number 5655-TM4), and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 1974, 2022.**

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

Contents

About this information.....	xvii
Prerequisite knowledge.....	xvii
How new and changed information is identified.....	xvii
How to read syntax diagrams.....	xvii
Accessibility features for IMS 15.3.....	xix
How to send your comments.....	xix
 Part 1. Application programming design.....	1
Chapter 1. Designing an application: Introductory concepts.....	3
Storing and processing information in a database.....	3
Database hierarchy examples.....	5
Your program's view of the data.....	9
Processing a database record.....	11
Tasks for developing an application.....	12
Chapter 2. Designing an application: Data and local views.....	15
An overview of application design.....	15
Identifying application data.....	17
Listing data elements.....	17
Naming data elements.....	19
Documenting application data.....	20
Designing a local view.....	21
Analyzing data relationships.....	22
Local view examples.....	28
Chapter 3. Analyzing IMS application processing requirements.....	35
Defining IMS application requirements.....	35
Accessing databases with your IMS application program.....	36
Accessing data: the types of programs you can write for your IMS application.....	39
DB batch processing.....	39
TM batch processing.....	40
Processing messages: Message Processing Programs.....	40
Processing messages: IMS Fast Path Programs.....	41
Batch message processing: BMPs.....	42
Java message processing: JMPs.....	45
Java batch processing: JBPs.....	45
IMS programming integrity and recovery considerations.....	45
How IMS protects data integrity: commit points.....	45
Planning for program recovery: checkpoint and restart.....	48
Data availability considerations.....	52
Use of STAE or ESTAE and SPIE in IMS programs.....	53
Dynamic allocation for IMS databases.....	54
Chapter 4. Analyzing CICS application processing requirements.....	55
Defining CICS application requirements.....	55
Accessing databases with your CICS application program.....	56
Writing a CICS program to access IMS databases.....	58
Writing a CICS online program.....	58
Using data sharing for your CICS program.....	59

Scheduling and terminating a PSB (CICS online programs only).....	59
Linking and passing control to other programs (CICS online programs only).....	60
How CICS distributed transactions access IMS.....	60
Maximizing the performance of your CICS system.....	61
Programming integrity and database recovery considerations for your CICS program.....	61
How IMS protects data integrity for CICS online programs.....	61
Recovering databases accessed by batch and BMP programs.....	62
Data availability considerations for your CICS program.....	65
Unavailability of a database.....	65
Unavailability of some data in a database.....	66
The SETS or SETU and ROLS functions.....	67
Use of STAE or ESTAE and SPIE in IMS batch programs.....	67
Dynamic allocation for IMS databases.....	67
 Chapter 5. Gathering requirements for database options.....	69
Analyzing data access.....	69
Direct access.....	70
Sequential access.....	74
Accessing z/OS files through IMS: GSAM.....	75
Accessing IMS data through z/OS: SHSAM and SHISAM.....	75
Understanding how data structure conflicts are resolved.....	76
Using different fields: field-level sensitivity.....	76
Resolving processing conflicts in a hierarchy: secondary indexing.....	77
Creating a new hierarchy: logical relationships.....	81
Providing data security.....	84
Keeping a program from accessing the data: data sensitivity.....	85
Preventing a program from updating data: processing options.....	87
Read without integrity.....	89
 Chapter 6. Gathering requirements for message processing options.....	91
Identifying online security requirements.....	91
Analyzing screen and message formats.....	93
An overview of MFS.....	93
An overview of basic edit.....	94
Editing considerations in your application.....	94
Gathering requirements for conversational processing.....	95
What happens in a conversation.....	95
Designing a conversation.....	95
Important points about the scratchpad area (SPA).....	96
Recovery considerations in conversations.....	96
Identifying output message destinations.....	97
The originating terminal.....	97
To other programs and terminals.....	98
 Chapter 7. Designing an application for APPC.....	101
Overview of APPC and LU 6.2.....	101
Application program types.....	101
Application objectives.....	102
Conversation type.....	104
Conversation state.....	104
Synchronization level.....	104
Introduction to resource recovery.....	105
Summary of z/OS Resource Recovery Services support.....	108
Distributed sync point.....	109
Application programming interface for LU type 6.2.....	110
LU 6.2 partner program design.....	111
LU 6.2 flow diagrams.....	111
Integrity tables.....	131

DFSAPPC message switch.....	133
Chapter 8. Testing an IMS application program.....	135
Recommendations for testing an IMS program.....	135
Testing DL/I call sequences (DFSDDLTO) before testing your IMS program.....	135
Using BTS to test your IMS program.....	136
Tracing DL/I calls with image capture for your IMS program.....	136
Using image capture with DFSDDLTO.....	136
Restrictions on using image capture output.....	137
Running image capture online.....	137
Running image capture as a batch job.....	138
Retrieving image capture data from the log data set.....	138
Requests for monitoring and debugging your IMS program.....	139
Retrieving database statistics: the STAT call.....	139
Writing Information to the system log: the LOG request.....	151
What to do when your IMS program terminates abnormally.....	151
Chapter 9. Testing a CICS application program.....	155
Recommendations for testing a CICS program.....	155
Testing your CICS program.....	155
Tracing DL/I calls with image capture.....	156
Requests for monitoring and debugging your CICS program.....	159
What to do when your CICS program terminates abnormally.....	159
Chapter 10. Documenting your application program.....	163
Documentation for other programmers.....	163
Documentation for end users.....	163
Part 2. Application programming for IMS DB.....	165
Chapter 11. Writing your application programs for IMS DB.....	167
Programming guidelines.....	167
Segment search arguments (SSAs).....	168
SSA guidelines.....	170
Multiple qualification statements.....	171
SSAs and command codes.....	174
Considerations for coding DL/I calls and data areas.....	175
Preparing to run your CICS DL/I call program.....	176
Examples of how to code DL/I calls and data areas.....	177
Coding a batch program in assembler language.....	177
Coding a CICS online program in assembler language.....	178
Coding a batch program in C language.....	180
Coding a batch program in COBOL.....	182
Coding a CICS online program in COBOL.....	184
Coding a program in Java.....	188
Coding a batch program in Pascal.....	188
Coding a batch program in PL/I.....	190
Coding a CICS online program in PL/I.....	192
Chapter 12. Defining application program elements for IMS DB.....	195
Formatting DL/I calls for language interfaces.....	195
Assembler language application programming.....	195
C language application programming.....	198
COBOL application programming.....	201
Java application programming for IMS.....	203
Pascal application programming.....	204
Application programming for PL/I.....	206

Specifying the I/O PCB mask.....	208
Specifying the DB PCB mask.....	212
Specifying the AIB mask.....	214
Specifying the AIB mask for ODBA applications.....	216
Specifying the UIB (CICS online programs only).....	219
Specifying the I/O areas.....	221
Formatting segment search arguments (SSAs).....	222
SSA coding rules.....	222
SSA coding formats.....	223
Data areas in GSAM databases.....	226
AIBTDLI interface.....	226
Language specific entry points.....	227
Program communication block (PCB) lists.....	230
The AERTDLI interface.....	231
Language environments.....	232
Special DL/I situations for IMS DB programming.....	233
Application programming with the IMS catalog.....	234
 Chapter 13. Database versioning and application programming.....	 237
 Chapter 14. Establishing a DL/I interface from COBOL or PL/I.....	 239
 Chapter 15. Current position in the database after each call.....	 241
Current position after successful calls.....	241
Position after retrieval calls.....	243
Position after DLET.....	243
Position after REPL.....	245
Position after ISRT.....	245
Current position after unsuccessful calls.....	246
Multiple processing.....	249
Advantages of using multiple positioning.....	253
Multiple DB PCBs.....	255
 Chapter 16. Using IMS application program sync points.....	 257
Commit process.....	257
Two-phase commit in the synchronization process.....	258
Unit of recovery.....	260
DBCTL single-phase commit.....	260
Sync-point log records.....	261
Sync points with a data-propagation manager.....	261
 Chapter 17. Recovering databases and maintaining database integrity.....	 263
Issuing checkpoints.....	263
Restarting your program from the latest checkpoint.....	263
Maintaining database integrity (IMS batch, BMP, and IMS online regions).....	264
Backing out to a prior commit point: ROLL, ROLB, and ROLS.....	264
Backing out to an intermediate backout point: SETS, SETU, and ROLS.....	267
Reserving segments for the exclusive use of your program.....	270
 Chapter 18. Secondary indexing and logical relationships.....	 271
How secondary indexing affects your program.....	271
SSAs with secondary indexes.....	271
Multiple qualification statements with secondary indexes.....	272
DL/I returns with secondary indexes.....	274
Status codes for secondary indexes.....	274
Processing segments in logical relationships.....	274
How logical relationships affect your programming.....	276
Status codes for logical relationships.....	277

Chapter 19. HALDB selective partition processing.....	279
Chapter 20. Processing GSAM databases.....	283
Accessing GSAM databases.....	283
PCB masks for GSAM databases.....	283
Retrieving and inserting GSAM records.....	285
Explicit open and close calls to GSAM.....	287
GSAM record formats.....	287
GSAM I/O areas.....	288
GSAM status codes.....	288
Symbolic CHKP and XRST with GSAM.....	289
GSAM coding considerations.....	289
Origin of GSAM data set characteristics.....	290
DD statement DISP parameter for GSAM data sets.....	291
Extended checkpoint restart for GSAM data sets.....	292
Concatenated data sets used by GSAM.....	293
Specifying GSAM data set attributes.....	293
DLI, DBB, and BMP region types and GSAM.....	294
Chapter 21. Processing Fast Path databases.....	295
Fast Path database calls.....	296
Main storage databases (MSDBs).....	297
Restrictions on using calls for MSDBs.....	297
Data entry databases (DEDBs).....	298
Updating segments: REPL, DLET, ISRT, and FLD.....	298
Checking the contents of a field: FLD/VERIFY.....	299
Changing the contents of a field: FLD/CHANGE.....	301
Example of using FLD/VERIFY and FLD/CHANGE.....	302
Commit-point processing in MSDBs and DEDBs.....	302
Processing DEDBs (IMS and CICS with DBCTL).....	304
Processing Fast Path DEDBs with subset pointer command codes.....	304
Processing DEDBs with a secondary index.....	308
Retrieving location with the POS call (for DEDB only).....	316
Commit-point processing in a DEDB.....	319
P processing option.....	319
H processing option.....	319
Calls with dependent segments for DEDBs.....	320
DEDB DL/I calls to extract DEDB information.....	321
AL_LEN Call.....	325
DI_LEN Call.....	325
DS_LEN Call.....	325
AREALIST Call.....	326
DEDBINFO Call.....	326
DEDSTR Call.....	327
Fast Path coding considerations.....	327
Chapter 22. Writing ODBA application programs.....	329
General application program flow of ODBA application programs.....	329
Server program structure.....	332
Db2 for z/OS stored procedures use of ODBA.....	333
Best practices for Db2 for z/OS stored procedures with ODBA.....	334
Design best practices for ODBA Db2 for z/OS stored procedures.....	334
Writing Db2 for z/OS stored procedures that use ODBA.....	336
Stopping Db2 for z/OS stored procedure threads.....	336
Testing an ODBA application program.....	337
Tracing DL/I calls with image capture to test your ODBA program.....	338
Using image capture with DFSDDLTO to test your ODBA program.....	338

Running image capture online.....	339
Retrieving image capture data from the log data set.....	339
Requests for monitoring and debugging your ODBA program.....	340
What to do when your ODBA program terminates abnormally.....	340
Recommended actions after an abnormal termination of an ODBA program.....	340
Diagnosing an abnormal termination of an ODBA program.....	341
Chapter 23. Programming with the IMS support for DRDA.....	343
DDM commands for data operations with the IMS support for DRDA.....	344

Part 3. Application programming for IMS TM.....347

Chapter 24. Defining application program elements for IMS TM.....	349
Formatting DL/I calls for language interfaces.....	349
Application programming for assembler language.....	349
Application programming for C language.....	352
Application programming for COBOL.....	355
Java application programming for IMS.....	357
Application programming for Pascal.....	357
Application programming for PL/I.....	360
Relationship of calls to PCB types.....	362
Specifying the I/O PCB mask.....	363
Specifying the alternate PCB mask.....	367
Specifying the AIB mask.....	367
Specifying the I/O areas.....	370
AIBTDLI interface.....	370
Specifying language-specific entry points.....	371
Program communication block (PCB) lists.....	373
Language environments.....	374
Special DL/I situations for IMS TM programming.....	375
Chapter 25. Message processing with IMS TM.....	377
How your program processes messages.....	377
Message types.....	377
When a message is processed.....	380
Results of a message: I/O PCB.....	381
How IMS TM edits messages.....	381
Printing output messages.....	382
Using Basic Edit.....	382
Using Intersystem Communication Edit.....	383
Using Message Format Service.....	383
Using LU 6.2 User Edit exit routine (optional).....	390
Message processing considerations for DB2.....	390
Sending messages to other terminals and programs.....	391
Sending messages to other terminals.....	391
Sending messages to other IMS application programs.....	393
How the VTAM I/O facility affects your VTAM terminal.....	395
Communicating with other IMS TM systems using Multiple Systems Coupling.....	395
Implications of MSC for program coding.....	395
Receiving messages from other IMS TM systems.....	396
Sending messages to alternate destinations in other IMS TM systems.....	397
IMS conversational processing.....	398
A conversational example.....	398
Conversational structure.....	399
Replying to the terminal.....	403
Conversational processing using ROLB, ROLL, and ROLS.....	403
Passing the conversation to another conversational program.....	404

Message switching in APPC conversations.....	406
Processing conversations with APPC.....	407
Ending the APPC conversation.....	408
Coding a conversational program.....	408
Standard IMS application programs.....	408
Modified IMS application programs.....	409
CPI-C driven application programs.....	409
Processing conversations with OTMA.....	410
Backing out to a prior commit point: ROLL, ROLB, and ROLS calls.....	410
Comparison of ROLB, ROLL, and ROLS.....	411
ROLL.....	412
ROLB.....	412
ROLS.....	413
Backing out to an intermediate backout point: SETS/SETU and ROLS.....	414
Writing message-driven programs.....	416
Coding DC calls and data areas.....	416
Before coding your program.....	417
MPP code examples.....	417
Message processing considerations for DB2.....	423
 Chapter 26. IMS Spool API.....	425
Managing the IMS Spool API overall design.....	425
IMS Spool API design.....	425
Sending data to the JES spool data sets.....	425
IMS Spool API performance considerations.....	426
IMS Spool API application coding considerations.....	427
Understanding parsing errors.....	429
Diagnosis examples.....	430
Understanding allocation errors.....	432
Understanding dynamic output for print data sets.....	433
Sample programs using the Spool API.....	434
 Chapter 27. IMS Message Format Service.....	437
Advantages of using MFS.....	437
MFS control blocks.....	438
MFS examples.....	438
Relationship between MFS control blocks and screen format.....	442
Overview of MFS components.....	444
Devices and logical units that operate with MFS.....	444
Using distributed presentation management (DPM).....	446
 Chapter 28. Callout requests for services or data.....	449
Callout request approaches.....	450
Resume tpipe protocol.....	452
Implementing the synchronous callout function.....	453
Control data in synchronous callout requests.....	456
Implementing the asynchronous callout function.....	457
 Part 4. Application programming for EXEC DLI.....	459
 Chapter 29. Writing your application programs for EXEC DLI.....	461
Programming guidelines.....	461
Coding a program in assembler language.....	462
Coding a program in COBOL.....	465
Coding a program in PL/I.....	468
Coding a program in C.....	471
Preparing your EXEC DLI program for execution.....	475

Translator, compiler, and binder options required for EXEC DLI.....	476
Chapter 30. Defining application program elements.....	477
Specifying an application interface block (AIB).....	477
Specifying the DL/I interface block (DIB).....	477
Defining a key feedback area.....	480
Defining I/O areas.....	481
Chapter 31. EXEC DLI commands for an application program.....	483
PCBs and PSB.....	483
Chapter 32. Recovering databases and maintaining database integrity.....	485
Issuing checkpoints in a batch or BMP program.....	485
Restarting your program and checking for position.....	486
Backing out database updates dynamically: the ROLL and ROLB commands.....	486
Using intermediate backout points: the SETS and ROLS commands.....	486
Chapter 33. Processing Fast Path databases.....	487
Processing Fast Path DEDBs with subset pointer options.....	487
Preparing to use subset pointers.....	489
Designating subset pointers.....	489
Subset pointer options.....	489
Subset pointer status codes.....	496
The POS command.....	496
Locating a specific sequential dependent segment.....	497
Locating the last inserted sequential dependent segment.....	497
Identifying free space with the POS command.....	498
The P processing option.....	498
Chapter 34. Comparing command-level and call-level programs.....	499
DL/I calls for IMS and CICS.....	499
Comparing EXEC DLI commands and DL/I calls.....	499
Comparing command codes and options.....	501
Chapter 35. Data availability enhancements.....	503
Part 5. Application programming for SQL.....	505
Chapter 36. SQL considerations and restrictions for COBOL.....	507
Chapter 37. Writing application programs for SQL.....	509
Coding SQL statements in application programs: General information.....	509
Defining the items that your program can use to check whether an SQL statement executed successfully.....	509
Defining SQL descriptor areas.....	510
Declaring host variables and indicator variables.....	510
Using SQL statements in your application.....	511
Checking the execution of SQL statements.....	521
Coding SQL statements in COBOL application programs.....	524
Defining the SQL communications area in COBOL.....	524
Defining SQL descriptor areas in COBOL.....	524
Declaring host variables and indicator variables in COBOL.....	525
Equivalent SQL and COBOL data types.....	531
SQL statements in COBOL programs.....	533
SQL aggregate functions supported for COBOL.....	536
Adding and modifying data.....	538
Inserting rows.....	538

Updating segment data.....	539
Deleting data from segments.....	540
Accessing data.....	541
Retrieving data by using the SELECT statement.....	541
Retrieving a set of rows by using a cursor.....	545
Commit or roll back data.....	548
Preparing an application to run on IMS.....	548
Processing SQL statements.....	548
Part 6. Java application development for IMS.....	551
Chapter 38. IMS solutions for Java development overview.....	553
Chapter 39. Comparison of hierarchical and relational databases.....	555
Chapter 40. Programming with the IMS Universal drivers.....	559
IMS Universal drivers overview.....	559
Distributed and local connectivity with the IMS Universal drivers.....	560
Configuring JAXB with JDK 9 or higher.....	563
Comparison of IMS Universal drivers programming approaches for accessing IMS.....	564
Support for variable-length database segments with the IMS Universal drivers.....	565
Support for flattening complex structures.....	566
Generating the runtime Java metadata class.....	567
Hospital database example.....	568
Programming using the IMS Universal Database resource adapter.....	571
Overview of the IMS Universal Database resource adapter.....	571
Transaction types and programming interfaces supported by the IMS Universal Database resource adapter.....	571
Software configurations supported by the IMS Universal Database resource adapter.....	572
Connecting to IMS with the IMS Universal Database resource adapter.....	573
Sample EJB application using the IMS Universal Database resource adapter CCI programming interface.....	584
Accessing IMS data with the DLIIInteractionSpec class.....	585
Accessing IMS data with the SQLInteractionSpec class.....	589
Accessing IMS data with the IMS Universal JCA/JDBC driver.....	592
Programming with the IMS Universal JDBC driver.....	594
Supported drivers for JDBC.....	594
Connecting to IMS using the IMS Universal JDBC driver.....	595
Sample application for the IMS Universal JDBC driver.....	607
Using ByteBuffers with the IMS Universal JDBC Driver.....	608
Using the removeInvalidCaseFields property.....	610
Using the expandArrayResultSet property.....	611
Writing SQL queries to access an IMS database with the IMS Universal JDBC driver.....	613
Writing DL/I calls to access an IMS database with the IMS Universal JDBC driver.....	628
IMS Universal JDBC driver support for XML.....	631
Data transformation support for JDBC.....	635
Programming with the IMS Universal DL/I driver.....	641
Basic steps in writing a IMS Universal DL/I driver application.....	642
Java packages for IMS Universal DL/I driver support.....	642
Connecting to an IMS database by using the IMS Universal DL/I driver.....	643
IMS Universal DL/I driver interfaces for executing DL/I operations.....	645
Inspecting the PCB status code and related information using the com.ibm.ims.dli.AIB interface.....	664
Committing or rolling back DL/I transactions.....	665
Accessing dynamic arrays with the Universal Drivers using the DBArrayElementSet class.....	667
Configuring the IMS Universal drivers for SSL support.....	668

Configuring the IMS Universal Database resource adapter for SSL support in a container-managed environment.....	669
Configuring IMS Universal drivers for SSL support in a stand-alone environment.....	669
Tracing IMS Universal drivers applications.....	670
Chapter 41. Programming Java dependent regions.....	673
Overview of the IMS Java dependent regions.....	673
Programming with the IMS Java dependent region resource adapter.....	674
Developing JMP applications with the IMS Java dependent region resource adapter.....	675
Developing JBP applications with the IMS Java dependent region resource adapter.....	683
IMS Java dependent region resource adapter support for ICAL callout with control data.....	690
Programming with the Callout API.....	692
Program switching in JMP and JBP applications.....	694
IBM Enterprise PL/I for z/OS and Java language interoperability.....	701
IBM Enterprise COBOL for z/OS interoperability with JMP and JBP applications.....	702
IBM Enterprise COBOL for z/OS backend applications in a JMP or JBP region.....	703
IBM Enterprise COBOL for z/OS frontend applications in a JMP or JBP region.....	703
Accessing Db2 for z/OS databases from JMP or JBP applications.....	703
Chapter 42. 31-bit COBOL and 64-bit Java interoperability.....	705
Part 7. PL/I top-down development for IMS Enterprise Suite SOAP Gateway web services.....	711
Chapter 43. WSDL-to-PL/I segmentation APIs for adding business logic in generated PL/I templates.....	713
Chapter 44. Sample of a generated PL/I application template.....	715
Chapter 45. Trace output for WSDL-to-PL/I segmentation APIs.....	717
Chapter 46. Limitations and restrictions of the segmentation APIs.....	719
Part 8. IMS Transaction Manager Resource Adapter	721
Chapter 47. IMS Transaction Manager Resource Adapter overview.....	723
Components of the IMS TM resource adapter.....	723
Runtime process of the IMS TM resource adapter	724
IMS TM resource adapter features.....	724
New features in IMS TM Resource Adapter Version 15.....	725
Supported platforms.....	726
Supported software configurations.....	726
Requirements for the IMS TM resource adapter.....	727
Restrictions for the IMS TM resource adapter.....	727
WebSphere Application Server platform configurations and communication protocol considerations.....	728
Chapter 48. Installing the IMS TM resource adapter runtime component.....	731
Preparing to use the IMS TM resource adapter.....	731
Potential migration issues for the IMS TM resource adapter.....	732
Updating the IMS TM resource adapter.....	732
Extracting the compressed file for installation on distributed platforms.....	733
Extracting the compressed file for installation on z/OS.....	734
Verifying file contents for the IMS TM resource adapter runtime component.....	734
Installing the resource adapter on WebSphere Application Server.....	735
Creating a connection factory in WebSphere Application Server.....	736
Installing the resource adapter on WebSphere Liberty servers.....	737

Configuring the connection factory for WebSphere Liberty servers.....	738
Verifying installation by using the installation verification program.....	739
Prerequisites for running the IVP.....	739
Deploying the IVP EAR file in the Java EE application server.....	740
Running the IMS TM resource adapter IVP.....	742
Running the IMS TM resource adapter callout IVP sample	742
Deploying the sample application on WebSphere Application Server to process callout requests	744
Deploying the sample application on WebSphere Liberty servers to process callout requests	746
Running the IMS host callout IVP application.....	747
Installing IMS TM resource adapter service and updates.....	748
Configuring for resource workload routing.....	749
 Chapter 49. Developing an application for use with the IMS TM resource adapter.....	751
Interacting with the IMS Transaction Manager.....	751
Programming models.....	752
Commit mode and sync level processing.....	782
Socket connections.....	783
IMSInteractionSpec property configuration.....	790
Submitting commands to IMS.....	791
Configuring IMS connection factories.....	791
TCP/IP connections to IMS Connect.....	792
IMS connection factory.....	793
Input and output message formats.....	793
Measuring IMS TM resource adapter performance.....	797
Securing interactions with the IMS Transaction Manager.....	797
IMS TM resource adapter security.....	797
Container-managed EIS signon.....	799
Component-managed EIS signon.....	802
Secure Sockets Layer (SSL) support.....	804
Changing RACF passwords.....	808
Securing message retrieval from IMS hold queues.....	809
Enabling support for distributed network security credentials.....	810
IMS TM resource adapter timeouts.....	814
Execution timeouts.....	815
Socket timeouts.....	817
Other types of timeouts.....	819
Conversational programs.....	820
Client-managed and IMS Connect-managed conversation state programming models.....	821
Orphaned IMS conversation.....	821
Business process choreography applications.....	822
Enabling your Java client for IMS conversational transactions.....	822
Processing global transactions.....	824
Global transaction support with two-phase commit.....	824
Global transaction and two-phase commit support process.....	825
Global transaction support in client applications.....	826
Two-phase commit environment recommendations.....	827
Other transaction support.....	827
Common Client Interface (CCI).....	828
Sample CCI application code.....	830
Samples and tutorials.....	831
 Chapter 50. Running your application on a stand-alone WebSphere Application Server.....	833
Installing your EAR file on WebSphere servers.....	833
 Chapter 51. Diagnosing problems.....	835
Diagnosing IVP failures.....	835
Diagnosing problems accessing IMS from Java applications.....	836

Diagnosing problems with callout requests.....	837
Java exceptions that involve output messages.....	837
Logging and tracing IMS TM resource adapter information.....	838
Logging and tracing in WebSphere Application Server.....	838
Logging and tracing in WebSphere Liberty.....	839
Creating a stand-alone Logger with output sent to a file.....	839
Analyzing the trace data.....	840
IMS TM resource adapter messages and exceptions.....	842
Other exceptions and error messages.....	866
J2CA0056I.....	866
WLTC0017E.....	867
HWSP1445E.....	867
HWSSSL00E.....	868
Chapter 52. Reference information.....	869
IMS connection factory properties.....	869
Client ID (clientID).....	869
CM0 dedicated (CM0Dedicated).....	869
Data store name (dataStoreName).....	870
Group name (groupName).....	870
Host name (hostName).....	870
Password (password).....	870
Password phrase (passwordPhrase).....	870
Port number (portNumber).....	870
SSL enabled (SSLEnabled).....	870
SSL encryption type (SSLEncryptionType).....	870
SSL keystore name (SSLKeyStoreName).....	871
SSL keystore password (SSLKeyStorePassword).....	871
SSL truststore name (SSLTrustStoreName).....	872
SSL truststore password (SSLTrustStorePassword).....	872
SSL protocol (SSLProtocol).....	872
User name (userName).....	872
IMS interaction specification properties.....	872
Alternate client ID (altClientID).....	872
Async output available (asyncOutputAvailable).....	873
Callout request type (calloutRequestType).....	873
Conversation ended (convEnded).....	873
Conversation ID (convID).....	873
Commit mode (commitMode).....	874
CM0 response (CM0Response).....	874
Execution timeout (executionTimeout).....	874
Ignore PURG call (ignorePURGCall).....	875
IMS request type (imsRequestType).....	875
Interaction verb (interactionVerb).....	876
Lterm name (ltermName).....	877
Map name (mapName).....	877
Purge async output (purgeAsyncOutput).....	878
Reroute (reRoute).....	878
Reroute name (reRouteName).....	878
Resume tpipe network security credentials (resumeTpipeNSC).....	879
Socket timeout (socketTimeout).....	879
Synchronous callout correlator token (syncCalloutCorrelatorToken).....	879
Synchronous callout status code (syncCalloutStatusCode).....	879
Sync level (syncLevel).....	879
Transaction expiration (transExpiration).....	880
Transaction tracking ID (trckID).....	880
Use conversation ID (useConvID).....	880
Java API specifications.....	881

Notices.....	883
Programming interface information.....	884
Trademarks.....	884
Terms and conditions for product documentation.....	885
IBM Online Privacy Statement.....	885
Bibliography.....	887
Index.....	889

About this information

These topics provide guidance information for writing application programs that access IMS databases or IMS transactions. The following topics describe how to gather and analyze program requirements, and how to develop and debug IMS application programs. They also describe how to use different programming languages to issue DL/I calls, and include information about the IMS solutions for SQL and Java™ development. They also describe how to use different programming languages to issue EXEC DL/I calls. Application programming interface (API) information is in *IMS Version 15.3 Application Programming APIs*.

This information is available in [IBM® Documentation](#).

Prerequisite knowledge

This information is a guide to IMS application programming for any of the following environments:

- IMS Database Manager (IMS DB), including IMS Database Control (DBCTL)
- IMS Transaction Manager (IMS TM)
- CICS® EXEC DLI
- WebSphere® Application Server for z/OS®
- WebSphere Application Server for distributed platforms
- Java dependent regions (JMP and JBP)
- Any environment for stand-alone Java application development

This book provides guidance information for writing application programs that access IMS databases or process IMS messages. It also describes how to use different programming languages to make DL/I, EXEC DLI, or JDBC calls that interact with IMS. API (application programming interface) information is in *IMS Version 15.3 Application Programming APIs*.

To learn about z/OS, see [z/OS Basic Skills](#). For more resources, see [IBM Z Education and Training](#).

To learn about IMS, see the IBM Press publication *An Introduction to IMS*, the resources listed for [IBM Information Management System](#), and the variety of options available in [IBM Training](#).

How new and changed information is identified

For most IMS library PDF publications, information that is added or changed after the PDF publication is first published is denoted by a character (revision marker) in the left margin. 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.




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.




- 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.



- 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 15.3

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 15.3. 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 15.3 ISPF panel functions by using a keyboard or keyboard shortcut keys.

For information about navigating the IMS 15.3 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 15.3 is available in IBM Documentation.

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:

- Submit a comment by using the DISQUS commenting feature at the bottom of any [IBM Documentation](#) topic.
- Send an email to imspubs@us.ibm.com. Be sure to include the book title.
- Click the **Contact Us** tab at the bottom of any [IBM Documentation](#) topic.

To help us respond quickly and accurately, please include as much information as you can about the content you are commenting on, where we can find it, and what your suggestions for improvement might be.

Part 1. Application programming design

To design an application program for IMS, you need to identify the application data and analyze requirements for application processing. You may also need to perform other tasks, such as gathering requirements for database and message processing options, and testing an application program.

Chapter 1. Designing an application: Introductory concepts

This section provides an introduction to designing application programs. It explains some basic concepts about processing a database, and gives an overview of the tasks covered in this information.

Storing and processing information in a database

The advantages of storing and processing data in a database are that all of the data needs to appear only once and that each program must process only the data that it needs.

One way to understand this is to compare three ways of storing data: in separate files, in a combined file, and in a database.

Storing data in separate files

If you keep separate files of data for each part of your organization, you can ensure that each program uses only the data it needs, but you must store a lot of data in multiple places simultaneously. Problems with keeping separate files are:

- Redundant data takes up space that could be put to better use
- Maintaining separate files can be difficult and complex

For example, suppose that a medical clinic keeps separate files for each of its departments, such as the clinic department, the accounting department, and the ophthalmology department:

- The clinic department keeps data about each patient who visits the clinic, such as:

- Identification number
- Name
- Address
- Illnesses
- Date of each illness
- Date patient came to clinic for treatment
- Treatment given for each illness
- Doctor that prescribed treatment
- Charge for treatment

- The accounting department also keeps information about each patient. The information that the accounting department might keep for each patient is:

- Identification number
- Name
- Address
- Charge for treatment
- Amount of payments

- The information that the ophthalmology department might keep for each patient is:

- Identification number
- Name
- Address
- Illnesses relating to ophthalmology
- Date of each illness
- Names of members in patient's household
- Relationship between patient and each household member

If each of these departments keeps separate files, each department uses only the data that it needs, but much of the data is redundant. For example, every department in the clinic uses at least the patient's number, name, and address. Updating the data is also a problem, because if a department changes a piece of data, the same data must be updated in each separate file. Therefore, it is difficult to keep the data in each department's files current. Current data might exist in one file while defunct data remains in another file.

Storing data in a combined file

Another way to store data is to combine all the files into one file for all departments to use. In the medical example, the patient record that would be used by each department would contain these fields:

- Identification number
- Name
- Address
- Illnesses
- Date of each illness
- Date patient came to clinic for treatment
- Treatment given for each illness
- Doctor that prescribed treatment
- Charge for treatment
- Amount of payments
- Names of members in patient's household
- Relationship between patient and each household member

Using a combined file solves the updating problem, because all the data is in one place, but it creates a new problem: the programs that process this data must access the entire file record to get to the part that they need. For example, to process only the patient's number, charges, and payments, an accounting program must access all of the other fields also. In addition, changing the format of any of the fields within the patient's record affects all the application programs, not just the programs that use that field.

Using combined files can also involve security risks, because all of the programs have access to all of the fields in a record.

Storing data in a database

Storing data in a database gives you the advantages of both separate files and combined files: all the data appears only once, and each program has access to the data that it needs. This means that:

- When you update a field, you do it in one place only.
- Because you store each piece of information only in one place, you cannot have an updated version of the information in one place and an out-of-date version in another place.
- Each program accesses only the data it needs.
- You can prevent programs from accessing private or secured information.

In addition, storing data in a database has two advantages that neither of the other ways has:

- If you change the format of part of a database record, the change does not affect the programs that do not use the changed information.
- Programs are not affected by how the data is stored.

Because the program is independent of the physical data, a database can store all the data only once and yet make it possible for each program to use only the data that it needs. In a database, what the data looks like when it is stored is different from what it looks like to an application program.

Database hierarchy examples

In an IMS DB, a record is stored and accessed in a hierarchy. A hierarchy shows how each piece of data in a record relates to other pieces of data in the record.

IMS connects the pieces of information in a database record by defining the relationships between the pieces of information that relate to the same subject. The result is a database hierarchy.

Medical hierarchy example

The medical database shown in following figure contains information that a medical clinic keeps about its patients. The hierarchies used in the medical hierarchy example are used with full-function databases and Fast Path data entry databases (DEDBs).




Figure 1. Medical hierarchy

Each piece of data represented in the figure above is called a *segment* in the hierarchy. Each segment contains one or more *fields* of information. The PATIENT segment, for example, contains all the information that relates strictly to the patient: the patient's identification number, name, and address.

Definitions: A *segment* is the smallest unit of data that an application program can retrieve from the database. A *field* is the smallest unit of a segment.

The PATIENT segment in the medical database is the *root segment*. The segments below the root segment are the *dependents*, or children, of the root. For example, ILLNESS, BILLING, and HOUSHOLD are all children of PATIENT. ILLNESS, BILLING, and HOUSHOLD are called direct dependents of PATIENT; TREATMNT and PAYMENT are also dependents of PATIENT, but they are not direct dependents, because they are at a lower level in the hierarchy.

A *database record* is a single root segment (root segment *occurrence*) and all of its dependents. In the medical example, a database record is all of the information about one patient.

Definitions: A *root segment* is the highest-level segment. A dependent is a segment below a root segment. A root segment occurrence is a database record and all of its dependents.

Each database record has only one root segment occurrence, but it might have several occurrences at lower levels. For example, the database record for a patient contains only one occurrence of the PATIENT segment type, but it might contain several ILLNESS and TREATMNT segment occurrences for that patient.

The tables that follow show the layouts of each segment in the hierarchy.

The segment's field names are in the first row of each table. The number below each field name is the length in bytes that has been defined for that field.

• PATIENT Segment

The following table shows the PATIENT segment.

It has three fields:

- The patient's number (PATNO)
- The patient's name (NAME)
- The patient's address (ADDR)

PATIENT has a unique key field: PATNO. PATIENT segments are stored in ascending order based on the patient number. The lowest patient number in the database is 00001 and the highest is 10500.

Table 1. PATIENT segment

Field name	Field length
PATNO	10
NAME	5
ADDR	30

• ILLNESS Segment

The following figure shows the ILLNESS segment.

It has two fields:

- The date when the patient came to the clinic with the illness (ILLDATE)
- The name of the illness (ILLNAME)

The key field is ILLDATE. Because it is possible for a patient to come to the clinic with more than one illness on the same date, this key field is non-unique, that is, there may be more than one ILLNESS segment with the same (an equal) key field value.

Usually during installation, the database administrator (DBA) decides the order in which to place the database segments with equal or no keys. The DBA can use the RULES keyword of the SEGM statement of the DBD to specify the order of the segments.

For segments with equal keys or no keys, RULES determines where the segment is inserted. Where RULES=LAST, ILLNESS segments that have equal keys are stored on a first-in-first-out basis among those with equal keys. ILLNESS segments with unique keys are stored in ascending order on the date field, regardless of RULES. ILLDATE is specified in the format YYYYMMDD.

Table 2. ILLNESS segment

Field name	Field length
ILLDATE	8
ILLNAME	10

• TREATMNT Segment

The following table shows the TREATMNT segment.

It contains four fields:

- The date of the treatment (DATE)
- The medicine that was given to the patient (MEDICINE)
- The quantity of the medicine that the patient received (QUANTITY)
- The name of the doctor who prescribed the treatment (DOCTOR)

The TREATMNT segment's key field is DATE. Because a patient may receive more than one treatment on the same date, DATE is a non-unique key field. TREATMNT, like ILLNESS, has been specified as having RULES=LAST. TREATMNT segments are also stored on a first-in-first-out basis. DATE is specified in the same format as ILLDATE—YYYYMMDD.

Table 3. TREATMNT segment

Field name	Field length
DATE	8
MEDICINE	10

Table 3. TREATMNT segment (continued)

Field name	Field length
QUANTITY	4
DOCTOR	10

- **BILLING Segment**

The following table shows the BILLING segment. It has only one field: the amount of the current bill. BILLING has no key field.

Table 4. BILLING segment

Field name	Field length
BILLING	6

- **PAYMENT Segment**

The following table shows the PAYMENT segment. It has only one field: the amount of payments for the month. The PAYMENT segment has no key field.

Table 5. PAYMENT segment

Field name	Field length
PAYMENT	6

- **HOUSHOLD Segment**

The following table shows the HOUSHOLD segment.

It contains two fields:

- The names of the members of the patient's household (RELNAME)
- How each member of the household is related to the patient (RELATN)

The HOUSHOLD segment's key field is RELNAME.

Table 6. HOUSHOLD segment

Field name	Field length
RELNAME	10
RELATN	8

Bank account hierarchy example

The bank account hierarchy is an example of an application program that is used with main storage databases (MSDBs). In the medical hierarchy example, the database record for a particular patient comprises the PATIENT segment and all of the segments underneath the PATIENT segment. In an MSDB, such as the one in the bank account example, the segment is the whole database record. The database record contains only the fields that the segment contains.

The two types of MSDBs are *related* and *nonrelated*. In related MSDBs, each segment is "owned" by one logical terminal. The "owned" segment can only be updated by the terminal that owns it. In nonrelated MSDBs, the segments are not owned by logical terminals. The following examples of a related MSDB and a nonrelated MSDB illustrate the differences between the two types of databases.

Related MSDBs

Related MSDBs can be fixed or dynamic. In a fixed related MSDB, you can store summary data about a particular teller at a bank. For example, you can have an identification code for the teller's terminal. Then

you can keep a count of that teller's transactions and balance for the day. This type of application requires a segment with three fields:

TELLERID

A two-character code that identifies the teller

TRANCNT

The number of transactions the teller has processed

TELLBAL

The balance for the teller

The following table shows what the segment for this type of application program looks like.

Table 7. Teller segment in a fixed related MSDB		
TELLERID	TRANCNT	TELLBAL

Some of the characteristics of fixed related MSDBs include:

- You can only read and replace segments. You cannot delete or insert segments. In the bank teller example, the teller can change the number of transactions processed, but you cannot add or delete any segments. You never need to add or delete segments.
- Each segment is assigned to one logical terminal. Only the owning terminal can change a segment, but other terminals can read the segment. In the bank teller example, you do not want tellers to update the information about other tellers, but you allow the tellers to view each other's information. Tellers are responsible for their own transactions.
- The name of the logical terminal that owns the segment is the segment's key. Unlike non-MSDB segments, the MSDB key is not a field of the segment. It is used as a means of storing and accessing segments.
- A logical terminal can only own one segment in any one MSDB.

In a dynamic related MSDB, you can store data summarizing the activity of all bank tellers at a single branch. For example, this segment contains:

BRANCHNO

The identification number for the branch

TOTAL

The bank branch's current balance

TRANCNT

The number of transactions for the branch on that day

DEPBAL

The deposit balance, giving the total dollar amount of deposits for the branch

WTHBAL

The withdrawal balance, giving the dollar amount of the withdrawals for the branch

The following table shows what the branch summary segment looks like in a dynamic related MSDB.

Table 8. Branch summary segment in a dynamic related MSDB				
BRANCHNO	TOTAL	TRANCNT	DEPBAL	WTHBAL

How dynamic related MSDBs differ from fixed related MSDBs:

- The owning logical terminal can delete and insert segments in a dynamic related MSDB.
- The MSDB can have a pool of unassigned segments. This kind of segment is assigned to a logical terminal when the logical terminal inserts it, and is returned to the pool when the logical terminal deletes it.

Nonrelated MSDBs

A nonrelated MSDB is used to store data that is updated by several terminals during the same time period. For example, you might store data about an individuals' bank accounts in a nonrelated MSDB segment, so that the information can be updated by a teller at any terminal. Your program might need to access the data in the following segment fields:

ACCNTNO

The account number

BRANCH

The name of the branch where the account is

TRANCNT

The number of transactions for this account this month

BALANCE

The current balance

The following table shows what the account segment in a nonrelated MSDB application program looks like.

Table 9. Account segment in a nonrelated MSDB			
ACCNTNO	BRANCH	TRANCNT	BALANCE

The characteristics of nonrelated MSDBs include:

- Segments are not owned by terminals as they are in related MSDBs. Therefore, IMS programs and Fast Path programs can update these segments. Updating segments is not restricted to the owning logical terminal.
- Your program cannot delete or insert segments.
- Segment keys can be the name of a logical terminal. A nonrelated MSDB exists with terminal-related keys. The segments are not owned by the logical terminals, and the logical terminal name is used to identify the segment.
- If the key is not the name of a logical terminal, it can be any value, and it is in the first field of the segment. Segments are loaded in key sequence.

Your program's view of the data

IMS uses two kinds of control blocks to enable application programs to be independent of your method of storing data in the database, the database description (DBD), and the database program communication block (DB PCB).

Database Description (DBD)

A *database description (DBD)* is physical structure of the database. The DBD also defines the appearance and contents, or fields, that make up each of the segment types in the database.

For example, the DBD for the medical database hierarchy shown in "Medical hierarchy example" describes the physical structure of the hierarchy and each of the six segment types in the hierarchy: PATIENT, ILLNESS, TREATMNT, BILLING, PAYMENT, and HOUSHOLD.

Related Reading: For more information on generating DBDs, see *IMS Version 15.3 Database Utilities*.

Database Program Communication Block (DB PCB)

A *database program communication block (DB PCB)* is a control block that defines an application program's view of the database. An application program often needs to process only some of the segments in a database. A PCB defines which of the segments in the database the program is allowed to access—which segments the program is sensitive to.

The data structures that are available to the program contain only segments that the program is sensitive to. The PCB also defines how the application program is allowed to process the segments in the data structure: whether the program can only read the segments, or whether it can also update them.

To obtain the highest level of data availability, your PCBs should request the fewest number of sensitive segments and the least capability needed to complete the task.

All the DB PCBs for a single application program are contained in a *program specification block* (PSB). A program might use only one DB PCB (if it processes only one data structure) or it might use several DB PCBs, one for each data structure.

Related Reading: For more information on generating PSBs, see *IMS Version 15.3 Database Utilities*.

The following figure illustrates the concept of defining a view for an application program. An accounting program that calculates and prints bills for the clinic's patients would need only the PATIENT, BILLING, and PAYMENT segments. You could define the data structure shown in the following figure in a DB PCB for this program.




Figure 2. Accounting program's view of the database

A program that updates the database with information on patients' illnesses and treatments, in contrast, would need to process the PATIENT, ILLNESS, and TREATMNT segments. You could define the data structure shown in the following figure in a DB PCB for this program.




Figure 3. Patient illness program's view of the database

Sometimes a program needs to process all of the segments in the database. When this is true, the program's view of the database as defined in the DB PCB is the same as the database hierarchy that is defined in the DBD.

An application program processes only the segments in a database that it requires; therefore, if you change the format of a segment that is not processed, you do not change the program. A program is affected only by the segments that it accesses. In addition to being sensitive to only certain segments in a database, a program can also be sensitive to only certain fields within a segment. If you change a segment or field that the program is not sensitive to, it does not affect the program. You define segment and *field-level* sensitivity during PSBGEN.

Definition: *Field-level* sensitivity is when a program is sensitive to only certain fields within a segment.

Related Reading: For more information, see *IMS Version 15.3 Database Administration*.

Processing a database record

To process the information in the database, your application program communicates with IMS in three ways: by passing control, by communicating processing requests, and by exchanging information using DL/I calls.

- Passing control—IMS passes control to your application program through an entry statement in your program. Your program returns control to IMS when it has finished its processing.

When you are running a CICS online program, CICS passes control to your application program, and your program schedules a PSB to make IMS requests. Your program returns control to CICS. If you are running a batch or BMP program, IMS passes control to your program with an existing PSB scheduled.

- Communicating processing requests—You communicate processing requests to IMS in one of two ways:
 - In IMS, you issue DL/I calls to process the database.
 - In CICS, you can issue either DL/I calls or EXEC DLI commands. EXEC DLI commands more closely resemble a higher-level language than do DL/I calls.
- Exchanging information using DL/I calls—Your program exchanges information in two areas:
 - A DL/I call reports the results of your request in a control block and the AIB communication block when using one of the AIB interfaces. For programs written using DL/I calls, this control block is the DB PCB. For programs written using EXEC DLI commands, this control block is the DLI interface block (DIB). The contents of the DIB reflect the status of the last DL/I command executed in the program. Your program includes a mask of the appropriate control block and uses this mask to check the results of the request.

- When you request a segment from the database, IMS returns the segment to your I/O area. When you want to update a segment in the database, you place the new value of the segment in the I/O area.

An application program can read and update a database. When you update a database, you can replace, delete, or add segments. In IMS, you indicate in the DL/I call the segment you want to process, and whether you want to read or update it. In CICS, you can indicate what you want using either a DL/I call or an EXEC DLI command.

Tasks for developing an application

The following tasks are involved in developing an IMS application, and the programs that are part of the application.

Designing the application

Application program design varies from place to place, and from one application to another.

Therefore, this information does not try to cover the early tasks that are part of designing an application program. Instead, it covers only the tasks that you are concerned with after the early specifications for the application have been developed. The tasks for designing the application are:

- **Analyzing Application Data Requirements**

Two important parts of application design are defining the data that each of the business processes in the application requires and designing a local view for each of the business processes.

- **Analyzing Application Processing Requirements**

When you understand the business processes that are part of the application, you can analyze the requirements of each business process in terms of the processing that is available with different types of application programs.

- **Gathering Requirements for Database Options**

You then need to look at the database options that can most efficiently meet the requirements, and gather information about your application's data requirements that relates to each of the options.

- **Gathering Requirements for Message Processing Options**

If your application communicates with terminals and other application programs, look at the message processing options and the requirements they satisfy.

For more information about designing a CICS application, see *CICS Transaction Server for z/OS CICS Application Programming Guide*.

Developing specifications

Developing specifications involves defining what your application will do, and how it will be done. The task of developing specifications is not described in this information because it depends entirely on the specific application and your standards.

Implementing the design

When the specifications for each of the programs in the application are developed, you can structure and code the programs according to those specifications. The tasks of implementing the design are:

- **Writing the Database Processing Part of the Program**

When the program design is complete, you can structure and code your requests and data areas based on the programming specifications that have been developed.

- **Writing the Message Processing Part of the Program**

If you are writing a program that communicates with terminals and other programs, you need to structure and code the message processing part of the program.

- **Analyzing APPC/IMS Requirements**

The LU 6.2 feature of IMS TM enables your application to be distributed throughout the network.

- **Testing an Application Program**

When you finish coding your program, test it by itself and then as part of a system.

- **Documenting an Application Program**

Documenting a program continues throughout the project and is most effective when done incrementally. When the program is completely tested, information must be supplied to those who use and maintain your program.

Chapter 2. Designing an application: Data and local views

Designing an application that meets the requirements of end users involves a variety of tasks and, usually, people from several departments. Application design begins when a department or business area communicates a need for some type of processing. Application design ends when each of the parts of the application system—for example, the programs, the databases, the display screens, and the message formats—have been designed.

An overview of application design

The application design process varies from place to place and from application to application. The overview that is given in this section and the suggestions about documenting application design and converting existing applications are not the only way that these tasks are performed.

The purpose of this overview is to give you a frame of reference so that you can understand where the techniques and guidelines explained in this section fit into the process. The order in which you perform the tasks described here, and the importance you give to each one, depend on your settings. Also, the individuals involved in each task, and their titles, might differ depending on the site. The tasks are as follows:

- Establish your standards

Throughout the design process, be aware of your established standards. Some of the areas that standards are usually established for are:

- Naming conventions (for example, for databases and terminals)
- Formats for screens and messages
- Control of and access to the database
- Programming and conventions (for common routines and macros)

Setting up standards in these areas is usually an ongoing task that is the responsibility of database and system administrators.

- Follow your security standards

Security protects your resources from unauthorized access and use. As with defining standards, designing an adequate security system is often an ongoing task. As an application is modified or expanded, often the security must be changed in some way also. Security is an important consideration in the initial stages of application design.

Establishing security standards and requirements is usually the responsibility of system administration. These standards are based on the requirements of your applications.

Some security concerns are:

- Access to and use of the databases
- Access to terminals
- Distribution of application output
- Control of program modification
- Transaction and command entry

- Define application data

Identifying the data that an application requires is a major part of application design. One of the tasks of data definition is learning from end users what information will be required to perform the required processing.

- Provide input for database design

To design a database that meets the requirements of all the applications that will process it, the database administrator (DBA) needs information about the data requirements of each application. One way to gather and supply this information is to design a local view for each of the business processes in your application. A local view is a description of the data that a particular business process requires.

- Design application programs

When the overall application flow and system externals have been defined, you define the programs that will perform the required processing. Some of the most important considerations involved in this task are: standards, security requirements, privacy requirements, and performance requirements. The specifications you develop for the programs should include:

- Security requirements
- Input and output data formats and volumes
- Data verification and validation requirements
- Logic specifications
- Performance requirements
- Recovery requirements
- Linkage requirements and conventions
- Data availability considerations

In addition, you might be asked to provide some information about your application to the people responsible for network and user interface design.

- Document the application design process

Recording information about the application design process is valuable to others who work with the application now and in the future. One kind of information that is helpful is information about why you designed the application the way you did. This information can be helpful to people who are responsible for the database, your IMS system, and the programs in the application—especially if any part of the application must be changed in the future. Documenting application design is done most thoroughly when it is done during the design process, instead of at the end of it.

- Convert an existing application

One of the main aspects in converting an existing application to IMS is to know what already exists. Before starting to convert the existing system, find out everything you can about the way it works currently. For example, the following information can be of help to you when you begin the conversion:

- Record layouts of all records used by the application
- Number of data element occurrences for each data element
- Structure of any existing related databases

Related concepts

[“Providing data security” on page 84](#)

You can control the security of data accessed by your IMS application programs through data sensitivity and processing options.

[“Identifying online security requirements” on page 91](#)

Security in an online system means protecting the data from unauthorized use through terminals. It also means preventing unauthorized use of both the IMS system and the application programs that access the database. For example, you do not want a program that processes paychecks to be available to everyone who can access the system.

[“Identifying application data” on page 17](#)

Two important aspects of application design are identifying the application data and describing the data that a particular business process requires.

[“Designing a local view” on page 21](#)

A *local view* is a description of the data that an individual business process requires.

Identifying application data

Two important aspects of application design are identifying the application data and describing the data that a particular business process requires.

One of the steps of identifying application data is to thoroughly understand the processing the user wants performed. You need to understand the input data and the required output data in order to define the data requirements of the application. You also need to understand the business processes that are involved in the user's processing needs. Three of the tasks involved in identifying application data are:

- Listing the data required by the business process
- Naming the data
- Documenting the data

When analyzing the required application data, you can categorize the data as either an entity or a data element.

Definitions: An *entity* is anything about which information can be stored. A *data element* is the smallest named unit of data pertaining to an entity. It is information that describes the entity.

Example: In an education application, "students" and "courses" are both entities; these are two subjects about which you collect and process data. The following table shows some data elements that relate to the student and course entities. The entity is listed with its related data elements.

Table 10. Entities and data elements	
Entity	Data elements
Student	Student Name
	Student Number
Course	Course Name
	Course Number
	Course Length

When you store this data in an IMS database, groups of data elements are potential segments in the hierarchy. Each data element is a potential field in that segment.

Related concepts

[“An overview of application design” on page 15](#)

The application design process varies from place to place and from application to application. The overview that is given in this section and the suggestions about documenting application design and converting existing applications are not the only way that these tasks are performed.

Listing data elements

To identify application data, you list its data elements.

For example, to identify application data, consider a company that provides technical education to its customers. The education company has one headquarters office, called Headquarters, and several local education centers, called Ed Centers.

A class is a single offering of a course on a specific date at a particular Ed Center. One course might have several offerings at different Ed Centers; each of these is a separate class. Headquarters is responsible for developing all the courses that will be offered, and each Ed Center is responsible for scheduling classes and enrolling students for its classes.

Suppose that one of the education company's requirements is for each Ed Center to print weekly current rosters for all classes at the Ed Center. The current roster is to give information about the class and the

students enrolled in the class. Headquarters wants the current rosters to be in the format shown in the following figure.

```

CHICAGO                                01/04/04

  TRANSISTOR THEORY                     41837
  10 DAYS
  INSTRUCTOR(S): BENSON, R.J.          DATE: 01/14/04

  STUDENT      CUST  LOCATION          STATUS  ABSENT  GRADE
1. ADAMS, J.W.  XYZ   SOUTH BEND, IND    CONF
2. BAKER, R.T.  ACME  BENTON HARBOR, MICH  WAIT
3. DRAKE, R.A.  XYZ   SOUTH BEND, IND    CANC
.
.
.
33. WILLIAMS, L.R. BEST  CHICAGO, ILL    CONF

CONFIRMED = 30
WAIT-LISTED = 1
CANCELED = 2

```

Figure 4. Current roster for technical education example

To list the data elements for a particular business process, look at the required output. The current roster shown in the previous figure is the roster for the class, "Transistor Theory" to be given in the Chicago Ed Center, starting on January 14, 2004, for ten days. Each course has a course code associated with it—in this case, 41837. The code for a particular course is always the same. For example, if Transistor Theory is also offered in New York, the course code is still 41837. The roster also gives the names of the instructors who are teaching the course. Although the example only shows one instructor, a course might require more than one instructor.

For each student, the roster keeps the following information: a sequence number for each student, the student's name, the student's company (CUST), the company's location, the student's status in the class, and the student's absences and grade. All the above information on the course and the students is input information.

The current date (the date that the roster is printed) is displayed in the upper right corner (01/04/04). The current date is an example of data that is output only data; it is generated by the operating system and is not stored in the database.

The bottom-left corner gives a summary of the class status. This data is not included in the input data. These values are determined by the program during processing.

When you list the data elements, abbreviating them is helpful, because you will be referring to them frequently when you design the local view.

The data elements list for current roster is:

EDCNTR

Name of Ed Center giving class

DATE

Date class starts

CRSNAME

Name of course

CRSCODE

Course code

LENGTH

Length of course

INSTRS

Names of instructors teaching class

STUSEQ#

Student's sequence number

STUNAME

Student's name

CUST

Name of student's company

LOCTN

Location of student's company

STATUS

Student's status in class—confirmed, wait list, or cancelled

ABSENCE

Number of days student was absent

GRADE

Student's grade for the course

After you have listed the data elements, choose the major entity that these elements describe. In this case, the major entity is class. Although a lot of information exists about each student and some information exists about the course in general, together all this information relates to a specific class. If the information about each student (for example, status, absence, and grade) is not related to a particular class, the information is meaningless. This holds true for the data elements at the top of the list as well: The Ed Center, the date the class starts, and the instructor mean nothing unless you know what class they describe.

Naming data elements

Some of the data elements your application uses might already exist and be named. After you have listed the data elements, find out if any of them exist by checking with your database administrator (DBA).

Before you begin naming data elements, be aware of the naming standards that you are subject to. When you name data elements, use the most descriptive names possible. Remember that, because other applications probably use at least some of the same data, the names should mean the same thing to everyone. Try not to limit the name's meaning only to your application.

Recommendation: Use global names rather than local names. A *global name* is a name whose meaning is clear outside of any particular application. A *local name* is a name that, to be understood, must be seen in the context of a particular application.

One of the problems with using local names is that you can develop synonyms, two names for the same data element.

For example, in the current roster example, suppose the student's company was referred to simply as "company" instead of "customer". But suppose the accounting department for the education company used the same piece of data in a billing application—the name of the student's company—and referred to it as "customer". This would mean that two business processes were using two different names for the same piece of data. At worst, this could lead to redundant data if no one realized that "customer" and "company" contained the same data. To solve this, use a global name that is recognized by both departments using this data element. In this case, "customer" is more easily recognized and the better choice. This name uniquely identifies the data element and has a specific meaning within the education company.

When you choose data element names, use qualifiers so that each name can mean only one thing.

For example, suppose Headquarters, for each course that is taught, assigns a number to the course as it is developed and calls this number the "sequence number". The Ed Centers, as they receive student enrollments for a particular class, assign a number to each student as a means of identification within the class. The Ed Centers call **this** number the "sequence number". Thus Headquarters and the Ed Centers are using the same name for two separate data elements. This is called a *homonym*. You can solve the homonym problem by qualifying the names. The number that Headquarters assigns to each course can be called "course code" (CRSCODE), and the number that the Ed Centers assign to their students can be called "student sequence number" (STUSEQ#).

Homonym

One word for two different things.

Choose data element names that identify the element and describe it precisely. Make your data element names:

Unique

The name is clearly distinguishable from other names.

Self-explanatory

The name is easily understood and recognized.

Concise

The name is descriptive in a few words.

Universal

The name means the same thing to everyone.

Documenting application data

After you have determined what data elements a business process requires, record as much information about each of the data elements as possible.

This information is useful to the DBA. Be aware of any standards that you are subject to regarding data documentation. Many places have standards concerning what information should be recorded about data and how and where that information should be recorded. The amount and type of this information varies from place to place. The following list is the type of information that is often recorded.

The descriptive name of the data element

Data element names should be precise, yet they should be meaningful to people who are familiar and also to those who are unfamiliar with the application.

The length of the data element

The length of the data element determines segment size and segment format.

The character format

The programmer needs to know if the data is alphanumeric, hexadecimal, packed decimal, or binary.

The range of possible values for the element

The range of possible values for the element is important for validity checking.

The default value

The programmer also needs the default value.

The number of data element occurrences

The number of data element occurrences helps the DBA to determine the required space for this data, and it affects performance considerations.

How the business process affects the data element

Whether the data element is read or updated determines the processing option that is coded in the PSB for the application program.

You should also record control information about the data. Such information should address the following questions:

- What action should the program take when the data it attempts to access is not available?
- If the format of a particular data element changes, which business processes does that affect? For example, if an education database has as one of its data elements a five-digit code for each course, and the code is changed to six digits, which business processes does this affect?
- Where is the data now? Know the sources of the data elements required by the application.
- Which business processes make changes to a particular data element?
- Are there security requirements about the data in your application? For example, you would not want information such as employees' salaries available to everyone?
- Which department owns and controls the data?

One way to gather and record this information is to use a form similar to the one shown in the following table. The amount and type of data that you record depends on the standards that you are subject to. For example, the following table lists the ID number, data element name, length, the character format, the allowed, null, default values, and the number of occurrences.

Table 11. Example of data elements information form

ID #	Data element name	Length	Char. format	Allowed values	Null values	Default value	Number of occurrences
5	Course Code	5 bytes	Hexa-decimal	0010090000	00000	N/A	There are 200 courses in the curriculum. An average of 10 are new or revised per year. An average of 5 are dropped per year.
25	Status	4 bytes	Alpha-numeric	CONF WAIT CANC	blanks	WAIT	1 per student
36	Student Name	20 bytes	Alpha-numeric	Alpha only	blanks	N/A	There are 3 to 100 students per class with an average of 40 per class.

A *data dictionary* is a good place to record the facts about the application's data. When you are analyzing data, a dictionary can help you find out whether a particular data element already exists, and if it does, its characteristics. With the IBM OS/VS DB/DC Data Dictionary, you can determine online what segments exist in a particular database and what fields those segments contain. You can use either tool to create reports involving the same information.

Designing a local view

A *local view* is a description of the data that an individual business process requires.

It includes the following:

- A list of the data elements
- A conceptual data structure that shows how you have grouped data elements by the entities that they describe
- The relationships between each of the groups of data elements

Definitions: A *data aggregate* is a group of data elements. When you have grouped data elements by the entity they describe, you can determine the relationships between the data aggregates. These relationships are called *mappings*. Based on the mappings, you can design a conceptual data structure for the business process. You should document this process as well.

Related concepts

[“An overview of application design” on page 15](#)

The application design process varies from place to place and from application to application. The overview that is given in this section and the suggestions about documenting application design and converting existing applications are not the only way that these tasks are performed.

Analyzing data relationships

When you analyze data relationships, you are developing conceptual data structures for the business processes in your application.

This process, called *data structuring*, is a way to analyze the relationships among the data elements a business process requires, not a way to design a database. The decisions about segment formats and contents belong to the DBA. The information you develop is input for designing a database.

Data structuring can be done in many different ways.

Grouping data elements into hierarchies

The data elements that describe a data aggregate, the student, might be represented by the descriptive names STUSEQ#, STUNAME, CUST, LOCTN, STATUS, ABSENCE, and GRADE. We call this group of data elements the student data aggregate.

Data elements have values and names. In the student data elements example, the values are a particular student's sequence number, the student's name, company, company location, the student's status in the class, the student's absences, and grade. The names of the data aggregate are not unique—they describe all the students in the class in the same terms. The combined values, however, of a data aggregate occurrence are unique. No two students can have the same values in each of these fields.

As you group data elements into data aggregates and data structures, look at the data elements that make up each group and choose one or more data elements that uniquely identify that group. This is the data aggregate's *controlling key*, which is the data element or group of data elements in the aggregate that uniquely identifies the aggregate. Sometimes you must use more than one data element for the key in order to uniquely identify the aggregate.

By following the three steps explained in this section, you can develop a conceptual data structure for a business process's data. However, you are not developing the logical data structure for the program that performs the business process. The three steps are:

1. Separate repeating data elements in a single occurrence of the data aggregate.
2. Separate duplicate values in multiple occurrences of the data aggregate.
3. Group each data element with its controlling keys.

Step 1. separating repeating data elements

Look at a single occurrence of the data aggregate. The following table shows what this looks like for the class aggregate; the data element is listed with the class aggregate occurrence.

Table 12. Single occurrence of class aggregate	
Data element	Class aggregate occurrence
EDCNTR	CHICAGO
DATE(START)	1/14/96
CRSNAME	TRANSISTOR THEORY
CRS CODE	41837
LENGTH	10 DAYS
INSTRS	multiple
STUSEQ#	multiple

Table 12. Single occurrence of class aggregate (continued)

Data element	Class aggregate occurrence
STUNAME	multiple
CUST	multiple
LOCTN	multiple
STATUS	multiple
ABSENCE	multiple
GRADE	multiple

The data elements defined as multiple are the elements that repeat. Separate the repeating data elements by shifting them to a lower level. Keep data elements with their controlling keys.

The data elements that repeat for a single class are: STUSEQ#, STUNAME, CUST, LOCTN, STATUS, ABSENCE, and GRADE. INSTRS is also a repeating data element, because some classes require two instructors, although this class requires only one.

When you separate repeating data elements into groups, you have the structure shown in the following figure.

In the following figure, the data elements in each box form an aggregate. The entire figure depicts a data structure. The data elements include the Course aggregate, the Student aggregate, and the Instructor aggregate.

The following figure shows these aggregates with the keys indicated with leading asterisks (*).




Figure 5. Current roster after step 1

The keys for the data aggregates are shown in the following table.

Table 13. Data aggregates and keys for current roster after step 1

Data aggregate	Keys
Course aggregate	EDCNTR, DATE, CRSCODE

Table 13. Data aggregates and keys for current roster after step 1 (continued)

Data aggregate	Keys
Student aggregate	EDCNTR, DATE, CRSCODE, STUSEQ#
Instructor aggregate	EDCNTR, DATE, CRSCODE, INSTRS

The asterisks in the previous figure identify the key data elements. For the Class aggregate, it takes multiple data elements to identify the course, so you need multiple data elements to make up the key. The data elements that comprise the Class aggregate are:

- Controlling key element, STUSEQ#
- STUNAME
- CUST
- LOCTN
- STATUS
- ABSENCE
- GRADE

The data elements that comprise the Instructor aggregate are:

- Key element, INSTRS

The Course aggregate and the Instructor aggregate inherit the following keys from the root segment, Course aggregate:

- EDCNTR
- DATE
- CRSCODE

After you have shifted repeating data elements, make sure that each element is in the same group as its controlling key. INSTRS is separated from the group of data elements describing a student because the information about instructors is unrelated to the information about the students. The student sequence number does not control who the instructor is.

In the example shown in the previous figure, the Student aggregate and Instructor aggregate are both dependents of the Course aggregate. A dependent aggregate's key includes the concatenated keys of all the aggregates above the dependent aggregate. This is because a dependent's controlling key does not mean anything if you do not know the keys of the higher aggregates. For example, if you knew that a student's sequence number was 4, you would be able to find out all the information about the student associated with that number. This number would be meaningless, however, if it were not associated with a particular course. But, because the key for the Student aggregate is made up of Ed Center, date, and course code, you can deduce which class the student is in.

Step 2. isolating duplicate aggregate values

Look at multiple occurrences of the aggregate—in this case, the values you might have for two classes. The following table shows multiple occurrences (2) of the same data elements. As you look at this table, check for duplicate values. Remember that both occurrences describe one course.

Table 14. Multiple occurrences of class aggregate

Data element list	Occurrence 1	Occurrence 2
EDCNTR	CHICAGO	NEW YORK
DATE(START)	1/14/96	3/10/96
CRSNAME	TRANS THEORY	TRANS THEORY

Table 14. Multiple occurrences of class aggregate (continued)

Data element list	Occurrence 1	Occurrence 2
CRSCODE	41837	41837
LENGTH	10 DAYS	10 DAYS
INSTRS	multiple	multiple
STUSEQ#	multiple	multiple
STUNAME	multiple	multiple
CUST	multiple	multiple
LOCTN	multiple	multiple
STATUS	multiple	multiple
ABSENCE	multiple	multiple
GRADE	multiple	multiple

The data elements defined as multiple are the data elements that repeat. The values in these elements are not the same. The aggregate is always unique for a particular class.

In this step, compare the two occurrences and shift the fields with duplicate values (TRANS THEORY and so on) to a higher level. If you need to, choose a controlling key for aggregates that do not yet have keys.

In the previous table, CRSNAME, CRSCODE, and LENGTH are the fields that have duplicate values. Much of this process is intuitive. Student status and grade, although they can have duplicate values, should not be separated because they are not meaningful values by themselves. These values would not be used to identify a particular student. This becomes clear when you remember to keep data elements with their controlling keys. When you separate duplicate values, you have the structure shown in the following figure.




Figure 6. Current roster after step 2

Step 3. grouping data elements with their controlling keys

This step is often a check on the first two steps. (Sometimes the first two steps have already done what this step instructs you to do.)

At this stage, make sure that each data element is in the group that contains its controlling key. The data element should depend on the full key. If the data element depends only on part of the key, separate the data element along with the partial (controlling) key on which it depends.

In this example, CUST and LOCTN do not depend on the STUSEQ#. They are related to the student, but they do not depend on the student. They identify the company and company address of the student.

CUST and LOCTN are not dependent on the course, the Ed Center, or the date, either. They are separate from all of these things. Because a student is only associated with one CUST and LOCTN, but a CUST and LOCTN can have many students attending classes, the CUST and LOCTN aggregate should be above the student aggregate.

The following figure shows these aggregates and keys indicated with leading asterisks (*) and shows what the structure looks like when you separate CUST and LOCTN.




Figure 7. Current roster after step 3

The keys for the data aggregates are shown in the following table.

Table 15. Data aggregates and keys for current roster after step 3

Data aggregate	Keys
Course aggregate	CRSCODE
Class aggregate	CRSCODE, EDCNTR, DATE
Customer aggregate	CUST, LOCTN
Student aggregate	(when viewed from the customer aggregate in "Current roster after step 3" instead of from the course aggregate, in "Current roster after step 2") CUST, LOCTN, STUSEQ, CRSCODE, EDCNTR, DATE
Instructor aggregate	CRSCODE, EDCNTR, DATE, INSTRS

Deciding on the arrangement of the customer and location information is part of designing a database. Data structuring should separate any inconsistent data elements from the rest of the data elements.

Determining mappings

When you have arranged the data aggregates into a conceptual data structure, you can examine the relationships between the data aggregates. A mapping between two data aggregates is the quantitative relationship between the two.

The reason you record mappings is that they reflect relationships between segments in the data structure that you have developed. If you store this information in an IMS database, the DBA can construct a database hierarchy that satisfies all the local views, based on the mappings. In determining mappings, it is easier to refer to the data aggregates by their keys, rather than by their collected data elements.

The two possible relationships between any two data aggregates are:

- One-to-many

For each segment A, one or more occurrences of segment B exist. For example, each class maps to one or more students.

Mapping notation shows this in the following way:

Class  **Student**

- Many-to-many

Segment B has many A segments associated with it and segment A has many B segments associated with it. In a hierarchic data structure, a parent can have one or more children, but each child can be associated with only one parent. The many-to-many association does not fit into a hierarchy, because in a many-to-many association each child can be associated with more than one parent.

Related Reading: For more information about analyzing data requirements, see *IMS Version 15.3 Database Administration*.

Many-to-many relationships occur between segments in two business processes. A many-to-many relationship indicates a conflict in the way that two business processes need to process those data aggregates. If you use the IMS full-function database, you can solve this kind of processing conflict by using secondary indexing or logical relationships.

The mappings for the current roster are:

- **Course**  **Class**

For each course, there might be several classes scheduled, but a class is associated with only one course.

- **Class**  **Student**

A class has many students enrolled in it, but a student might be in only one class offering of this course.

- **Class**  **Instructor**

A class might have more than one instructor, but an instructor only teaches one class at a time.

- **Customer/location**  **Student**

A customer might have several students attending a particular class, but each student is only associated with one customer and location.

Related concepts

[“Understanding how data structure conflicts are resolved” on page 76](#)

The order in which application programs need to process fields and segments within hierarchies is frequently not the same for each application. When the DBA finds a conflict in the way that two or more programs need to access the data, three options are available to solve these problems. Each of the following options solves a different kind of conflict.

Local view examples

The following examples show how to design local views including the schedule of courses, the instructor skills report, and the instructor schedules.

Each example shows the following parts of designing a local view:

1. Gather the data. For each example, the data elements are listed and two occurrences of the data aggregate are shown. Two occurrences are shown because you need to look at both occurrences when you look for repeating fields and duplicate values.
2. Analyze the data relationships. First, group the data elements into a conceptual data structure using these three steps:
 - a. Separate repeating data elements in a single occurrence of the data aggregate by shifting them to a lower level. Keep data elements with their keys.

- b. Separate duplicating values in two occurrences of the data aggregate by shifting those data elements to a higher level. Again, keep data elements with their keys.
 - c. Group data elements with their keys. Make sure that all the data elements within one aggregate have the same key. Separate any that do not.
3. Determine the mappings between the data aggregates in the data structure you have developed.

Example 1: schedule of courses

Headquarters keeps a schedule of all the courses given each quarter and distributes it monthly. Headquarters wants the schedule to be sorted by course code and printed in the format shown in the following figure.

COURSE SCHEDULE			
COURSE:	TRANSISTOR THEORY	COURSE CODE:	418737
LENGTH:	10 DAYS	PRICE:	\$280
<u>DATE</u>		<u>LOCATION</u>	
APRIL 14		BOSTON	
APRIL 21		CHICAGO	
.			
.			
NOVEMBER 18		LOS ANGELES	

Figure 8. Schedule of courses

1. Gather the data. The following table lists the data elements and two occurrences of the data aggregate.

Table 16. Course schedule data elements

Data elements	Occurrence 1	Occurrence 2
CRSNAME	TRANS THEORY	MICRO PROG
CRSCODE	41837	41840
LENGTH	10 DAYS	5 DAYS
PRICE	\$280	\$150
DATE	multiple	multiple
EDCNTR	multiple	multiple

2. Analyze the data relationships. First, group the data elements into a conceptual data structure.
 - a. Separate repeating data elements in one occurrence of the data aggregate by shifting them to a lower level, as shown in the following table




Figure 9. Course schedule after step 1

- b. Next, separate duplicate values in two occurrences of the data aggregate by shifting the data elements to a higher level.

This data aggregate does not contain duplicate values.

- c. Group data elements with their controlling keys.

Data elements are grouped with their keys in the present structure. No changes are necessary for this step.

The keys for the data aggregates are shown in the following table.

Table 17. Data aggregates and keys for course schedule after step 1

Data aggregate	Keys
Course aggregate	CRSCODE
Class aggregate	CRSCODE, EDCNTR, DATE

3. When you have developed a conceptual data structure, determine the mappings for the data aggregates.

The mapping for this local view is: **Course** $\longleftrightarrow$ **Class**

Example 2: instructor skills report

Each Ed Center needs to print a report showing the courses that its instructors are qualified to teach. The report format is shown in the following figure.

INSTRUCTOR SKILLS REPORT		
<u>INSTRUCTOR</u>	<u>COURSE_CODE</u>	<u>COURSE_NAME</u>
BENSON, R. J.	41837	TRANS THEORY
MORRIS, S. R.	41837	TRANS THEORY
	41850	CIRCUIT DESIGN
	41852	LOGIC THEORY
.		
.		
REYNOLDS, P. W.	41840	MICRO PROG
	41850	CIRCUIT DESIGN

Figure 10. Instructor skills report

1. Gather the data. The following table lists the data elements and two occurrences of the data aggregate.

Table 18. Instructor skills data elements

Data elements	Occurrence 1	Occurrence 2
INSTR	REYNOLDS, P.W.	MORRIS, S. R.
CRSCODE	multiple	multiple
CRSNAME	multiple	multiple

2. Analyze the data relationships. First, group the data elements into a conceptual data structure.
 - a. Separate repeating data elements in one occurrence of the data aggregate by shifting to a higher level as shown in the following figure.

Instructor aggregate




Figure 11. Instructor skills after step 1

- b. Separate any duplicate values in the two occurrences of the data aggregate.
No duplicate values exist in this data aggregate.
 - c. Group data elements with their keys.
All data elements are grouped with their keys in the current data structure. There are no changes to this data structure.
3. Determine the mappings for the data aggregates.

The mapping for this local view is: **Instructor** $\longleftrightarrow$ **Course**

Example 3: instructor schedules

Headquarters wants to produce a report showing the schedules for all the instructors. The following figure shows the report format.

INSTRUCTOR SCHEDULES				
INSTRUCTOR	COURSE	CODE	ED CENTER	DATE
BENSON, R. J.	TRANS THEORY	41837	CHICAGO	1/14/96
MORRIS, S. R.	TRANS THEORY	41837	NEW YORK	3/10/96
	LOGIC THEORY	41852	BOSTON	3/27/96
	CIRCUIT DES	41840	CHICAGO	4/21/96
REYNOLDS, B. H.	MICRO PROG	41850	NEW YORK	2/25/96
	CIRCUIT DES	41850	LOS ANGELES	3/10/96

Figure 12. Instructor schedules

1. Gather the data. The following table lists the data elements and two occurrences of the data aggregate.

Table 19. Instructor schedules data elements

Data elements	Occurrence 1	Occurrence 2
INSTR	BENSON, R. J.	MORRIS, S. R.
CRSNAME	multiple	multiple
CRSCODE	multiple	multiple
EDCNTR	multiple	multiple
DATE(START)	multiple	multiple

2. Analyze the data relationships. First, group the data elements into a conceptual data structure.
 - a. Separate repeating data elements in one occurrence of the data aggregate by shifting data elements to a lower level as shown in the following figure.

Instructor aggregate




Figure 13. Instructor schedules step 1

- b. Separate duplicate values in two occurrences of the data aggregate by shifting data elements to a higher level as shown in the following figure.

In this example, CRSNAME and CRSCODE can be duplicated for one instructor or for many instructors, for example, 41837 for Benson and 41850 for Morris and Reynolds.

Instructor aggregate




Figure 14. Instructor schedules step 2

- c. Group data elements with their keys.

All data elements are grouped with their controlling keys in the current data structure. No changes to the current data structure are required.

3. Determine the mappings for the data aggregates.

The mappings for this local view are: **Instructor**  **Course Course**  **Class**

An analysis of data requirements is necessary to combine the requirements of the three examples presented in this section and to design a hierarchic structure for the database based on these requirements.

Related Reading: For more information on analyzing data requirements, see *IMS Version 15.3 Database Administration*.

Chapter 3. Analyzing IMS application processing requirements

Use the following information to plan for writing application programs for IMS environments.

Defining IMS application requirements

One of the steps of application design is to decide how the business processes, or tasks, that the end user wants performed can be best grouped into a set of programs that efficiently performs the required processing.

To analyze processing requirements, consider:

- **When the task must be performed**

- Will the task be scheduled unpredictably (for example, on terminal demand) or periodically (for example, weekly)?

- **How the program that performs the task is executed**

- Will the program be executed online, where response time is crucial, or by batch job submission, where a slower response time is acceptable?

- **The consistency of the processing components**

- Does the action the program is to perform involve more than one type of program logic? For example, does it involve mostly retrievals and only one or two updates? If so, you should consider separating the updates into a separate program.
- Does this action involve several large groups of data? If it does, it might be more efficient to separate the programs by the data they access.

- **Any special requirements about the data or processing**

- **Security**

- Should access to the program be restricted?

- **Recovery**

- Are there special recovery considerations in the program's processing?

- **Availability**

- Does your application require high data availability?

- **Integrity**

- Do other departments use the same data?

Answers to questions like these can help you decide on the number of application programs that the processing will require, and on the types of programs that perform the processing most efficiently. Although rules dealing with how many programs can most efficiently do the required processing do not exist, here are some suggestions:

- As you look at each programming task, examine the data and processing that each task involves. If a task requires different types of processing and has different time limitations (for example, daily as opposed to different times throughout the month), that task might be more efficiently performed by several programs.
- As you define each program, it is a good idea for maintenance and recovery reasons to keep it as simple as possible. The simpler a program is—the less it does—the easier it is to maintain, and to restart after a program or system failure. The same is true with data availability—the less data that is accessed, the more likely the data is to be available. The more limited the access requested, the more likely the data is to be available.

Similarly, if the data that the application requires is physically in one place, it might be more efficient to have one program do more of the processing than usual. These are considerations that depend upon the processing and the data of each application.

- Documenting each of the user tasks is helpful during the design process, and in the future when others will work with your application. Be sure you are aware of standards in this area. The kind of information that is typically kept is when the action is to be executed, a functional description, and requirements for maintenance, security, and recovery.

For example, for the current roster process described previously, you might record the information shown in the following form. How frequently the program is run is determined by the number of classes (20) needed by the Education Center each week.

Documenting user task descriptions: current roster example

USER TASK DESCRIPTION

NAME: Current Roster

ENVIRONMENT: Batch **FREQUENCY:** 20 per week

INVOKING EVENT OR DOCUMENT: Time period (one week)

REQUIRED RESPONSE TIME: 24 hours

FUNCTION DESCRIPTION: Print weekly, a current student roster, in student number sequence for each class offered at the Education Center.

MAINTENANCE: Included in Education DB maintenance.

SECURITY: None.

RECOVERY: After a failure, the ability to start printing a particular class roster starting from a particular sequential student number.

Accessing databases with your IMS application program

When designing your program, consider the type of database it must access. The type of database depends on the operating environment.

The program types you can run and the different types of databases you can access in a DB batch, TM batch, DB/DC, DBCTL, or DCCTL environment are shown in the following table.

Table 20. Program and database options in IMS environments

Environment	Type of program you can run	Type of database that can be accessed
DB/DC	BMP	Db2 for z/OS DEDB and MSDB Full function z/OS files
	IFP	Db2 for z/OS DEDB Full function
	JBP	Db2 for z/OS DEDB Full function
	JMP	Db2 for z/OS DEDB Full function
	MPP	Db2 for z/OS DEDB and MSDB Full function
DB Batch	DB Batch	Db2 for z/OS Full function GSAM z/OS files
DBCTL	BMP (Batch-oriented)	Db2 for z/OS DEDB Full function GSAM z/OS files
	JBP	Db2 for z/OS DEDB Full function
DCCTL	BMP	Db2 for z/OS GSAM z/OS files
	IFP	Db2 for z/OS
	JMP	Db2 for z/OS
	MPP	Db2 for z/OS

Table 20. Program and database options in IMS environments (continued)

Environment	Type of program you can run	Type of database that can be accessed
TM Batch	TM Batch	Db2 for z/OS GSAM z/OS files

The types of databases that can be accessed are:

- **IMS Databases**

There are two types of IMS databases: full-function and Fast Path.

- **Full-function databases**

Full-function databases are hierarchic databases that are accessed through Data Language I (DL/I) call interface and can be processed by these types of application programs: IFP, JMP, JBP, MPP, BMP, and DB batch. DL/I calls make it possible for IMS application programs to retrieve, replace, delete, and add segments to full-function databases.

JMP and JBP applications use JDBC to access full-function databases in addition to DL/I.

If you use data sharing, online programs and batch programs can access the same full-function database concurrently.

Full-function database types include: HDAM, HIDAM, HSAM, HISAM, PHDAM, PHIDAM, SHSAM, and SHISAM.

- **Fast Path databases**

Fast Path databases are of two types: MSDBs and DEDBs.

- **Main storage databases** (MSDBs) are root-segment-only databases that reside in virtual storage during execution.
- **Data entry databases** (DEDBs) are hierarchic databases that provide a high level of availability for, and efficient access to, large volumes of detailed data.

MPP, BMP, and IFP programs can access Fast Path databases. In the DBCTL environment, BMP programs can access DEDBs but not MSDBs. JMP and JBP programs can access DEDBs but not MSDBs.

- **Db2 for z/OS databases**

Db2 for z/OS databases are relational databases that can be processed by IMS batch, BMP, IFP, JBP, JMP, and MPP programs. An IMS application program might access only DL/I databases, both DL/I and Db2 for z/OS databases, or only Db2 for z/OS databases. Relational databases are represented to application programs and users as tables, and are processed using a relational data language called Structured Query Language (SQL).

Note: Programs running in 64bit JMP and JBP regions cannot access Db2 for z/OS databases.

Related Reading: For information on processing Db2 for z/OS databases, see DB2® for z/OS Application Programming and SQL Guide.

- **z/OS Files**

BMPs (in DB/DC, DBCTL, and DCCTL environments) are the only type of online application program that can access z/OS files for their input or output. Batch programs can also access z/OS files.

- **GSAM Databases** (Generalized Sequential Access Method)

Generalized Sequential Access Method (GSAM) is an access method that makes it possible for BMPs and batch programs to access a sequential z/OS data set as a simple database. A GSAM database can be accessed by z/OS or by IMS.

Accessing data: the types of programs you can write for your IMS application

You must decide what type of program to use: batch programs, message processing programs (MPPs), IMS Fast Path (IFP) applications, batch message processing (BMP) applications, Java Message Processing (JMP) applications, or Java Batch Processing (JBP) applications. The types of programs you can use depend on whether you are running in the batch, DB/DC, or DBCTL environment.

DB batch processing

These topics describe DB batch processing and can help you decide if this batch program is appropriate for your application.

Data that a DB batch program can access

A DB batch program can access full-function databases, Db2 for z/OS databases, GSAM databases, and z/OS files. A DB batch program cannot access DEDBs or MSDBs.

Using DB batch processing

Batch programs are typically longer-running programs than online programs. You use a batch program when you have a large number of database updates to do or a report to print. Because a batch program runs by itself—it does not compete with any other programs for resources like databases—it can run independently of the control region. If you use data sharing, DB batch programs and online programs can access full-function databases concurrently. Batch programs:

- Typically produce a large amount of output, such as reports.
- Are not executed by another program or user. They are usually scheduled at specific time intervals (for example, weekly) and are started with JCL.
- Produce output that is not needed right away. The turnaround time for batch output is not crucial, as it usually is for online programs.

Recovering a DB batch program

Include checkpoints in your batch program to restart it in case of failure.

Issuing checkpoints

Issue checkpoints in a batch program to commit database changes and provide places from which to restart your program. Issuing checkpoints in a batch program is important, because commit points do not occur automatically, as they do in MPPs, transaction-oriented BMPs, and IFPs.

Issuing checkpoints is particularly important in a batch program that participates in data sharing with your online system. Checkpoints free up resources for use by online programs. You should initially include checkpoints in all batch programs that you write. Even though the checkpoint support might not be needed then, it is easier to incorporate checkpoints initially than to try to fit them in later. And it is possible that you might want to convert your batch program to a BMP or participate in data sharing.

To issue checkpoints (or other system service calls), you must specify an I/O PCB for your program. To obtain an I/O PCB, use the compatibility option by specifying `CMPAT=YES` in the `PSBGEN` statement in your program's PSB.

GSAM DB's are not backed out but are repositioned during the BMP's restart process through the `XRST` call. The `XRST` call repositions the dataset pointers to the checkpoint ID specified in the call. When the application starts-up, it will pick-up from that point and go forward. The checkpoint ID specified in the `XRST` call should be the same one that the non-GSAM DBs would have been backed out to, through either dynamic or batch backout.

Recommendation: For PSBs used by DB batch programs, always specify `CMPAT=YES`.

Backing out database changes

The type of storage medium for the system log determines what happens when a DB batch program terminates abnormally. You can specify that the system log be stored on either DASD (direct access storage device) or tape.

System log on DASD

If the system log is stored on DASD, using the BKO execution parameter you can specify that IMS is to dynamically back out the changes that the program has made to the database since its last commit point.

Related Reading: For information on using the BKO execution parameter, see *IMS Version 15.3 System Definition*.

Dynamically backing out database changes has the following advantages:

- Data accessed by the program that failed is available to other programs immediately. If batch backout is used, other programs cannot access the data until the IMS Batch Backout utility has been run to back out the database changes.
- If data sharing is being used and two programs are deadlocked, one of the programs can continue processing. Otherwise, if batch backout is used, both programs fail.

IMS performs dynamic backout for a batch program when an IMS-detected failure occurs, for example, when a deadlock is detected. Logging to DASD makes it possible for batch programs to issue the SETS, ROLB, and ROLS system service calls. These calls cause IMS to dynamically back out changes that the program has made.

Related Reading: For information on the SETS, ROLB, and ROLS calls, see the information about recovering databases and maintaining database integrity in *IMS Version 15.3 Database Administration*.

System log on tape

If a batch application program terminates abnormally and the batch system log is stored on tape, you must use the IMS Batch Backout utility to back out the program's changes to the database.

Related concepts

[“When to use checkpoint calls” on page 49](#)

Issuing Checkpoint calls is most important in programs that do not have built-in commit points.

TM batch processing

A TM batch program acts like a DB batch program with the following differences.

- It cannot access full-function databases, but it can access Db2 for z/OS databases, GSAM databases, and z/OS files.
- To issue checkpoints for recovery, you need not specify CMPAT=YES in your program's PSB. (The CMPAT parameter is ignored in TM batch.) The I/O PCB is always the first PCB in the list.
- You cannot dynamically back out a database because IMS does not own the databases.

The IEFRDER log DD statement is required in order to enable log synchronization with other external subsystems, such as DB2 for z/OS.

Processing messages: Message Processing Programs

A Message Processing Program (MPP) is an online program that can access full-function databases, DEDBs, MSDBs, and Db2 for z/OS databases. Unlike BMPs and batch programs, MPPs cannot access GSAM databases. MPPs can only run in DB/DC and DCCTL environments.

Using an MPP

The primary purpose of an MPP is to process requests from users at terminals and from other application programs. Ideally, MPPs are very small, and the processing they perform is tailored to respond to requests quickly. They process messages as their input, and send messages as responses.

Message

Data that is transmitted between any two terminals, application programs, or IMS systems. Each message has one or more segments.

MPPs are executed through transaction codes. When you define an MPP, you associate it with one or more transaction codes. Each transaction code represents a transaction the MPP is to process. To process a transaction, a user at a terminal enters a code for that transaction. IMS then schedules the MPP associated with that code, and the MPP processes the transaction. The MPP might need to access the database to do this. Generally, an MPP goes through these five steps to process a transaction:

1. Retrieve a message from IMS.
2. Process the message and access the database as necessary.
3. Respond to the message.
4. Repeat the process until no messages are forthcoming.
5. Terminate.

When an MPP is defined, a system administrator makes decisions about the program's scheduling and processing. For each MPP, a system administrator specifies:

- The transaction's priority
- The number of messages for a particular transaction code that the MPP can process in a single scheduling
- The amount of time (in seconds) in which the MPP is allowed to process a single transaction

Defining priorities and processing limits gives system administration some control over load balancing and processing.

Although the primary purpose of an MPP is to process and reply to messages quickly, it is flexible in how it processes a transaction and where it can send output messages. For example, an MPP can send output messages to other terminals and application programs.

Related concepts

[“Gathering requirements for database options” on page 69](#)

After designing hierarchies for the databases that your application will access, the DBA evaluates database options in terms of which options will best meet application requirements. Whether these options are used depends on the collected requirements of the applications. To design an efficient database, the DBA needs information about the individual applications.

Processing messages: IMS Fast Path Programs

An IMS Fast Path Program (IFP) is similar to an MPP: Its main purpose is to quickly process and reply to messages from terminals. Like an MPP, an IFP can access full-function databases, DEDBs, MSDBs, and Db2 for z/OS databases. IFPs can only be run in DB/DC and DCCTL environments.

Using an IFP

You should use an IFP if you need quick processing and can accept the characteristics and constraints associated with IFPs.

The main differences between IFPs and MPPs are as follows:

- Messages processed by IFPs must consist of only one segment. Messages that are processed by MPPs can consist of several segments.
- IFPs bypass IMS queuing, allowing for more efficient processing. Transactions that are processed by Fast Path's EMH (expedited message handler) are on a first-in, first-out basis.

IFPs also have the following characteristics:

- They run in **transaction response mode**. This means that they must respond to the terminal that sent the message before the terminal can enter any more requests.

- They process only **wait-for-input transactions**. When you define a program as processing wait-for-input transactions, the program remains in virtual storage, even when no additional messages are available for it to process.

Restrictions:

- An IMS program cannot send messages to an IFP transaction unless it is in another IMS system that is connected using Intersystem Communication (ISC).
- MPPs cannot pass conversations to an IFP transaction.

Recovering an IFP

IFPs must be defined as single mode. This means that a commit point occurs each time the program retrieves a message. Because of this, you do not need to issue checkpoint calls.

Batch message processing: BMPs

BMPs are application programs that can perform batch-type processing online and access the IMS message queues for their input and output. Because of this and because of the data available to them, BMPs are the most flexible of the IMS application programs. The two types of BMPs are: batch-oriented and transaction-oriented.

Batch processing online: batch-oriented BMPs

A batch-oriented BMP performs batch-type processing in any online environment. When run in the DB/DC or DCCTL environment, a batch-oriented BMP can send its output to the IMS message queue to be processed later by another application program. Unlike a transaction-oriented BMP, a batch-oriented BMP cannot access the IMS message queue for input.

Data a batch-oriented BMP can access

In the DBCTL environment, a batch-oriented BMP can access full-function databases, Db2 for z/OS databases, DEDBs, z/OS files, and GSAM databases. In the DB/DC environment, a batch-oriented BMP can access all of these types of databases, as well as Fast Path MSDBs. In the DCCTL environment, this program can access Db2 for z/OS databases, z/OS files, and GSAM databases.

Using a batch-oriented BMP

A batch-oriented BMP can be simply a batch program that runs online. (Online requests are processed by the IMS DB/DC, DBCTL, or DCCTL system rather than by a batch system.) You can even run the same program as a BMP or as a batch program.

Recommendation: If the program performs a large number of database updates without issuing checkpoints, consider running it as a batch program so that it does not degrade the performance of the online system.

To use batch-oriented BMPs most efficiently, avoid a large amount of batch-type processing online. If you have a BMP that performs time-consuming processing such as report writing and database scanning, schedule it during non-peak hours of processing. This will prevent it from degrading the response time of MPPs.

Because BMPs can degrade response times, your response time requirements should be the main consideration in deciding the extent to which you will use batch message processing. Therefore, use BMPs accordingly.

Recovering a batch-oriented BMP

Issuing checkpoint calls is an important part of batch-oriented BMP processing, because commit points do not occur automatically, as they do in MPPs, transaction-oriented BMPs, and IFPs. Unlike most batch programs, a BMP shares resources with MPPs. In addition to committing database changes and providing

places from which to restart (as for a batch program), checkpoints release resources that are locked for the program.

If a batch-oriented BMP fails, IMS and Db2 for z/OS back out the database updates the program has made since the last commit point. You then restart the program with JCL. If the BMP processes z/OS files, you must provide your own method of taking checkpoints and restarting.

Converting a batch program to a batch-oriented BMP

If you have IMS TM or are running in the DBCTL environment, you can convert a batch program to a batch-oriented BMP.

- If you have IMS TM, you might want to convert your programs for these reasons:
 - BMPs can send output to the message queues.
 - BMPs can access DEDBs and MSDBs.
 - BMPs simplify program recovery because logging goes to a single system log. If you use DASD for the system log in batch, you can specify that you want dynamic backout for the program. In that case, batch recovery is similar to BMP recovery, except, of course, with batch you need to manage multiple logs.
 - Restart can be done automatically from the last checkpoint without changing the JCL.
- If you are using DBCTL, you might want to convert your programs for these reasons:
 - BMPs can access DEDBs.
 - BMPs simplify program recovery because logging goes to a single system log. If you use DASD for the system log in batch, you can specify that you want dynamic backout for the program. In that case, batch recovery is similar to BMP recovery, except, of course, with batch you need to manage multiple logs.
- If you are running sysplex data sharing and you either have IMS TM or are using DBCTL, you might want to convert your program. This is because using batch-oriented BMPs helps you stay within the sysplex data-sharing limit of 32 connections for each OSAM or VSAM structure.

If you use data sharing, you can run batch programs concurrently with online programs. If you do not use data sharing, converting a batch program to a BMP makes it possible to run the program with BMPs and other online programs.

Also, if you plan to run your batch programs offline, converting them to BMPs enables you to run them with the online system, instead of waiting until the online system is not running. Running a batch program as a BMP can also keep the data more current.

- If you have IMS TM or are using DBCTL, you can have a program that runs as either a batch program or a BMP.

Recommendation: Code your checkpoints in a way that makes them easy to modify. Converting a batch program to a BMP or converting a batch program to use data sharing requires more frequent checkpoints. Also, if a program fails while running in a batch region, you must restart it in a batch region. If a program fails in a BMP region, you must restart it in a BMP region.

The requirements for converting a batch program to a BMP are:

- The program must have an I/O PCB. You can obtain an I/O PCB in batch by specifying the compatibility (CMPAT) option in the program specification block (PSB) for the program.

Related Reading: For more information on the CMPAT option in the PSB, see *IMS Version 15.3 System Utilities*.

- BMPs must issue checkpoint calls more frequently than batch programs.

Related concepts

[“When to use checkpoint calls” on page 49](#)

Issuing Checkpoint calls is most important in programs that do not have built-in commit points.

Batch message processing: transaction-oriented BMPs

Transaction-oriented BMPs can access z/OS files, GSAM databases, Db2 for z/OS databases, full-function databases, DEDBs, and MSDBs.

Data a transaction-oriented BMP can access

Unlike a batch-oriented BMP, a transaction-oriented BMP can access the IMS message queue for input and output, and it can only run in the DB/DC and DCCTL environments.

Using a transaction-oriented BMP

Unlike MPPs, transaction-oriented BMPs are not scheduled by IMS. You schedule them as needed and start them with JCL. For example, an MPP, as it processes each message, might send an output message giving details of the transaction to the message queue. A transaction-oriented BMP could then access the message queue to produce a daily activity report.

Typically, you use a transaction-oriented BMP to simulate direct update online: Instead of updating the database while processing its transactions, an MPP sends its updates to the message queue. A transaction-oriented BMP then performs the updates for the MPP. You can run the BMP as needed, depending on the number of updates. This improves response time for the MPP, and it keeps the data current. This can be more efficient than having the MPP process its transactions if the response time of the MPP is very important. One disadvantage in doing this, however, is that it splits the transaction into two parts which is not necessary.

If you have a BMP perform an update for an MPP, design the BMP so that, if the BMP terminates abnormally, you can reenter the last message as input for the BMP when you restart it. For example, suppose an MPP gathers database updates for three BMPs to process, and one of the BMPs terminates abnormally. You would need to reenter the message that the terminating BMP was processing to one of the other BMPs for reprocessing.

BMPs can process transactions defined as wait-for-input (WFI). This means that IMS allows the BMP to remain in virtual storage after it has processed the available input messages. IMS returns a QC status code, indicating that the program should terminate when one of the following occurs:

- The program reaches its time limit.
- The master terminal operator enters a command to stop processing.
- IMS is terminated with a checkpoint shutdown.

You specify WFI for a transaction on the WFI parameter of the TRANSACT macro during IMS system definition.

A batch message processing region (BMP) scheduled against WFI transactions returns a QC status code (no more messages) only for the following commands: /PSTOP REGION, /DBD, /DBR, or /STA.

Like MPPs, BMPs can send output messages to several destinations, including other application programs.

Recovering a transaction-oriented BMP

Like MPPs, with transaction-oriented BMPs, you can choose where commit points occur in the program. You can specify that a transaction-oriented BMP be single or multiple mode, just as you can with an MPP. If the BMP is single mode, issuing checkpoint calls is not as critical as in a multiple mode BMP. In a single mode BMP, a commit point occurs each time the program retrieves a message.

Related concepts

[“Identifying output message destinations” on page 97](#)

An application program can send messages to another application program or to IMS terminals. To send output messages, the program issues a call and references the I/O PCB or an alternate PCB. The I/O PCB

and alternate PCBs represent logical terminals and other application programs with which the application program communicates.

[“When to use checkpoint calls” on page 49](#)

Issuing Checkpoint calls is most important in programs that do not have built-in commit points.

Java message processing: JMPs

A JMP application program is similar to an MPP application program, except that JMP applications must be written in Java or object-oriented COBOL. Like an MPP application, a JMP application is started when there is a message in the message queue for the JMP application and IMS schedules the message for processing.

JMP applications can access IMS data or Db2 for z/OS data using JDBC. JMP applications run in JMP regions which have JVMs (Java Virtual Machines).

Related concepts

[“Overview of the IMS Java dependent regions” on page 673](#)

The IMS Java dependent regions are two types of IMS dependent regions that provide a Java Virtual Machine (JVM) environment for Java applications: Java message processing (JMP) regions and Java batch processing (JBP) regions.

Java batch processing: JBPs

A JBP application program is similar to a non-message-driven BMP application program, except that JBP applications must be written in Java, object-oriented COBOL, or object-oriented PL/I.

JBP applications can access IMS data or Db2 for z/OS data using JDBC. JBP applications run in JBP regions which have JVMs.

Related concepts

[“Overview of the IMS Java dependent regions” on page 673](#)

The IMS Java dependent regions are two types of IMS dependent regions that provide a Java Virtual Machine (JVM) environment for Java applications: Java message processing (JMP) regions and Java batch processing (JBP) regions.

IMS programming integrity and recovery considerations

IMS provides support for protecting data integrity for application programs.

How IMS protects data integrity: commit points

When an online program accesses the database, it is not necessarily the only program doing so. IMS and Db2 for z/OS make it possible for more than one application program to access the data concurrently without endangering the integrity of the data.

To access data concurrently while protecting data integrity, IMS and Db2 for z/OS prevent other application programs from accessing segments that your program deletes, replaces, or inserts, until your program reaches a *commit point*. A commit point is the place in the program's processing at which it completes a unit of work. When a unit of work is completed, IMS and Db2 for z/OS commit the changes that your program made to the database. Those changes are now permanent and the changed data is now available to other application programs.

What happens at a commit point

When an application program finishes processing one distinct unit of work, IMS and Db2 for z/OS consider that processing to be valid, even if the program later encounters problems. For example, an application program that is retrieving, processing, and responding to a message from a terminal constitutes a *unit of work*. If the program encounters problems while processing the next input message, the processing it has done on the first input message is not affected. These input messages are separate pieces of processing.

A commit point indicates to IMS that a program has finished a unit of work, and that the processing it has done is accurate. At that time:

- IMS releases segments it has locked for the program since the last commit point. Those segments are then available to other application programs.
- IMS and Db2 for z/OS make the program's changes to the database permanent.
- The current position in all databases except GSAM is reset to the start of the database.

If the program terminates abnormally before reaching the commit point:

- IMS and Db2 for z/OS back out all of the changes the program has made to the database since the last commit point. (This does not apply to batch programs that write their log to tape.)
- IMS discards any output messages that the program has produced since the last commit point.

Until the program reaches a commit point, IMS holds the program's output messages so that, if the program terminates abnormally, users at terminals and other application programs do not receive inaccurate information from the abnormally terminating application program.

If the program is processing an input message and terminates abnormally, the input message is **not** discarded if both of the following conditions exist:

1. You are not using the Non-Discardable Messages (NDM) exit routine.
2. IMS terminates the program with one of the following abend codes: U0777, U2478, U2479, U3303. The input message is saved and processed later.

Exception: The input message is discarded if it is not terminated by one of the abend codes previously referenced. When the program is restarted, IMS gives the program the next message.

If the program is processing an input message when it terminates abnormally, and you use the NDM exit routine, the input message might be discarded from the system regardless of the abend. Whether the input message is discarded from the system depends on how you have written the NDM exit routine.

Related Reading: For more information about the NDM exit routine, see *IMS Version 15.3 Exit Routines*.

- IMS notifies the MTO that the program terminated abnormally.
- IMS and Db2 for z/OS release any locks that the program has held on data it has updated since the last commit point. This makes the data available to other application programs and users.

Where commit points occur

A commit point can occur in a program for any of the following reasons:

- The program terminates normally. Except for a program that accesses Fast Path resources, normal program termination is always a commit point. A program that accesses Fast Path resources must reach a commit point before terminating.
- The program issues a checkpoint call. Checkpoint calls are a program's means of explicitly indicating to IMS that it has reached a commit point in its processing.
- If a program processes messages as its input, a commit point might occur when the program retrieves a new message. IMS considers this commit point the start of a new unit of work in the program. Retrieving a new message is not always a commit point. This depends on whether the program has been defined as **single mode** or **multiple mode**.
 - If you specify single mode, a commit point occurs each time the program issues a call to retrieve a new message. Specifying single mode can simplify recovery, because you can restart the program from the most recent call for a new message if the program terminates abnormally. When IMS restarts the program, the program begins by processing the next message.
 - If you specify multiple mode, a commit point occurs when the program issues a checkpoint call or when it terminates normally. At those times, IMS sends the program's output messages to their destinations. Because multiple-mode programs contain fewer commit points than do single mode programs, multiple mode programs might offer slightly better performance than single-mode programs. When a multiple mode program terminates abnormally, IMS can only restart it from a

checkpoint. Instead of reprocessing only the most recent message, a program might have several messages to reprocess, depending on when the program issued the last checkpoint call.

The following table lists the modes in which the programs can run. Because processing mode is not applicable to batch programs and batch-oriented BMPs, they are not listed in the table. The program type is listed, and the table indicates which mode is supported.

Table 21. Processing modes

Program type	Single mode only	Multiple mode only	Either mode
MPP			X
IFP	X		
Transaction-oriented BMP			X

You specify single or multiple mode on the MODE parameter of the TRANSACT macro.

Related Reading: For information on the TRANSACT macro, see *IMS Version 15.3 System Definition*.

See the following figure for an illustration of the difference between single-mode and multiple-mode programs. A single-mode program gets and processes messages, sends output, looks for more messages, and terminates if there are no more. A multiple-mode program gets and processes messages, sends output, but has a checkpoint before looking for more messages and terminating. For a single-mode program, the commit points are when the message is obtained and the program terminates. For multiple-mode, the commit point is at the checkpoint and when the program terminates.




Figure 15. Single mode and multiple mode

Db2 for z/OS does some processing with multiple- and single-mode programs that IMS does not. When a multiple-mode program issues a call to retrieve a new message, Db2 for z/OS performs an authorization check. If the authorization check is successful, Db2 for z/OS closes any SQL cursors that are open. This affects the design of your program.

The Db2 for z/OS SQL COMMIT statement causes Db2 for z/OS to make permanent changes to the database. However, this statement is valid only in TSO application programs. If an IMS application program issues this statement, it receives a negative SQL return code.

Planning for program recovery: checkpoint and restart

Recovery in an IMS application program that accesses Db2 for z/OS data is handled by both IMS and Db2 for z/OS. IMS coordinates the process, and Db2 for z/OS handles recovery of Db2 for z/OS data.

Related concepts

[“Introducing checkpoint calls” on page 48](#)

Checkpoint calls indicate to IMS that the program has reached a commit point. They also establish places in the program from which the program can be restarted. IMS has symbolic checkpoint calls and basic checkpoint calls.

[“When to use checkpoint calls” on page 49](#)

Issuing Checkpoint calls is most important in programs that do not have built-in commit points.

[“Specifying checkpoint frequency” on page 51](#)

You should specify checkpoint frequency in your program so that you can easily modify it when the frequency needs to be adjusted.

Introducing checkpoint calls

Checkpoint calls indicate to IMS that the program has reached a commit point. They also establish places in the program from which the program can be restarted. IMS has symbolic checkpoint calls and basic checkpoint calls.

A program might issue only one type of checkpoint call.

- MPPs and IFPs must use basic checkpoint calls.
- BMP, JMP, and batch programs can use either symbolic checkpoint calls or basic checkpoint calls.

Programs that issue symbolic checkpoint calls can specify as many as seven data areas in the program to be checkpointed. When IMS restarts the program, the Restart call restores these areas to the condition they were in when the program issued the symbolic checkpoint call. Because symbolic checkpoint calls do not support z/OS files, if your program accesses z/OS files, you must supply your own method of establishing checkpoints.

You can use symbolic checkpoint for either Normal Start or Extended Restart (XRST).

For example, typical calls for a Normal start would be as follows:

- XRST (I/O area is blank)
- CHKP (I/O area has checkpoint ID)
- Database Calls (including checkpoints)
- CHKP (final checkpoint)

For example, typical calls for an Extended Restart (XRST) would be as follows:

- XRST (I/O area has checkpoint ID)
- CHKP (I/O area has new checkpoint ID)
- Database Calls (including checkpoints)
- CHKP (final checkpoint)

The restart call, which you must use with symbolic checkpoint calls, provides a way of restarting a program after an abnormal termination. It restores the program's data areas to the way they were when the program issued the symbolic checkpoint call. It also restarts the program from the last checkpoint the program established before terminating abnormally.

All programs can use basic checkpoint calls. Because you cannot use the restart call with the basic checkpoint call, you must provide program restart. Basic checkpoint calls do not support either z/OS or GSAM files. IMS programs cannot use z/OS checkpoint and restart. If you access z/OS files, you must supply your own method of establishing checkpoints and restarting.

In addition to the actions that occur at a commit point, issuing a checkpoint call causes IMS to:

- Inform Db2 for z/OS that the changes your program has made to the database can be made permanent. Db2 for z/OS makes the changes to Db2 for z/OS data permanent, and IMS makes the changes to IMS data permanent.
- Write a log record containing the checkpoint identification given in the call to the system log, but only if the PSB contains a DB PCB. You can print checkpoint log records by using the IMS File Select and Formatting Print program (DFSERA10). With this utility, you can select and print log records based on their type, the data they contain, or their sequential positions in the data set. Checkpoint records are X'18' log records.

Related Reading: For more information about the DFSERA10 program, see *IMS Version 15.3 System Utilities*.

- Send a message containing the checkpoint identification that was given in the call to the system console operator and to the IMS master terminal operator.
- Return the next input message to the program's I/O area, if the program processes input messages. In MPPs and transaction-oriented BMPs, a checkpoint call acts like a call for a new message.

Restriction: Do not specify CHKPT=EOV on any DD statement in order to take an IMS checkpoint because of unpredictable results.

Related concepts

[“Planning for program recovery: checkpoint and restart” on page 48](#)

Recovery in an IMS application program that accesses Db2 for z/OS data is handled by both IMS and Db2 for z/OS. IMS coordinates the process, and Db2 for z/OS handles recovery of Db2 for z/OS data.

When to use checkpoint calls

Issuing Checkpoint calls is most important in programs that do not have built-in commit points.

The decision about whether your program should issue checkpoints, and if so, how often, depends on your program. Generally, these programs should issue checkpoint calls:

- Multiple-mode programs
- Batch-oriented BMPs (which can issue either SYNC or CHKP calls)
- Most batch programs
- Programs that run in a data sharing environment
- JMP applications

You do not need to issue checkpoint calls in:

- Single-mode BMP or MPP programs
- Database load programs
- Programs that access the database in read-only mode, as defined with the PROCOPT=GO option (during a PSBGEN), and are short enough to restart from the beginning
- Programs that have exclusive use of the database

Checkpoints in MPPs and transaction-oriented BMPs

The mode type of the program is specified on the MODE keyword of the TRANSACT macro during IMS system generation. The modes are single and multiple.

• In single-mode programs

In single mode programs (MODE=SNGL was specified on the TRANSACT macro during IMS system definition), a Get Unique to the message queue causes an implicit commit to be performed.

• In multiple-mode programs

In multiple-mode BMPs and MPPs, the only commit points are those that result from the checkpoint calls that the program issues and from normal program termination. If the program terminates abnormally and it has not issued checkpoint calls, IMS backs out the program's database updates

and cancels the messages it created since the beginning of the program. If the program has issued checkpoint calls, IMS backs out the program's changes and cancels the output messages it has created since the most recent checkpoint.

Consider the following when issuing checkpoint calls in multiple-mode programs:

- How long it would take to back out and recover that unit of processing. The program should issue checkpoints frequently enough to make the program easy to back out and recover.
- How you want the output messages grouped. checkpoint calls establish how a multiple-mode program's output messages are grouped. Programs should issue checkpoint calls frequently enough to avoid building up too many output messages.

Depending on the database organization, issuing a checkpoint call might reset your position in the database.

Related Reading: For more information about losing your position when a checkpoint is issued, see *IMS Version 15.3 Database Administration*.

Checkpoints in batch-oriented BMPs

Issuing checkpoint calls in a batch-oriented BMP is important for several reasons:

- In addition to committing changes to the database and establishing places from which the program can be restarted, checkpoint calls release resources that IMS has locked for the program.
- A batch-oriented BMP that uses DEDBs or MSDBs might terminate with abend U1008 if a SYNC or CHKP call is not issued before the application program terminates.
- If a batch-oriented BMP does not issue checkpoints frequently enough, it can be abnormally terminated, or it can cause another application program to be abnormally terminated by IMS for any of these reasons:
 - If a BMP retrieves and updates many database records between checkpoint calls, it can tie up large portions of the databases and cause long waits for other programs needing those segments.

Exception: For a BMP with a processing option of GO or exclusive, IMS does not lock segments for programs. Issuing checkpoint calls releases the segments that the BMP has locked and makes them available to other programs.

- The space needed to maintain lock information about the segments that the program has read and updated exceeds what has been defined for the IMS system. If a BMP locks too many segments, the amount of storage needed for the locked segments can exceed the amount of available storage. If this happens, IMS terminates the program abnormally. You must increase the program's checkpoint frequency before rerunning the program. The available storage is specified during IMS system definition.

Related Reading: For more information on specifying storage, see *IMS Version 15.3 System Definition*.

You can limit the number of locks for the BMP by using the LOCKMAX=*n* parameter on the PSBGEN statement. For example, a specification of LOCKMAX=5 means the application cannot obtain more than 5000 locks at any time. The value of *n* must be between 0 and 255. When a maximum lock limit does not exist, 0 is the default. If the BMP tries to acquire more than the specified number of locks, IMS terminates the application with abend U3301.

Related Reading: For more information about this abend, see *IMS Version 15.3 Messages and Codes, Volume 3: IMS Abend Codes*.

Checkpoints in batch programs

Batch programs that update databases should issue checkpoint calls. The main consideration in deciding how often to take checkpoints in a batch program is the time required to back out and reprocess the program after a failure. A general recommendation is to issue one checkpoint call every 10 or 15 minutes.

If you might need to back out the entire batch program, the program should issue the checkpoint call at the beginning of the program. IMS backs out the program to the checkpoint you specify, or to the most recent checkpoint, if you do not specify a checkpoint. If the database is updated after the beginning of the program and before the first checkpoint, IMS is not able to back out these database updates.

For a batch program to issue checkpoint calls, it must specify the compatibility option in its PSB (CMPAT=YES). This generates an I/O PCB for the program, which IMS uses as an I/O PCB in the checkpoint call.

Another important reason for issuing checkpoint calls in batch programs is that, although they may currently run in an IMS batch region, they might later need to access online databases. This would require converting them to BMPs. Issuing checkpoint calls in a BMP is important for reasons other than recovery—for example, to release database resources for other programs. So, you should initially include checkpoints in all batch programs that you write. Although the checkpoint support might not be needed then, it is easier to incorporate checkpoint calls initially than to try to fit them in later.

To free database resources for other programs, batch programs that run in a data-sharing environment should issue checkpoint calls more frequently than those that do not run in a data-sharing environment.

Related concepts

[“DB batch processing” on page 39](#)

These topics describe DB batch processing and can help you decide if this batch program is appropriate for your application.

[“Batch processing online: batch-oriented BMPs” on page 42](#)

A batch-oriented BMP performs batch-type processing in any online environment. When run in the DB/DC or DCCTL environment, a batch-oriented BMP can send its output to the IMS message queue to be processed later by another application program. Unlike a transaction-oriented BMP, a batch-oriented BMP cannot access the IMS message queue for input.

[“Batch message processing: transaction-oriented BMPs” on page 44](#)

Transaction-oriented BMPs can access z/OS files, GSAM databases, Db2 for z/OS databases, full-function databases, DEDBs, and MSDBs.

[“Planning for program recovery: checkpoint and restart” on page 48](#)

Recovery in an IMS application program that accesses Db2 for z/OS data is handled by both IMS and Db2 for z/OS. IMS coordinates the process, and Db2 for z/OS handles recovery of Db2 for z/OS data.

Specifying checkpoint frequency

You should specify checkpoint frequency in your program so that you can easily modify it when the frequency needs to be adjusted.

You can do this by:

- Using a counter in your program to keep track of elapsed time, and issuing a checkpoint call after a certain time interval.
- Using a counter to keep track of the number of root segments your program accesses, and issuing a checkpoint call after a certain number of root segments.
- Using a counter to keep track of the number of updates your program performs, and issuing a checkpoint call after a certain number of updates.

Related concepts

[“Planning for program recovery: checkpoint and restart” on page 48](#)

Recovery in an IMS application program that accesses Db2 for z/OS data is handled by both IMS and Db2 for z/OS. IMS coordinates the process, and Db2 for z/OS handles recovery of Db2 for z/OS data.

Data availability considerations

The following information describes the conditions that could cause data to become unavailable in a full-function database and the program calls that allow your program to manage data under these conditions.

Dealing with unavailable data

The conditions that make the database unavailable for both read and update are:

- The /LOCK command for a database was issued.
- The /STOP command for a database was issued.
- The /DBRECOVERY command was issued.
- Authorization for a database failed.

The conditions that make the database available only for read and not for update are:

- The /DBDUMP command has been issued.
- Database ACCESS value is RD (read).

In addition to unavailability of an entire database, other situations involving unavailability of a limited amount of data can also inhibit program access. One such example would be a failure situation involving data sharing. The active IMS system knows which locks were held by a sharing IMS system at the time the sharing IMS system failed. Although the active IMS system continues to use the database, it must reject access to the data which the failed IMS system locked upon failure. This situation occurs for both full-function and DEDB databases.

The two situations where the program might encounter unavailable data are:

- The program makes a call requiring access to a database that was unavailable at the time the program was scheduled.
- The database was available when the program was scheduled, but limited amounts of data are unavailable. The current call has attempted to access the unavailable data.

Regardless of the condition causing the data to be unavailable, the program has two possible approaches when dealing with unavailable data. The program can be **insensitive** or **sensitive** to data unavailability.

- When the program is insensitive, IMS takes appropriate action when the program attempts to access unavailable data.
- When the program is sensitive, IMS informs the program that the data it is attempting to access is not available.

If the program is insensitive to data unavailability, and attempts to access unavailable data, IMS aborts the program (3303 pseudo-abend), and backs out any updates the program has made. The input message that the program was processing is suspended, and the program is scheduled to process the input message when the data becomes available. However, if the database is unavailable because dynamic allocation failed, a call results in an AI (unable to open) status code.

If the program is sensitive to data unavailability and attempts to access unavailable data, IMS returns a status code indicating that it could not process the call. The program then takes the appropriate action. A facility exists for the program to initiate the same action that IMS would have taken if the program had been insensitive to unavailable data.

IMS does not schedule batch programs if the data that the program can access is unavailable. If the batch program is using block-level data sharing, it might encounter unavailable data if the sharing system fails and the batch system attempts to access data that was updated but not committed by the failed system.

The following conditions alone do not cause a batch program to fail during initialization:

- A PCB refers to a HALDB.

- The use of DBRC is suppressed.

However, without DBRC, a database call using a PCB for a HALDB is not allowed. If the program is sensitive to unavailable data, such a call results in the status code BA; otherwise, such a call results in message DFS3303I, followed by ABENDU3303.

Scheduling and accessing unavailable databases

By using the INIT, INQY, SETS, SETU, and ROLS calls, the program can manage a data environment where the program is scheduled with unavailable databases.

The INIT call informs IMS that the program is sensitive to unavailable data and can accept the status codes that are issued when the program attempts to access such data. The INIT call can also be used to determine the data availability for each PCB.

The INQY call is operable in both batch and online IMS environments. IMS application programs can use the INQY call to request information regarding output destination, session status, the current execution environment, the availability of databases, and the PCB address based on the PCBNAME. The INQY call is only supported by way of the AIB interface (AIBTDLI or CEETDLI using the AIB rather than the PCB address).

The SETS, SETU, and ROLS calls enable the application to define multiple points at which to preserve the state of full-function (except HSAM) databases and message activity. The application can then return to these points at a later time. By issuing a SETS or SETU call before initiating a set of DL/I calls to perform a function, the program can later issue the ROLS call if it cannot complete a function due to data unavailability.

The ROLS call allows the program to roll back its IMS full-function database activity to the state that it was in prior to a SETS or SETU call being issued. If the PSB contains an MSDB or a DEDB, the SETS and ROLS (with token) calls are invalid. Use the SETU call instead of the SETS call if the PSB contains a DEDB, MSDB, or GSAM PCB.

The ROLS call can also be used to undo all update activity (database and messages) since the last commit point and to place the current input message on the suspend queue for later processing. This action is initiated by issuing the ROLS call without a token or I/O area.

Restriction: With Db2 for z/OS, you cannot use ROLS (with a token) or SETS.

Related information

[3303 \(Messages and Codes\)](#)

Use of STAE or ESTAE and SPIE in IMS programs

IMS uses STAE or ESTAE routines in the control region, the dependent (MPP, IFP, BMP) regions, and the batch regions. In the control region, STAE or ESTAE routines ensure that database logging and various resource cleanup functions are complete.

In the dependent region, STAE or ESTAE routines are used to notify the control region of any abnormal termination of the application program or the dependent region itself. If the control region is not notified of the dependent region termination, resources are not properly released and normal checkpoint shutdown might be prevented.

In the batch region, STAE or ESTAE routines ensure that database logging and various resource cleanup functions are complete. If the batch region is not notified of the application program termination, resources might not be properly released.

Two important aspects of the STAE or ESTAE facility are that:

- IMS relies on its STAE or ESTAE facility to ensure database integrity and resource control.
- The STAE or ESTAE facility is also available to the application program.

Because of these two factors, be sure you clearly understand the relationship between the program and the STAE or ESTAE facility.

Generally, do not use the STAE or ESTAE facility in your application program. However, if you believe that the STAE or ESTAE facility is required, you must observe the following basic rules:

- When the environment supports STAE or ESTAE processing, the application program STAE or ESTAE routines always get control before the IMS STAE or ESTAE routines. Therefore, you must ensure that the IMS STAE or ESTAE exit routines receive control by observing the following procedures in your application program:
 - Establish the STAE or ESTAE routine only once and always before the first DL/I call.
 - When using the STAE or ESTAE facility, the application program should not alter the IMS abend code.
 - Do not use the RETRY option when exiting from the STAE or ESTAE routine. Instead, return a CONTINUE-WITH-TERMINATION indicator at the end of the STAE or ESTAE processing. If your application program specifies the RETRY option, be aware that IMS STAE or ESTAE exit routines will not get control to perform cleanup. Therefore, system and database integrity might be compromised.
- The application program STAE or ESTAE exit routine must not issue DL/I calls (DB or TM) because the original abend might have been caused by a problem between the application and IMS. A problem between the application and IMS could result in recursive entry to STAE or ESTAE with potential loss of database integrity, or in problems taking a checkpoint. This also could result in a hang condition or an ABENDU0069 during termination.

Related concepts

[“What to do when your IMS program terminates abnormally” on page 151](#)

When your program terminates abnormally, you can take the following actions to simplify the task of finding and fixing the problem.

Dynamic allocation for IMS databases

Use the dynamic allocation function to specify the JCL information for IMS databases in a library instead of in the JCL of each batch or online job.

If you use dynamic allocation, do not include JCL DD statements for any database data sets that have been defined for dynamic allocation. Check with the DBA or comparable specialist to determine which databases have been defined for dynamic allocation.

Related Reading: For additional information on the definitions for dynamic allocation, see the description of the DFSMDA macro in *IMS Version 15.3 System Definition*.

Chapter 4. Analyzing CICS application processing requirements

IMS supports application programs running in a CICS environment

Defining CICS application requirements

One of the steps of application design is to decide how the business processes, or tasks can be best grouped into a set of programs that will efficiently perform the required processing.

Some of the considerations in analyzing processing requirements are:

- **When the task must be performed**

- Will it be scheduled unpredictably (for example on terminal demand) or periodically (for example, weekly)?

- **How the program that performs the task is executed**

- Will it be executed online, where response time is more important, or by batch job submission, where a slower response time is acceptable?

- **The consistency of the processing components**

- Does this action the program is to perform involve more than one type of program logic? For example, does it involve mostly retrievals, and only one or two updates? If so, you should consider separating the updates into a separate program.
- Does this action involve several large groups of data? If it does, it might be more efficient to separate the programs by the data they access.

- **Any special requirements about the data or processing**

- Security**

- Should access to the program be restricted?

- Recovery**

- Are there special recovery considerations in the program's processing?

- Integrity**

- Do other departments use the same data?

Answers to questions like these can help you decide on the number of application programs that the processing will require, and on the types of programs that perform the processing most efficiently. Although rules dealing with how many programs can most efficiently do the required processing do not exist, here are some suggestions:

- As you look at each programming task, examine the data and processing that each task involves. If a task requires different types of processing and has different time limitations (for example, weekly as opposed to monthly), that task may be more efficiently performed by several programs.
- As you define each program, it is a good idea for maintenance and recovery reasons to keep programs as simple as possible. The simpler a program is—the less it does—the easier it is to maintain, and to restart after a program or system failure. The same is true with data availability—the less data that is accessed, the more likely the data is to be available; the more limited the data accessed, the more likely the data is to be available.

Similarly, if the data that the application requires is physically in one place, it might be more efficient to have one program do more of the processing than usual. These are considerations that depend on the processing and the data of each application.

- Documenting each of the user tasks is helpful during the design process, and in the future when others will work with your application. Be sure you are aware of the standards in this area. The kind

of information that is typically kept is when the task is to be executed, a functional description, and requirements for maintenance, security, and recovery.

For example, for the Current Roster process described previously, you might record the information shown in the following form. How frequently the program is run is determined by the number of classes (20) for which the Ed Center will print current rosters each week.

Example: Current roster task description

USER TASK DESCRIPTION	
NAME: <u>Current Roster</u>	
ENVIRONMENT: <u>Batch</u>	FREQUENCY: <u>20 per week</u>
INVOKING EVENT OR DOCUMENT: <u>Time period (one week)</u>	
REQUIRED RESPONSE TIME: <u>24 hours</u>	
FUNCTION DESCRIPTION: <u>Print weekly, a current student roster, in student number sequence for each class offered at the Education Center.</u>	
MAINTENANCE: <u>Included in Education DB maintenance.</u>	
SECURITY: <u>None.</u>	
RECOVERY: <u>After a failure, the ability to start printing a particular class roster starting from a particular sequential student number.</u>	

Accessing databases with your CICS application program

When designing your program, consider the type of data it must access. The type of data depends on the operating environment.

The data from IMS and Db2 for z/OS databases, and z/OS files, that is available to CICS online and IMS batch programs is shown in the following table.

Table 22. The data that your CICS program can access

Type of program	IMS databases	Db2 for z/OS databases	z/OS files
CICS online	Yes ¹	Yes ²	Yes ³
DB batch	Yes	Yes ³	Yes

Notes:

1. Except for Generalized Sequential Access Method (GSAM) databases. GSAM enables batch programs to access a sequential z/OS data set as a simple database.
2. IMS does not participate in the call process.
3. Access through CICS file control or transient data services.

Also, consider the type of database your program must access. As shown in the following table, the type of program you can write and database that can be accessed depends on the operating environment.

Table 23. Program and database options in the CICS environments

Environment ¹	Type of program you can write	Type of database that can be accessed
DB batch	DB batch	Db2 for z/OS ² DL/I Full-function GSAM z/OS files
DBCTL	BMP	Db2 for z/OS DEDBs Full-function GSAM z/OS files
	CICS online	Db2 for z/OS ² DEDBs Full-function z/OS files (access through CICS file control or transient data services)

Notes:

1. A CICS environment, or CICS remote DL/I environment also exists and is also referred to as function shipping. In this environment, a CICS system supports applications that issue DL/I calls but the CICS system does not service the requests itself. The CICS environment "function ships" the DL/I calls to another CICS system that is using DBCTL. For more information on remote DL/I, see *CICS Transaction Server for z/OS IMS Database Control Guide*.
2. IMS does not participate in the call process.

The types of databases that can be accessed are:

• **Full-Function Databases**

Full-function databases are hierarchic databases that are accessed through Data Language I (DL/I). DL/I calls enable application programs to retrieve, replace, delete, and add segments to full-function databases. CICS online and BMP programs can access the same database concurrently (if participating in IMS data sharing); an IMS batch program must have exclusive access to the database (if not participating in IMS data sharing).

All types of programs (batch, BMPs, and online) can access full-function databases.

• **Fast Path DEDBs**

Data entry databases (DEDBs) are hierarchic databases for, and efficient access to, large volumes of detailed data. In the DBCTL environment, CICS online and BMP programs can access DEDBs.

• **Db2 for z/OS Databases**

Db2 for z/OS databases are relational databases. Relational databases are represented to application programs and users as tables and are processed using a relational data language called Structured Query Language (SQL). Db2 for z/OS databases can be processed by CICS online transactions, and by IMS batch and BMP programs.

Related Reading: For information on processing Db2 for z/OS databases, see *DB2 for z/OS Application Programming and SQL Guide*.

• **GSAM Databases**

Generalized Sequential Access Method (GSAM) is an access method that enables BMPs and batch programs to access a "flat" sequential z/OS data set as a simple database. A GSAM database can be accessed by z/OS or CICS.

- **z/OS Files**

CICS online and IMS batch programs can access z/OS files for their input, processing, or output. Batch programs can access z/OS files directly; online programs must access them through CICS file control or transient data services.

Related concepts

[“Using data sharing for your CICS program” on page 59](#)

If you use data sharing, your programs can participate in IMS data sharing. Under data sharing, CICS online and BMP programs can access the same DL/I database concurrently.

Writing a CICS program to access IMS databases

The types of programs you can use depend on whether you are running in the DBCTL environment. Within the different environments, the type of program you write depends on the processing your application requires. Each type of program answers different application requirements.

Related concepts

[“IMS solutions for Java development overview” on page 553](#)

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

Writing a CICS online program

Use the following information to decide if an online program is appropriate for your application.

Data that a CICS online program can access

CICS online programs run in the DBCTL environment and can access IMS full-function databases, Fast Path DEDBs, Db2 for z/OS databases, and z/OS files.

Online programs that access IMS databases are executed in the same way as other CICS programs.

Using a CICS online program

An online program runs under the control of CICS, and it accesses resources concurrently with other online programs. Some of the application requirements online programs can answer are:

- Information in the database must be available to many users.
- Program needs to communicate with terminals and other programs.
- Programs must be available to users at remote terminals.
- Response time is important.

The structure of an online program, and the way it receives status information, depend on whether it is a call- or command-level program. However, both command- and call-level online programs:

- Schedule a PSB (for CICS online programs). A PSB is automatically scheduled for batch or BMP programs.
- Issue either commands or calls to access the database. Online programs cannot mix commands and calls in one logical unit of work (LUW).
- Optionally, terminate a PSB for CICS online programs.
- Issue an EXEC CICS RETURN statement when they have finished their processing. This statement returns control to the linking program. When the highest-level program issues the RETURN statement, CICS regains control and terminates the PSB if it has not yet been terminated.

Because an online application program can be used concurrently by several tasks, it must be quasi-reentrant.

An online program in the DBCTL environment can use many IMS system service requests.

DL/I database or system service requests must refer to one of the program communication blocks (PCBs) from the list of PCBs passed to your program by IMS. The PCB that must be used for making system service requests is called the I/O PCB. When present, it is the first PCB in the list of PCBs.

For an online program in the DBCTL environment, the I/O PCB is optional. To use the I/O PCB, you must indicate this in the application program when it schedules the PSB.

Before you run your program, the program specification blocks (PSBs) and database descriptions (DBDs) the program uses must be converted to internal control block format using the IMS ACBGEN utility. PSBs specify the characteristics of an application program. DBDs specify the physical and logical characteristics of IMS databases.

Related Reading: For more information on performing an ACBGEN and a PSBGEN, see *IMS Version 15.3 System Utilities*.

Because an online program shares a database with other online programs, it may affect the performance of your online system.

Related concepts

[“Maximizing the performance of your CICS system” on page 61](#)

When you write programs that share data with other programs (for example, a program that will participate in IMS data sharing or a BMP), be aware of how your program affects the performance of the online system.

[“Distributed and local connectivity with the IMS Universal drivers” on page 560](#)

The IMS Universal drivers support distributed (*type-4*) and local (*type-2*) connectivity to IMS databases.

Using data sharing for your CICS program

If you use data sharing, your programs can participate in IMS data sharing. Under data sharing, CICS online and BMP programs can access the same DL/I database concurrently.

Batch programs in a data-sharing environment can access databases used by other batch programs, and by CICS and IMS online programs. With data sharing, you can share data directly and your program's requests need not go through a mirror transaction.

Related Reading: For more information on sharing a database with an IMS system, see *IMS Version 15.3 System Administration*.

Related concepts

[“Accessing databases with your CICS application program” on page 56](#)

When designing your program, consider the type of data it must access. The type of data depends on the operating environment.

Scheduling and terminating a PSB (CICS online programs only)

Before your online program issues any DL/I calls, it must indicate to IMS its intent to use a particular PSB by issuing either a PCB call or a SCHD command. In addition to indicating which PSB your program will use, the PCB call obtains the address of the PCBs in the PSB. When you no longer need a PSB, you can terminate it using the TERM request.

In a CICS online program, you use a PCB call or SCHD command (for command-level programs) to obtain the PSB for your program. Because CICS releases the PSB your program uses when the transaction ends, your program need not explicitly terminate the PSB. Only use a terminate request if you want to:

- Use a different PSB
- Commit all the database updates and establish a logical unit of work for backing out updates
- Free IMS resources for use by other CICS tasks

A terminate request causes a CICS sync point, and a CICS sync point terminates the PSB. For more information about CICS recovery concepts, see the appropriate CICS publication.

Do not use terminate requests for other reasons because:

- A terminate request forces a CICS sync point. This sync point releases all recoverable resources and IMS database resources that were enqueued for this task.

If the program continues to update other CICS resources after the terminate request and then terminates abnormally, only those resources that were updated after the terminate request are backed out. Any IMS changes made by the program are not backed out.

- IMS lock management detects deadlocks that occur if two transactions are waiting for segments held by the other.

When a deadlock is detected, one transaction is abnormally terminated. Database changes are backed out to the last TERM request. If a TERM request or CICS sync point was issued prior to the deadlock, CICS does not restart the transaction.

Related Reading: For a complete description of transaction restart considerations, see *CICS Transaction Server for z/OS Recovery and Restart Guide*.

- Issuing a terminate request causes additional logging.
- If the terminal output requests are issued after a terminate request and the transaction fails at this point, the terminal operator does not receive the message.

The terminal operator may assume that the entire transaction failed, and reenter the input, thus repeating the updates that were made before the terminate request. These updates were not backed out.

Linking and passing control to other programs (CICS online programs only)

Use CICS to link your program to other programs without losing access to the facilities acquired in the linking program.

For example:

- You could schedule a PSB and then link to another program using a LINK command. On return from that program, the PSB is still scheduled.
- Similarly, you could pass control to another program using the XCTL command, and the PSB remains scheduled until that program issues an EXEC CICS RETURN statement. However, when you pass control to another program using XCTL, the working storage of the program passing control is lost. If you want to retain the working storage for use by the program being linked to, you must pass the information in the COMMAREA.

Recommendation: To simplify your work, instead of linking to another program, you can issue all DL/I requests from one program module. This helps to keep the programming simple and easy to maintain.

Terminating a PSB or issuing a sync point affects the linking program. For example, a terminate request or sync point that is issued in the program that was linked causes the release of CICS resources enqueued in the linking program.

How CICS distributed transactions access IMS

CICS can divide a single, logical unit of work into separate CICS transactions and coordinate the sync point globally. If such CICS transactions access DBCTL, locking and buffer management issues might occur.

To IMS, the transactions are separate units of work, on different DBCTL threads, and they do not share locks or buffers. For example, if a global transaction runs, obtains a database lock, and reaches the commit point, CICS does not process the synchronization point until the other transactions in the CICS unit of recovery (UOR) are ready to commit. If a second transaction in the same CICS UOR requests the same lock as that held by the first transaction, the second transaction is held in a lock wait state. The first transaction cannot complete the sync point and release the lock until the second transaction also reaches

the commit point, but this cannot happen because the second transaction is in a lock wait state. You must ensure that this type of collision does not occur with CICS distributed transactions that access IMS.

Maximizing the performance of your CICS system

When you write programs that share data with other programs (for example, a program that will participate in IMS data sharing or a BMP), be aware of how your program affects the performance of the online system.

A BMP program, in particular, can affect the performance of the CICS online transactions. This is because BMP programs usually make a larger number of database updates than CICS online transactions, and a BMP program is more likely to hold segments that CICS online programs need. Limit the number of segments held by a BMP program, so CICS online programs need not wait to acquire them.

One way to limit the number of segments held by a BMP or batch program that participates in IMS data sharing is to issue checkpoint requests in your program to commit database changes and release segments held by the program. When deciding how often to issue checkpoint requests, you can use one or more of the following techniques:

- Divide the program into small logical units of work, and issue a checkpoint call at the end of each unit.
- Issue a checkpoint call after a certain number of DL/I requests have been issued, or after a certain number of transactions are processed.

In CICS online programs, release segments for use by other transactions to maximize the performance of your online system. (Ordinarily, database changes are committed and segments are released only when control is returned to CICS.) To more quickly free resources for use by other transactions, you can issue a TERM request to terminate the PSB. However, less processing overhead generally occurs if the PSB is terminated when control is returned to CICS.

Related concepts

[“Writing a CICS online program” on page 58](#)

Use the following information to decide if an online program is appropriate for your application.

[“Taking checkpoints in batch and BMP programs” on page 62](#)

You can take checkpoints in batch and BMP programs. Checkpoints are important for recovery and for integrity.

Programming integrity and database recovery considerations for your CICS program

IMS provides support for protecting data integrity for CICS online programs

How IMS protects data integrity for CICS online programs

IMS can protect the data integrity for CICS online programs.

IMS protects the integrity of the database for programs that share data by:

- Preventing other application programs with update capability from accessing any segments in the database record your program is processing, until your program finishes with that record and moves to a new database record in the same database.
- Preventing other application programs from accessing segments that your program deletes, replaces, or inserts, until your program reaches a sync point. When your program reaches a sync point, the changes your program has made to the database become permanent, and the changed data becomes available to other application programs.

Exception: If PROCOPT=GO has been defined during PSBGEN for your program, your program can access segments that have been updated but not committed by another program.

- Backing out database updates made by an application program that terminates abnormally.

You may also want to protect the data your program accesses by retaining segments for the sole use of your program until your program reaches a sync point—even if you do not update the segments. (Ordinarily, if you do not update the segments, IMS releases them when your program moves to a new database record.) You can use the Q command code to reserve segments for the exclusive use of your program. You should use this option only when necessary because it makes data unavailable to other programs and can have an impact on performance.

Recovering databases accessed by batch and BMP programs

You can plan for recovering databases accessed by batch or BMP programs.

CICS recovers databases accessed by CICS online programs in the same way it handles other recoverable CICS resources. For example, if an IMS transaction terminates abnormally, CICS and IMS back out all database updates to the last sync point.

For batch or BMP programs, do the following:

- Take checkpoints in your program to commit database changes and provide places from which your program can be restarted.
- Provide the code for or issue a request to restart your program.

You may also want to back out the database changes that have been made by a batch program that has not yet committed these changes.

To perform these tasks, you use system service calls, described in more detail in the appropriate application programming information for your environment.

Requesting an I/O PCB in batch programs

For your program to successfully issue any system service request, an I/O PCB must have been previously requested.

Related concepts

[“Developing JBP applications with the IMS Java dependent region resource adapter” on page 683](#)
JBP applications are similar to JMP applications, except that JBP applications do not receive input messages from the IMS message queue. Unlike batch message processing (BMP) applications, JBP applications must be non-message-driven applications.

Taking checkpoints in batch and BMP programs

You can take checkpoints in batch and BMP programs. Checkpoints are important for recovery and for integrity.

Taking checkpoints in batch and BMP programs is important for two reasons:

- Recovery: Checkpoints establish places in your program from which your program could be restarted, in the event of a program or system failure. If your program abnormally terminates after issuing a checkpoint request, database changes will be backed out to the point at which the checkpoint request was issued.
- Integrity: Checkpoints also commit the changes that your program has made to the database.

In addition to providing places from which to restart your program and committing database changes, issuing checkpoint calls in a BMP program or in a program participating in IMS data sharing releases database segments for use by other programs.

When a batch or BMP program issues a checkpoint request, IMS writes a record containing a checkpoint ID to the IMS system log.

When your application program reaches a point during its execution where you want to make sure that all changes made to that point have been physically entered in the database, issue a checkpoint request. If some condition causes your program to fail before its execution is complete, the database must be restored to its original state. The changes made to the database must be backed out so that the database is not left in a partially updated condition for access by other application programs.

If your program runs a long time, you can reduce the number of changes that must be backed out by taking checkpoints in your program. Then, if your program terminates abnormally, only the database updates that occurred after the checkpoint must be backed out. You can also restart the program from the point at which you issued the checkpoint request, instead of having to restart it from the beginning.

Issuing a checkpoint call cancels your position in the database.

Issue a checkpoint call just before issuing a Get Unique call, which reestablishes your position in the database record after the checkpoint is taken.

Types of checkpoints

The two types of checkpoint calls are basic and symbolic. Both types commit your program's changes to the database and establish places from which your program can be restarted:

Batch and BMP programs can issue basic checkpoint calls using the CHKP call. When you use basic checkpoint calls, you must provide the code for restarting the program after an abnormal termination.

Batch and BMP programs can also issue symbolic checkpoint calls. You can issue a symbolic checkpoint call by using the CHKP call. Like the basic checkpoint call, the symbolic checkpoint call commits changes to the database and establishes places from which the program can be restarted. In addition, the symbolic checkpoint call:

- Works with the Extended Restart call to simplify program restart and recovery.
- Lets you specify as many as seven data areas in the program to be checkpointed. When you restart the program, the restart call restores these areas to the way they were when the program terminated abnormally.

Specifying a checkpoint ID

Each checkpoint call your program issues must have an identification, or ID. Checkpoint IDs must be 8 bytes in length and contain printable EBCDIC characters.

When you want to restart your program, you can supply the ID of the checkpoint from which you want the program to be started. This ID is important because when your program is restarted, IMS searches for checkpoint information with an ID matching the one you have supplied. The first matching ID that IMS finds becomes the restart point for your program. This means that checkpoint IDs must be unique both within each application program and among application programs. If checkpoint IDs are not unique, you cannot be sure that IMS will restart your program from the checkpoint you specified.

One way to make sure that checkpoint IDs are unique within and among programs is to construct IDs in the following order:

- Three bytes of information that uniquely identifies your program.
- Five bytes of information that serves as the ID within the program, for example, a value that is increased by 1 for each checkpoint command or call, or a portion of the system time obtained at program start by issuing the TIME macro.

Specifying checkpoint frequency

To determine the frequency of checkpoint requests, you must consider the type of program and its performance characteristics.

In batch programs

When deciding how often to issue checkpoint requests in a batch program, you should consider the time required to back out and reprocess the program after a failure. For example, if you anticipate that the processing your program performs will take a long time to back out, you should establish checkpoints more frequently.

If you might back out of the entire program, issue the checkpoint request at the very beginning of the program. IMS backs out the database updates to the checkpoint you specify. If the database is updated

after the beginning of the program and before the first checkpoint, IMS is not able to back out these database updates.

In a data-sharing environment, also consider the impact of sharing resources with other programs on your online system. You should issue checkpoint calls more frequently in a batch program that shares data with online programs, to minimize resource contention.

It is a good idea to design all batch programs with checkpoint and restart in mind. Although the checkpoint support may not be needed initially, it is easier to incorporate checkpoint calls initially than to try to fit them in later. If the checkpoint calls are incorporated, it is easier to convert batch programs to BMP programs or to batch programs that use data sharing.

In BMP programs

When deciding how often to issue checkpoint requests in a BMP program, consider the performance of your CICS online system. Because these programs share resources with CICS online transactions, issue checkpoint requests to release segments so CICS online programs need not wait to acquire them.

Printing checkpoint log records

You can print checkpoint log records by using the IMS File Select and Formatting Print Program (DFSERA10). With this utility, you can select and print log records based on their type, the data they contain, or their sequential positions in the data set. Checkpoint records are type 18 log records. *IMS Version 15.3 System Utilities* describes this program.

Related concepts

[“Maximizing the performance of your CICS system” on page 61](#)

When you write programs that share data with other programs (for example, a program that will participate in IMS data sharing or a BMP), be aware of how your program affects the performance of the online system.

Backing out database changes

If your program terminates abnormally, the database must be restored to its previous state and uncommitted changes must be backed out. Changes made by a BMP or CICS online program are automatically backed out. Database changes made by a batch program might or might not be backed out, depending on whether your system log is on DASD.

For a batch program

What happens when a batch program terminates abnormally and how you recover the database depend on the storage medium for the system log. You can specify that the system log is to be stored on either DASD or on tape.

• When the system log is on DASD

You can specify that IMS is to dynamically back out the changes that a batch program has made to the database since its last commit point by coding BKO=Y in the JCL. IMS performs dynamic backout for a batch program when an IMS-detected failure occurs, such as when a deadlock is detected (for batch programs that share data).

DASD logging also makes it possible for batch programs to issue the rollback (ROLB) system service request, in addition to ROLL. The ROLB request causes IMS to dynamically back out the changes the program has made to the database since its last commit point, and then to return control to the application program.

Dynamically backing out database changes has the following advantages:

- Data accessed by the program that failed is immediately available to other programs. Otherwise, if batch backout is not used, data is not available to other programs until the IMS Batch Backout utility has been run to back out the database changes.

- If two programs are deadlocked, one of the programs can continue processing. Otherwise, if batch backout is not used, both programs will fail. (This applies only to batch programs that share data.)

Instead of using dynamic backout, you can run the IMS Batch Backout utility to back out changes.

- **When the system log is on tape**

If a batch application program terminates abnormally and the system log is stored on tape, you must use the IMS Batch Backout utility to back out the program's changes to the database.

Related Reading: For more information, see *IMS Version 15.3 Database Utilities*.

For BMP programs

If your program terminates abnormally, the changes the program has made since the last commit point are backed out. If a system failure occurs, or if the CICS control region or DBCTL terminates abnormally, DBCTL emergency restart backs out all changes made by the program since the last commit point. You need not use the IMS Batch Backout utility because DBCTL backs out the changes. If you need to back out all changes, you can use the ROLL system service call to dynamically back out database changes.

Restarting your program

If you issue symbolic checkpoint calls (for batch and BMP programs), you can use the Extended Restart system service request (XRST) to restart your program after an abnormal termination.

The XRST call restores the program's data areas to the way they were when the program terminated abnormally, and it restarts the program from the last checkpoint request the program issued before terminating abnormally.

If you use basic checkpoint calls (for batch and BMP programs), you must provide the necessary code to restart the program from the latest checkpoint in the event that it terminates abnormally.

One way to restart the program from the latest checkpoint is to store repositioning data in an HDAM database. Your program writes a database record containing repositioning information to the HDAM database. It updates this record at intervals. When the program terminates, the database record is deleted. At the completion of the XRST call, the I/O area always contains a checkpoint ID used by the restart. Normally, XRST will return the 8-byte symbolic checkpoint ID, followed by 4 blanks. If the 8-byte ID consists of all blanks, then XRST will return the 14-byte time-stamp ID. Also, check the status code in the PCB. The only successful status code for an XRST call is a row of blanks.

Related concepts

“Developing JBP applications with the IMS Java dependent region resource adapter” on page 683
JBP applications are similar to JMP applications, except that JBP applications do not receive input messages from the IMS message queue. Unlike batch message processing (BMP) applications, JBP applications must be non-message-driven applications.

Data availability considerations for your CICS program

The data that a program needs to access may sometimes be unavailable. Use the following functions when data is not available.

Unavailability of a database

The conditions that make an entire database unavailable for both read and update are the following.

- A STOP command has been issued for the database.
- A DBRECOVERY (DBR) command has been issued for the database.
- DBRC authorization for the database has failed.

The conditions that make a database available for read but not for update are:

- A DBDUMP command has been issued for the database.

- The database access value is RD (read).

In a data-sharing environment, the command or error that created any of these conditions may have originated on the other system which is sharing data.

Whether a program is scheduled or whether an executing program can schedule a PSB when the database is unavailable depends on the type of program and the environment:

- A batch program

IMS does not schedule a batch program when one of the databases that the program can access is not available.

In a non-data sharing environment, DBRC authorization for a database may fail because the database is currently authorized to a DB/DC environment. In a data-sharing environment, a CICS or a DBCTL master terminal global command to recover a database or to dump a database may make the database unavailable to a batch program.

The following conditions alone do not cause a batch program to fail during initialization:

- A PCB refers to a HALDB.
- The use of DBRC is suppressed.

However, without DBRC, a database call using a PCB for a HALDB is not allowed. If the program is sensitive to unavailable data, such a call results in the status code BA; otherwise, such a call results in message DFS3303I, followed by ABENDU3303.

- An online or BMP program in the DBCTL environment.

When a program executing in this environment attempts to schedule with a PSB containing one or more full-function databases that are unavailable, the scheduling is allowed. If the program does not attempt to access the unavailable database, it can function normally. If it does attempt to access the database, the result is the same as when the database is available but some of the data in it is not available.

Unavailability of some data in a database

In addition to the situation where the entire database is unavailable, there are other situations where a limited amount of data is unavailable. One example is a failure situation involving data sharing where the IMS system knows which locks were held by a sharing IMS at the time the sharing IMS system failed. This IMS system continues to use the database but rejects access to the data that the failed IMS system held locked at the time of failure.

A batch program, an online program, or a BMP program can be operating in the DBCTL environment. If so, the online or BMP programs may have been scheduled when an entire database was not available. The following options apply to these programs when they attempt to access data and either the entire database is unavailable or only some of the data in the database is unavailable.

Programs executing in these environments have an option of being sensitive or insensitive to data unavailability.

- When the program is insensitive to data unavailability and attempts to access unavailable data, the program fails with a 3303 abend. For online programs, this is a pseudo-abend. For batch programs, it is a real abend. However, if the database is unavailable because dynamic allocation failed, a call results in an AI (unable to open) status code.
- When the program is sensitive to data unavailability and attempts to access unavailable data, IMS returns a status code indicating that it could not process the request. The program can then take the appropriate action. A facility exists for the program to then initiate the same action that IMS would have taken if the program had been insensitive to unavailable data.

The program issues the INIT call or ACCEPT STATUS GROUP A command to inform IMS that it is sensitive to unavailable data and can accept the status codes issued when the program attempts to access such data. The INIT request can also be used to determine data availability for each PCB in the PSB.

The SETS or SETU and ROLS functions

The SETS or SETU and ROLS requests allow an application to define multiple points at which to preserve the state of full-function databases.

The application can then return to these points at a later time. By issuing a SETS or SETU request before initiating a set of DL/I requests to perform a function, the program can later issue the ROLS request if it cannot complete the function due possibly to data unavailability.

ROLS allows the program to roll back its IMS activity to the state prior to the SETS or SETU call.

Restriction: SETS or SETU and ROLS only roll back the IMS updates. They do not roll back the updates made using CICS file control or transient data.

Additionally, you can use the ROLS call or command to undo all database update activity since the last checkpoint.

Use of STAE or ESTAE and SPIE in IMS batch programs

IMS uses STAE or ESTAE routines in the IMS batch regions to ensure that database logging and various resource cleanup functions are completed.

Two important aspects of the STAE or ESTAE facility are that:

- IMS relies on its STAE or ESTAE facility to ensure database integrity and resource control.
- The STAE or ESTAE facility is also available to the application program.

Because of these two factors, be sure you clearly understand the relationship between the program and the STAE or ESTAE facility.

Generally, do not use the STAE or ESTAE facility in your batch application program. However, if you believe that the STAE or ESTAE facility is required, you must observe the following basic rules:

- When the environment supports STAE or ESTAE processing, the application program STAE or ESTAE routines always get control before the IMS STAE or ESTAE routines. Therefore, you must ensure that the IMS STAE or ESTAE exit routines receive control by observing the following procedures in your application program:
 - Establish the STAE or ESTAE routine only once and always before the first DL/I call.
 - When using the STAE or ESTAE facility, the application program must not alter the IMS abend code.
 - Do not use the RETRY option when exiting from the STAE or ESTAE routine. Instead, return a CONTINUE-WITH-TERMINATION indicator at the end of the STAE or ESTAE processing. If your application program does specify the RETRY option, be aware that IMS STAE or ESTAE exit routines will not get control to perform cleanup. Therefore, system and database integrity may be compromised.
- The application program STAE/ESTAE exit routine must not issue DL/I calls because the original abend may have been caused by a problem between the application and IMS. This would result in recursive entry to STAE/ESTAE with potential loss of database integrity or in problems taking a checkpoint.

Dynamic allocation for IMS databases

Use the dynamic allocation function to specify the JCL information for IMS databases in a library instead of in the JCL of each batch job or in the JCL for DBCTL.

If you use dynamic allocation, do not include JCL DD statements for any database data sets that have been defined for dynamic allocation. Check with the database administrator (DBA) or comparable specialist at to determine which databases have been defined for dynamic allocation.

Related Reading: For more information on the definitions for dynamic allocation, see the DFSMDA macro in *IMS Version 15.3 System Definition*.

Chapter 5. Gathering requirements for database options

After designing hierarchies for the databases that your application will access, the DBA evaluates database options in terms of which options will best meet application requirements. Whether these options are used depends on the collected requirements of the applications. To design an efficient database, the DBA needs information about the individual applications.

Related concepts

[“Processing messages: Message Processing Programs” on page 40](#)

A Message Processing Program (MPP) is an online program that can access full-function databases, DEDBs, MSDBs, and Db2 for z/OS databases. Unlike BMPs and batch programs, MPPs cannot access GSAM databases. MPPs can only run in DB/DC and DCCTL environments.

Analyzing data access

The DBA chooses a type of database, based on how the majority of programs that use the database will access the data.

IMS databases are categorized according to the access method used. The types of databases that can be defined are as follows:

- Hierarchical Direct Access Method (HDAM)
- Partitioned Hierarchical Direct Access Method (PHDAM)
- Hierarchical Indexed Direct Access Method (HIDAM)
- Partitioned Hierarchical Indexed Direct Access Method (PHIDAM)
- Main Storage Database (MSDB)
- Data Entry Database (DEDB)
- Hierarchical Sequential Access Method (HSAM)
- Hierarchical Indexed Sequential Access Method (HISAM)
- Generalized Sequential Access Method (GSAM)
- Simple Hierarchical Sequential Access Method (SHSAM)
- Simple Hierarchical Indexed Sequential Access Method (SHISAM)

Important: PHDAM and PHIDAM are the partitioned versions of the HDAM and HIDAM database types, respectively. The corresponding descriptions of the HDAM and HIDAM database types therefore apply to PHDAM and PHIDAM.

Some of the information that you can gather to help the DBA to decide the type of database to set up are as follows:

- To access a database record, a program must first access the root of the record. How will each program access root segments?

Is it:

- Directly
- Sequentially
- Both

- The segments within the database record are the dependents of the root segment. How will each program access the segments within each database record?

Is it:

- Directly

- Sequentially
- Both

It is important to note the distinction between accessing a database record and accessing segments within the record. A program might access database records sequentially, but after the program is within a record, the program might access the segments directly. These are different, and can influence the choice of the access method to use.

- To what extent will the program update the database?

Is it:

- By adding new database records?
- By adding new segments to existing database records?
- By deleting segments or database records?

Again, note the difference between updating a database record and updating a segment within the database record.

Direct access

The advantage of direct access processing is that you can get good results for both direct and sequential processing. Direct access means that by using a randomizing routine or an index, IMS can find any database record that you want, regardless of the sequence of database records in the database.

IMS full function has four direct access methods.

- HDAM and PHDAM process data directly by using a randomizing routine to store and locate root segments.
- HIDAM and PHIDAM use an index to help them provide direct processing of root segments.

The direct access methods use pointers to maintain the hierarchic relationships between segments of a database record. By following pointers, IMS can access a path of segments without passing through all the segments in the preceding paths.

Some of the requirements that direct access satisfies are:

- Fast direct processing of roots using an index or a randomizing routine
- Sequential processing of database records with HIDAM and PHIDAM using the index
- Fast access to a path of segments using pointers

In addition, when you delete data from a direct-access database, the new space is available almost immediately. This gives you efficient space utilization; therefore, reorganization of the database is often unnecessary. Direct access methods internally maintain their own pointers and addresses.

A disadvantage of direct access is that you have a larger IMS overhead because of the pointers. But if direct access fulfills your data access requirements, it is more efficient than using a sequential access method.

Primarily direct processing: HDAM

HDAM is efficient for a database that is usually accessed directly but sometimes sequentially. HDAM uses a randomizing routine to locate its root segments and then chains dependent segments together according to the pointer options chosen. The z/OS access methods that HDAM can use are Virtual Storage Access Method (VSAM) and Overflow Storage Access Method (OSAM).

Important: PHDAM is the partitioned version of the HDAM database type. The corresponding descriptions of the HDAM database type therefore apply to PHDAM.

The requirements that HDAM satisfies are:

- Direct access of roots by root keys because HDAM uses a randomizing routine to locate root segments
- Direct access of paths of dependents

- Adding new database records and new segments because the new data goes into the nearest available space
- Deleting database records and segments because the space created by a deletion can be used by any new segment

HDAM characteristics

An HDAM database:

- Can store root segments anywhere. Root segments do not need to be in sequence because the randomizing routine locates them.
- Uses a randomizing routine to locate the relative block number and root anchor point (RAP) within the block that points to the root segment.
- Accesses the RAPs from which the roots are chained in physical sequence. Then the root segments that are chained from the root anchors are returned. Therefore, sequential retrieval of root segments from HDAM is not based on the results of the randomizing routine and is not in key sequence unless the randomizing routine put them into key sequence.
- May not give the desired result for some calls unless the randomizing module causes the physical sequence of root segments to be in the key sequence. For example, a GU call for a root segment that is qualified as less than or equal to a root key value would scan in physical sequence for the first RAP of the first block. This may result in a not-found condition, even though segments meeting the qualification do exist.

For dependent segments, an HDAM database:

- Can store them anywhere
- Chains all segments of one database record together with pointers

An Overview of how HDAM works

This topic contains Diagnosis, Modification, and Tuning information.

When a database record is stored in an HDAM database, HDAM keeps one or more RAPs at the beginning of each physical block. The RAP points to a root segment. HDAM also keeps a pointer at the beginning of each physical block that points to any free space in the block. When you insert a segment, HDAM uses this pointer to locate free space in the physical block. To locate a root segment in an HDAM database, you give HDAM the root key. The randomizing routine gives it the relative physical block number and the RAP that points to the root segment. The specified RAP number gives HDAM the location of the root within a physical block.

Although HDAM can place roots and dependents anywhere in the database, it is better to choose HDAM options that keep roots and dependents close together.

HDAM performance depends largely on the randomizing routine you use. Performance can be very good, but it also depends on other factors such as:

- The block size you use
- The number of RAPs per block
- The pattern for chaining together different segments. You can chain segments of a database record in two ways:
 - In hierarchic sequence, starting with the root
 - In parent-to-dependent sequence, with parents having pointers to each of their paths of dependents

To use HDAM for sequential access of database records by root key, you need to use a secondary index or a randomizing routine that stores roots in physical key sequence.

Direct and sequential processing: HIDAM

HIDAM is the access method that is most efficient for an approximately equal amount of direct and sequential processing.

Important: PHIDAM is the partitioned version of the HIDAM database type. The corresponding descriptions of the HIDAM database type therefore apply to PHIDAM.

The z/OS access methods it can use are VSAM and OSAM. The specific requirements that HIDAM satisfies are:

- Direct and sequential access of records by their root keys
- Direct access of paths of dependents
- Adding new database records and new segments because the new data goes into the nearest available space
- Deleting database records and segments because the space created by a deletion can be used by any new segment

HIDAM can satisfy most processing requirements that involve an even mixture of direct and sequential processing. However, HIDAM is not very efficient with sequential access of dependents.

HIDAM characteristics

For root segments, a HIDAM database:

- Initially loads them in key sequence
- Can store new root segments wherever space is available
- Uses an index to locate a root that you request and identify by supplying the root's key value

For dependent segments, a HIDAM database:

- Can store segments anywhere, preferably fairly close together
- Chains all segments of a database record together with pointers

An overview of how HIDAM works

This topic contains Diagnosis, Modification, and Tuning information.

HIDAM uses two databases. The primary database holds the data. An index database contains entries for all of the root segments in order by their key fields. For each key entry, the index database contains the address of that root segment in the primary database.

When you access a root, you supply the key to the root. HIDAM looks up the key in the index to find the address of the root and then goes to the primary database to find the root.

HIDAM chains dependent segments together so that when you access a dependent segment, HIDAM uses the pointer in one segment to locate the next segment in the hierarchy.

When you process database records directly, HIDAM locates the root through the index and then locates the segments from the root. HIDAM locates dependents through pointers.

If you plan to process database records sequentially, you can specify special pointers in the DBD for the database so that IMS does not need to go to the index to locate the next root segment. These pointers chain the roots together. If you do not chain roots together, HIDAM always goes to the index to locate a root segment. When you process database records sequentially, HIDAM accesses roots in key sequence in the index. This only applies to sequential processing; if you want to access a root segment directly, HIDAM uses the index, and not pointers in other root segments, to find the root segment you have requested.

Main storage database: MSDB

Use MSDBs to store the most frequently-accessed data. MSDBs are suitable for applications such as general ledger applications in the banking industry.

Recommendation: Use DEDBs instead of MSDBs when you develop new Fast Path databases. Terminal-related MSDBs and non-terminal-related MSDBs with terminal-related keys are no longer supported. Although non-terminal-related MSDBs with non-terminal-related-keys are still supported, you should consider converting any existing MSDBs to DEDBs. You can use the MSDB-to-DEDB Conversion utility.

MSDB characteristics

MSDBs reside in virtual storage, enabling application programs to avoid the I/O activity that is required to access them. The two kinds of MSDBs are terminal-related and non-terminal-related.

In a terminal-related MSDB, each segment is owned by one terminal, and each terminal owns only one segment. One use for this type of MSDB is an application in which each segment contains data associated with a logical terminal. In this type of application, the program can read the data (perhaps for reporting purposes), but cannot update it. A non-terminal-related MSDB stores data that is needed by many users during the same time period. It can be updated and read from all terminals (for example, a real time inventory control application, where reduction of inventory can be noted from many cash registers).

An overview of how MSDBs work

This topic contains Diagnosis, Modification, and Tuning information.

MSDB segments are stored as root segments only. Only one type of pointer, the forward chain pointer, is used. This pointer connects the segment records in the database.

Data entry database: DEDB

DEDBs are designed to provide access to and efficient storage for large volumes of data. The primary requirement a DEDB satisfies is a high level of data availability.

DEDB characteristics

DEDBs are hierarchic databases that can have as many as 15 hierarchic levels, and as many as 127 segment types. They can contain both direct and sequential dependent segments. Because the sequential dependent segments are stored in chronological order as they are committed to the database, they are useful in journaling applications.

DEDBs support a subset of functions and options that are available for a HIDAM or HDAM database. For example, a DEDB does not support logically related segments or access with primary indexes. Access with secondary indexes is supported.

An overview of how DEDBs work

This topic contains Diagnosis, Modification, and Tuning information.

A DEDB can be partitioned into multiple areas, with each area containing a different collection of database records. The data in a DEDB area is stored in a VSAM data set. Root segments are stored in the root-addressable part of an area, with direct dependents stored close to the roots for fast access. Direct dependents that cannot be stored close to their roots are stored in the independent overflow portion of the area. Sequential dependents are stored in the sequential dependent portion at the end of the area so that they can be quickly inserted. Each area data set can have up to seven copies, making the data easily available to application programs.

Sequential access

When you use a sequential access method, the segments in the database are stored in hierarchic sequence, one after another, with no pointers.

IMS full-function has two sequential access methods. Like the direct access methods, one has an index and the other does not:

- HSAM only processes root segments and dependent segments sequentially.
- HISAM processes data sequentially but has an index so that you can access records directly. HISAM is primarily for sequentially processing dependents, and directly processing database records.

Some of the general requirements that sequential access satisfies are:

- Fast sequential processing
- Direct processing of database records with HISAM
- Small IMS overhead on storage because sequential access methods relate segments by adjacency rather than with pointers

The three disadvantages of using sequential access methods are:

- Sequential access methods give slower access to the right-most segments in the hierarchy, because HSAM and HISAM must read through all other segments to get to them.
- HISAM requires frequent reorganization to reclaim space from deleted segments and to keep the logical records of a database record physically adjoined.
- You cannot update HSAM databases. You must create a new database to change any of the data.

Sequential processing only: HSAM

HSAM is a hierarchic access method that can handle only sequential processing. You can retrieve data from HSAM databases, but you cannot update any of the data. The z/OS access methods that HSAM can use are QSAM and BSAM.

HSAM is ideal for the following situations:

- You are using the database to collect (but not update) data or statistics.
- You only plan to process the data sequentially.

HSAM characteristics

HSAM stores database records in the sequence in which you submit them. You can only process records and dependent segments sequentially, which means the order in which you have loaded them. HSAM stores dependent segments in hierarchic sequence.

An overview of how HSAM works

This topic contains Diagnosis, Modification, and Tuning information.

HSAM databases are very simple databases. The data is stored in hierarchic sequence, one segment after the other, and no pointers or indexes are used.

Primarily sequential processing: HISAM

HISAM is an access method that stores segments in hierarchic sequence with an index to locate root segments. It also has an overflow data set. Store segments in a logical record until you reach the end of the logical record. When you run out of space on the logical record, but you still have more segments belonging to the database record, you store the remaining segments in an overflow data set. The access methods that HISAM can use are VSAM and OSAM.

HISAM is well-suited for:

- Direct access of record by root keys
- Sequential access of records
- Sequential access of dependent segments

The situations in which your processing has some of these characteristics but where HISAM is not necessarily a good choice, occur when:

- You must access dependents directly.
- You have a high number of inserts and deletes.
- Many of the database records exceed average size and must use the overflow data set. The segments that overflow into the overflow data set require additional I/O.

HISAM characteristics

For database records, HISAM databases:

- Store records in key sequence
- Can locate a particular record with a key value by using the index

For dependent segments, HISAM databases:

- Start each HISAM database record in a new logical record in the primary data set
- Store the remaining segments in one or more logical records in the overflow data set if the database record does not fit in the primary data set

An overview of how HISAM works

This topic contains Diagnosis, Modification, and Tuning information.

HISAM does not immediately reuse space. When you insert a new segment, HISAM databases shift data to make room for the new segment, and this leaves unused space after deletions. HISAM space is reclaimed when you reorganize a HISAM database.

Accessing z/OS files through IMS: GSAM

GSAM enables IMS batch application programs and BMPs to access a sequential z/OS data set as a simple database. The z/OS access methods that GSAM can use are BSAM and VSAM. A GSAM database is a z/OS data set record that is defined as a database record. The record is handled as one unit; it contains no segments or fields and the structure is not hierarchic. GSAM databases can be accessed by z/OS, IMS, and CICS.

In a CICS environment, an application program can access a GSAM database from either a Call DL/I (or EXEC DLI) batch or batch-oriented BMP program. A CICS application cannot, however, use EXEC DLI to process GSAM databases; it must use IMS calls.

You commonly use GSAM to send input to and receive output from batch-oriented BMPs or batch programs. To process a GSAM database, an application program issues calls similar to the ones it issues to process a full-function database. The program can read data sequentially from a GSAM database, and it can send output to a GSAM database.

GSAM is a sequential access method. You can only add records to an output database sequentially.

Accessing IMS data through z/OS: SHSAM and SHISAM

Two database access methods give you simple hierarchic databases that z/OS can use as data sets, SHSAM and SHISAM.

These access methods can be particularly helpful when you are converting data from z/OS files to an IMS database. SHISAM is indexed and SHSAM is not.

When you use these access methods, you define an entire database record as one segment. The segment does not contain any IMS control information or pointers; the data format is the same as it is in z/OS data sets. The z/OS access methods that SHSAM can use are BSAM and QSAM. SHISAM uses VSAM.

SHSAM and SHISAM databases can be accessed by z/OS access methods without IMS, which is useful during transitions.

Understanding how data structure conflicts are resolved

The order in which application programs need to process fields and segments within hierarchies is frequently not the same for each application. When the DBA finds a conflict in the way that two or more programs need to access the data, three options are available to solve these problems. Each of the following options solves a different kind of conflict.

- When an application program does not need access to all the fields in a segment, or if the program needs to access them in a different order, the DBA can use **field level sensitivity** for that program. *Field-level sensitivity* makes it possible for an application program to access only a subset of the fields that a segment contains, or for an application program to process a segment's fields in an order that is different from their order in the segment.
- When an application program needs to access a particular segment by a field other than the segment's key field, the DBA can use a **secondary index** for that database.
- When the application program needs to relate segments from different hierarchies, the DBA can use **logical relationships**. Using logical relationships can give the application program a logical hierarchy that includes segments from several hierarchies.

Related concepts

[“Determining mappings” on page 27](#)

When you have arranged the data aggregates into a conceptual data structure, you can examine the relationships between the data aggregates. A mapping between two data aggregates is the quantitative relationship between the two.

Using different fields: field-level sensitivity

Field-level sensitivity applies the same kind of security for fields within a segment that segment sensitivity does for segments within a hierarchy: An application program can access only those fields within a segment, and those segments within a hierarchy to which it is sensitive.

Field-level sensitivity also makes it possible for an application program to use a subset of the fields that make up a segment, or to use all the fields in the segment but in a different order. If a segment contains fields that the application program does not need to process, using field-level sensitivity enables the program not to process them.

Example of field-level sensitivity

Suppose that a segment containing data about an employee contains the fields shown in the following table. These fields are:

- Employee number: EMPNO
- Employee name: EMPNAME
- Birthdate: BIRTHDAY
- Salary: SALARY
- Address: ADDRESS

Table 24. Physical employee segment				
Employee Number	Employee Name	Birthdate	Salary	Address
EMPNO	EMPNAME	BIRTHDAY	SALARY	ADDRESS

A program that printed mailing labels for employees' checks each week would not need all the data in the segment. If the DBA decided to use field-level sensitivity for that application, the program would receive only the fields it needed in its I/O area. The I/O area would contain the EMPNAME and ADDRESS fields. The following table shows what the program's I/O area would contain.

<i>Table 25. Employee segment with field-level sensitivity</i>	
Employee Name	Address
EMPNAME	ADDRESS

Field-level sensitivity makes it possible for a program to receive a subset of the fields that make up a segment, the same fields but in a different order, or both.

Another situation in which field-level sensitivity is very useful is when new uses of the database involve adding new fields of data to an existing segment. In this situation, you want to avoid re-coding programs that use the current segment. By using field-level sensitivity, the old programs can see only the fields that were in the original segment. The new program can see both the old and the new fields.

Specifying field-level sensitivity

You specify field-level sensitivity in the PSB for the application program by using a sensitive field (SENFLD) statement for each field to which you want the application program to be sensitive.

Related reference

[SENFLD statement \(System Utilities\)](#)

Resolving processing conflicts in a hierarchy: secondary indexing

Sometimes a database hierarchy does not meet all the processing requirements of the application programs that will process it.

Secondary indexing can be used to solve two kinds of processing conflicts:

- When an application program needs to retrieve a segment in a sequence other than the one that has been defined by the segment's key field
- When an application program needs to retrieve a segment based on a condition that is found in a dependent of that segment

To understand these conflicts and how secondary indexing can resolve them, consider the examples of two application programs that process the patient hierarchy, shown in the following figure. Three segment types in this hierarchy are:

- PATIENT contains three fields: the patient's identification number, name, and address. The patient number field is the key field.
- ILLNESS contains two fields: the date of the illness and the name of the illness. The date of the illness is the key field.
- TREATMNT contains four fields: the date the medication was given; the name of the medication; the quantity of the medication that was given; and the name of the doctor who prescribed the medication. The date that the medication was given is the key field.




Figure 16. Patient hierarchy

Retrieving segments based on a different key

When an application program retrieves a segment from the database, the program identifies the segment by the segment's key field. But sometimes an application program needs to retrieve a segment in a sequence other than the one that has been defined by the segment's key field. Secondary indexing makes this possible.

Note: A new database type, the Partitioned Secondary Index (PSINDEX), is supported by the High Availability Large Database (HALDB). PSINDEX is the partitioned version of the secondary index database type. The corresponding descriptions of the secondary index database type therefore apply to PSINDEX.

For example, suppose you have an online application program that processes requests about whether an individual has ever been to the clinic. If you are not sure whether the person has ever been to the clinic, you will not be able to supply the identification number for the person. But the key field of the PATIENT segment is the patient's identification number.

Segment occurrences of a segment type (for example, the segments for each of the patients) are stored in a database in order of their keys (in this case, by their patient identification numbers). If you issue a request for a PATIENT segment and identify the segment you want by the patient's name instead of the patient's identification number, IMS must search through all of the PATIENT segments to find the PATIENT segment you have requested. IMS does not know where a particular PATIENT segment is just by having the patient's name.

To make it possible for this application program to retrieve PATIENT segments in the sequence of patients' names (rather than in the sequence of patients' identification numbers), you can index the PATIENT segment on the patient name field and store the index entries in a separate database. The separate database is called a *secondary index database*.

Then, if you indicate to IMS that it is to process the PATIENT segments in the patient hierarchy in the sequence of the index entries in the secondary index database, IMS can locate a PATIENT segment if you supply the patient's name. IMS goes directly to the secondary index and locates the PATIENT index entry with the name you have supplied; the PATIENT index entries are in alphabetical order of the patient names. The index entry is a pointer to the PATIENT segment in the patient hierarchy. IMS can determine whether a PATIENT segment for the name you have supplied exists, and then it can return the segment to the application program if the segment exists. If the requested segment does not exist, IMS indicates this to the application program by returning a not-found status code.

Related reading: For more information on HALDB, see *IMS Version 15.3 Database Administration*.

Three terms involved in secondary indexing are:

Pointer segment

The index entry in the secondary index database that IMS uses to find the segment you have requested. In the previous example, the pointer segment is the index entry in the secondary index database that points to the PATIENT segment in the patient hierarchy.

Source segment

The segment that contains the field that you are indexing. In the previous example, the source segment is the PATIENT segment in the patient hierarchy, because you are indexing on the name field in the PATIENT segment.

Target segment

The segment in the database that you are processing to which the secondary index points; it is the segment that you want to retrieve.

In the previous example, the target segment and the source segment are the same segment—the PATIENT segment in the patient hierarchy. When the source segment and the target segment are different segments, secondary indexing solves the processing conflict.

The PATIENT segment that IMS returns to the application program's I/O area looks the same as it would if secondary indexing had not been used.

The key feedback area is different. When IMS retrieves a segment without using a secondary index, IMS places the concatenated key of the retrieved segment in the key feedback area. The concatenated key contains all the keys of the segment's parents, in order of their positions in the hierarchy. The key of the root segment is first, followed by the key of the segment on the second level in the hierarchy, then the third, and so on—with the key of the retrieved segment last.

But when you retrieve a segment from an indexed database, the contents of the key feedback area after the request are a little different. Instead of placing the key of the root segment in the left-most bytes of the key feedback area, DL/I places the key of the pointer segment there. Note that the term "key of the pointer segment," as used here, refers to the key as perceived by the application program—that is, the key does not include subsequence fields.

For example, suppose index segment A shown in the following figure is indexed on a field in segment C. Segment A is the target segment, and segment C is the source segment.




Figure 17. Indexing a root segment

When you use the secondary index to retrieve one of the segments in this hierarchy, the key feedback area contains one of the following:

- If you retrieve segment A, the key feedback area contains the key of the pointer segment from the secondary index.
- If you retrieve segment B, the key feedback area contains the key of the pointer segment, concatenated with the key of segment B.
- If you retrieve segment C, the key of the pointer segment, the key of segment B, and the key of segment C are concatenated in the key feedback area.

Although this example creates a secondary index for the root segment, you can index dependent segments as well. If you do this, you create an inverted structure: the segment you index becomes the root segment, and its parent becomes a dependent.

For example, suppose you index segment B on a field in segment C. In this case, segment B is the target segment, and segment C is the source field. The following figure shows the physical database structure and the structure that is created by the secondary index.




Figure 18. Indexing a dependent segment

When you retrieve the segments in the secondary index data structure on the right, IMS returns the following to the key feedback area:

- If you retrieve segment B, the key feedback area contains the key of the pointer segment in the secondary index database.
- If you retrieve segment A, the key feedback area contains the key of the pointer segment, concatenated with the key of segment A.
- If you retrieve segment C, the key feedback area contains the key of the pointer segment, concatenated with the key of segment C.

Retrieving segments based on the qualification of a dependent segment

Sometimes an application program needs to retrieve a segment, but only if one of the dependents of the segment meet certain qualifications.

For example, suppose that the medical clinic wants to print a monthly report of the patients who have visited the clinic during that month. If the application program that processes this request does not use a secondary index, the program has to retrieve each PATIENT segment, and then retrieve the ILLNESS segment for each PATIENT segment. The program tests the date in the ILLNESS segment to determine whether the patient has visited the clinic during the current month, and prints the patient's name if the answer is yes. The program continues retrieving PATIENT segments and ILLNESS segments until it has retrieved all the PATIENT segments.

But with a secondary index, you can make the processing of the program simpler. To do this, you index the PATIENT segment on the date field in the ILLNESS segment. When you define the PATIENT segment in the DBD, you give IMS the name of the field on which you are indexing the PATIENT segment, and the name of the segment that contains the index field. The application program can then request a PATIENT segment and qualify the request with the date in the ILLNESS segment. The PATIENT segment that is returned to the application program looks just as it would if you were not using a secondary index.

In this example, the PATIENT segment is the target segment; it is the segment that you want to retrieve. The ILLNESS segment is the source segment; it contains the information that you want to use to qualify your request for PATIENT segments. The index segment in the secondary database is the pointer segment. It points to the PATIENT segments.

Creating a new hierarchy: logical relationships

When an application program needs to associate segments from different hierarchies, logical relationships can make that possible.

Logical relationships can solve the following conflicts:

- When two application programs need to process the same segment, but they need to access the segment through different hierarchies
- When a segment's parent in one application program's hierarchy acts as that segment's child in another application program

Accessing a segment through different paths

Sometimes an application program needs to process the data in a different order than the way it is arranged in the hierarchy.

For example, an application program that processes data in a purchasing database also requires access to a segment in a patient database:

- Program A processes information in the patient database about the patients at a medical clinic: the patients' illnesses and their treatments.
- Program B is an inventory program that processes information in the purchasing database about the medications that the clinic uses: the item, the vendor, information about each shipment, and information about when and under what circumstances each medication is given.

The following figure shows the hierarchies that Program A and Program B require for their processing. Their processing requirements conflict: they both need to have access to the information that is contained in the TREATMNT segment in the patient database. This information is:

- The date that a particular medication was given
- The name of the medication
- The quantity of the medication given
- The doctor that prescribed the medication

To Program B this is not information about a patient's treatment; it is information about the disbursement of a medication. To the purchasing database, this is the disbursement segment (DISBURSE).

The following figure shows the hierarchies for Program A and Program B. Program A needs the PATIENT segment, the ILLNESS segment, and the TREATMNT segment. Program B needs the ITEM segment, the VENDOR segment, the SHIPMENT segment, and the DISBURSE segment. The TREATMNT segment and the DISBURSE segment contain the same information.




Figure 19. Patient and inventory hierarchies

Instead of storing this information in both hierarchies, you can use a logical relationship. A logical relationship solves the problem by storing a pointer from where the segment is needed in one hierarchy to where the segment exists in the other hierarchy. In this case, you can have a pointer in the DISBURSE segment to the TREATMNT segment in the medical database. When IMS receives a request for information in a DISBURSE segment in the purchasing database, IMS goes to the TREATMNT segment in the medical database that is pointed to by the DISBURSE segment. The following figure shows the physical hierarchy that Program A would process and the logical hierarchy that Program B would process. DISBURSE is a pointer segment to the TREATMNT segment in Program A's hierarchy.




Figure 20. Logical relationships example

To define a logical relationship between segments in different hierarchies, you use a logical DBD. A logical DBD defines a hierarchy that does not exist in storage, but can be processed as though it does. Program B would use the logical structure shown in the previous figure as though it were a physical structure.

Inverting a parent-child relationship

Another type of conflict that logical relationships can resolve occurs when a segment's parent in one application program acts as that segment's child in another application program.

- The inventory program, Program B, needs to process information about medications using the medication as the root segment.
- A purchasing application program, Program C, processes information about which vendors have sold which medications. Program C needs to process this information using the vendor as the root segment.

The following figure shows the hierarchies for each of these application programs.




Figure 21. Supplies and purchasing hierarchies

Logical relationships can solve this problem by using pointers. Using pointers in this example would mean that the ITEM segment in the purchasing database would contain a pointer to the actual data stored in the ITEM segment in the supplies database. The VENDOR segment, however, would actually be stored in the purchasing database. The VENDOR segment in the supplies database would point to the VENDOR segment that is stored in the purchasing database.

The following figure shows the hierarchies of these two programs.




Figure 22. Program B and program C hierarchies

If you did not use logical relationships in this situation, you would:

- Keep the same data in both paths, which means that you would be keeping redundant data.
- Have the same disadvantages as separate files of data:
 - You would need to update multiple segments each time one piece of data changed.
 - You would need more storage.

Providing data security

You can control the security of data accessed by your IMS application programs through data sensitivity and processing options.

Data sensitivity

Controls what data a particular program can access.

Processing options

Controls how a particular program can process data that it can access.

Providing data availability

Specifying segment sensitivity and processing options also affects data availability. You should set the specifications so that the PCBs request the fewest SENSEGS and limit the possible processing options. With data availability, a program can continue to access and update segments in the database successfully, even though some parts of the database are unavailable.

The SENSEG statement defines a segment type in the database to which the application program is sensitive. A separate SENSEG statement must exist for each segment type. The segments can physically exist in one database or they can be derived from several physical databases. If an application program is sensitive to a segment that is below the root segment, it must also be sensitive to all segments in the path from the root segment to the sensitive segment.

Related Reading: For more information on using field-level sensitivity for data security and using the SENSEG statement to limit the scope of the PCBs, see *IMS Version 15.3 Database Administration*.

Related concepts

[“An overview of application design” on page 15](#)

The application design process varies from place to place and from application to application. The overview that is given in this section and the suggestions about documenting application design and converting existing applications are not the only way that these tasks are performed.

Keeping a program from accessing the data: data sensitivity

An IMS program can only access data to which it is sensitive.

You can control the data to which your program is sensitive on three levels:

- **Segment sensitivity** can prevent an application program from accessing all the segments in a particular hierarchy. Segment sensitivity tells IMS which segments in a hierarchy the program is allowed to access.
- **Field-level sensitivity** can keep a program from accessing all the fields that make up a particular segment. Field-level sensitivity tells IMS which fields within a particular segment a program is allowed to access.
- **Key sensitivity** means that the program can access segments below a particular segment, but it cannot access the particular segment. IMS returns only the key of this type of segment to the program.

You define each of these levels of sensitivity in the PSB for the application program. Key sensitivity is defined in the processing option for the segment. Processing options indicate to IMS exactly what a particular program may or may not do to the data. You specify a processing option for each hierarchy that the application program processes; you do this in the DB PCB that represents each hierarchy. You can specify one processing option for all the segments in the hierarchy, or you can specify different processing options for different segments within the hierarchy.

Segment sensitivity and field-level sensitivity are defined using special statements in the PSB.

Segment sensitivity

You define what segments an application program is sensitive to in the DB PCB for the hierarchy that contains those segments.

For example, suppose that the patient hierarchy shown in the following figures. The patient hierarchy is like a subset of the medical database.




Figure 23. Medical database hierarchy

PATIENT is the root segment and the parent of the three segments below it: ILLNESS, BILLING, and HOUSHOLD. Below ILLNESS is TREATMNT. Below BILLING is PAYMENT.

To make it possible for an application program to view only the segments PATIENT, ILLNESS, and TREATMNT from the medical database, you specify in the DB PCB that the hierarchy you are defining has these three segment types, and that they are from the medical database. You define the database hierarchy in the DBD; you define the application program's view of the database hierarchy in the DB PCB.

Field-level sensitivity

In addition to providing data independence for an application program, field-level sensitivity can also act as a security mechanism for the data that the program uses.

If a program needs to access some of the fields in a segment, but one or two of the fields that the program does not need to access are confidential, you can use field-level sensitivity. If you define that segment for the application program as containing only the fields that are not confidential, you prevent the program from accessing the confidential fields. Field-level sensitivity acts as a mask for the fields to which you want to restrict access.

Key sensitivity

To access a segment, an application program must be sensitive to all segments at a higher level in the segment's path. In other words, in the following figure, a program must be sensitive to segment B in order to access segment C.

For example, suppose that an application program needs segment C to do its processing. But if segment B contains confidential information (such as an employee's salary), the program is not able to access that segment. Using key sensitivity lets you withhold segment B from the application program while giving the program access to the dependents of segment B.

When a sensitive segment statement has a processing option of K specified for it, the program cannot access that segment, but the program can pass beyond that segment to access the segment's dependents. When the program does access the segment's dependents, IMS does not return that segment; IMS returns only the segment's key with the keys of the other segments that are accessed.




Figure 24. Sample hierarchy for key sensitivity example

Preventing a program from updating data: processing options

During PCB generation, you can use five options of the PROCOPT parameter (in the DATABASE macro) to indicate to IMS whether your program can read segments in the hierarchy, or whether it can also update segments.

From most restrictive to least restrictive, these options are:

G

Your program can read segments.

R

Your program can read and replace segments.

I

Your program can insert segments.

D

Your program can read and delete segments.

A

Your program can perform all the processing options. It is equivalent to specifying G, R, I, and D.

Related Reading: For a thorough description of the processing options see, *IMS Version 15.3 System Utilities*.

Processing options provide data security because they limit what a program can do to the hierarchy or to a particular segment. Specifying only the processing options the program requires ensures that the program cannot update any data it is not supposed to. For example, if a program does not need to delete segments from a database, the D option need not be specified.

When an application program retrieves a segment and has any of the just-described processing options, IMS locks the database record for that application. If PROCOPT=G is specified, other programs with the option can concurrently access the database record. If an update processing option (R, I, D, or A) is specified, no other program can concurrently access the same database record. If no updates are performed, the lock is released when the application moves to another database record or, in the case of HDAM, to another anchor point.

The following locking protocol allows IMS to make this determination. If the root segment is updated, the root lock is held at update level until commit. If a dependent segment is updated, it is locked at update level. When exiting the database record, the root segment is demoted to read level. When a program enters the database record and obtains the lock at either read or update level, the lock manager

provides feedback indicating whether or not another program has the lock at read level. This determines if dependent segments will be locked when they are accessed. For HISAM, the primary logical record is treated as the root, and the overflow logical records are treated as dependent segments.

When using block-level or database-level data sharing for online and batch programs, you can use additional processing options.

Related Reading:

- For a special case involving HISAM delete byte with parameter ERASE=YES see, *IMS Version 15.3 Database Administration*.
- For more information on database and block-level data sharing, see *IMS Version 15.3 System Administration*.

E option

With the E option, your program has exclusive access to the hierarchy or to the segment you use it with. The E option is used in conjunction with the options G, I, D, R, and A. While the E program is running, other programs cannot access that data, but may be able to access segments that are not in the E program's PCB. No dynamic enqueue by program isolation is done, but dynamic logging of database updates will be done.

GO option

When your program retrieves a segment with the GO option, IMS does not lock the segment. While the read without integrity program reads the segment, it remains available to other programs. This is because your program can only read the data (termed *read-only*); it is not allowed to update the database. No dynamic enqueue is done by program isolation for calls against this database. Serialization between the program with PROCOPT=GO and any other update program does not occur; updates to the same data occur simultaneously.

If a segment has been deleted and another segment of the same type has been inserted in the same location, the segment data and all subsequent data that is returned to the application may be from a different database record.

A read-without-integrity program can also retrieve a segment even if another program is updating the segment. This means that the program need not wait for segments that other programs are accessing. If a read-without-integrity program reads data that is being updated by another program, and that program terminates abnormally before reaching the next commit point, the updated segments might contain invalid pointers. If an invalid pointer is detected, the read-without-integrity program terminates abnormally, unless the N or T options were specified with GO. Pointers are updated during insert, delete and backout functions.

N option

When you use the N option with GO to access a full-function database or a DEDB, and the segment you are retrieving contains an invalid pointer, IMS returns a GG status code to your program. Your program can then terminate processing, continue processing by reading a different segment, or access the data using a different path. The N option must be specified as PROCOPT=GON, GON, or GONP.

T option

When you use the T option with GO and the segment you are retrieving contains an invalid pointer, the response from an application program depends on whether the program is accessing a full-function or Fast Path database.

For calls to full-function databases, the T option causes DL/I to automatically retry the operation. You can retrieve the updated segment, but only if the updating program has reached a commit point or has had its updates backed out since you last tried to retrieve the segment. If the retry fails, a GG status code is returned to your program.

For calls to Fast Path DEDBs, option T does not cause DL/I to retry the operation. A GG status code is returned. The T option must be specified as PROCOPT=GOT, GOT, or GOTP.

GOx and data integrity

For a very small set of applications and data, PROCOPT=GOx offers some performance and parallelism benefits. However, it does not offer application data integrity. For example, using PROCOPT=GOT in an online environment on a full-function database can cause performance degradation. The T option forces a re-read from DASD, negating the advantage of very large buffer pools and VSAM hiperspace for all currently running applications and shared data. For more information on the GOx processing option for DEDBs, see *IMS Version 15.3 System Utilities*.

Related concepts

“Read without integrity” on page 89

Database-level sharing of IMS databases provides for sharing of databases between a single update-capable batch or online IMS system and any number of other IMS systems that are reading data that are without integrity.

Read without integrity

Database-level sharing of IMS databases provides for sharing of databases between a single update-capable batch or online IMS system and any number of other IMS systems that are reading data that are without integrity.

A GE status code might be returned to a program using PROCOPT=GOx for a segment that exists in a HIDAM database during control interval (CI) splits.

In IMS, programs that use database-level sharing include PROCOPT=GOx in their DBPCBs for that data. For batch jobs, the DBPCB PROCOPTs establish the batch job's access level for the database. That is, a batch job uses the highest declared intent for a database as the access level for DBRC database authorization. In an online IMS environment, database ACCESS is specified on the DATABASE macro during IMS system definition, and it can be changed using the /START DB ACCESS=R0 command. Online IMS systems schedule programs with data availability determined by the PROCOPTs within those program PSBs being scheduled. That data availability is therefore limited by the online system's database access.

The PROCOPT=GON and GOT options provide certain limited PCB status code retry for some recognizable pointer errors, within the data that is being read without integrity. In some cases, dependent segment updates, occurring asynchronously to the read-without-integrity IMS instance, do not interfere with the program that is reading that data without integrity. However, update activity to an average database does not always allow a read-without-integrity IMS system to recognize a data problem.

What read without integrity means

Each IMS batch or online instance has OSAM and VSAM buffer pools defined for it. Without locking to serialize concurrent updates that are occurring in another IMS instance, a read without integrity from a database data set fetches a copy of a block or CI into the buffer pool in storage. Blocks or CIs in the buffer pool can remain there a long time. Subsequent read without integrity of other blocks or CIs can then fetch more recent data. Data hierarchies and other data relationships between these different blocks or CIs can be inconsistent.

For example, consider an index database (VSAM KSDS), which has an index component and a data component. The index component contains only hierarchic control information, relating to the data component CI where a given keyed record is located. Think of this as the way that the index component CI maintains the high key in each data component CI. Inserting a keyed record into a KSDS data component CI that is already full causes a CI split. That is, some portion of the records in the existing CI are moved to a new CI, and the index component is adjusted to point to the new CI.

For example, suppose the index CI shows the high key in the first data CI as KEY100, and a split occurs. The split moves keys KEY051 through KEY100 to a new CI; the index CI now shows the high key in the first data CI as KEY050, and another entry shows the high key in the new CI as KEY100.

A program that is reading is without integrity, which already read the "old" index component CI into its buffer pool (high key KEY100), does not point to the newly created data CI and does not attempt to access it. More specifically, keyed records that exist in a KSDS at the time a read-without-integrity program starts might never be seen. In this example, KEY051 through KEY100 are no longer in the first data CI even though the "old" copy of the index CI in the buffer pool still indicates that any existing keys up to KEY100 are in the first data CI.

Hypothetical cases also exist where the deletion of a dependent segment and the insertion of that same segment type under a different root, placed in the same physical location as the deleted segment, can cause simple Get Next processing to give the appearance of only one root in the database. For example, accessing the segments under the first root in the database down to a level-06 segment (which had been deleted from the first root and is now logically under the last root) would then reflect data from the other root. The next and subsequent Get Next calls retrieve segments from the other root.

Read-only (PROCOPT=GO) processing does not provide data integrity.

Data set extensions

IMS instances with database-level sharing can open a database for read without integrity.

After the database is opened, another program that is updating that database can make changes to the data. These changes might result in logical and physical extensions to the database data set. Because the read-without-integrity program is not aware of these extensions, problems with the RBA (beyond end-of-data) can occur.

Related concepts

[“Preventing a program from updating data: processing options” on page 87](#)

During PCB generation, you can use five options of the PROCOPT parameter (in the DATABASE macro) to indicate to IMS whether your program can read segments in the hierarchy, or whether it can also update segments.

Chapter 6. Gathering requirements for message processing options

One of the tasks of application design is providing information about your application's requirements to the people in charge of designing and administering your IMS system.

Restriction: This information applies to DB/DC and DCCTL environments only.

Related concepts

[“Programming with the IMS Java dependent region resource adapter” on page 674](#)

IMS provides a set of Java APIs called the IMS Java dependent region resource adapter to develop Java applications to run on the IMS Java dependent regions.

Identifying online security requirements

Security in an online system means protecting the data from unauthorized use through terminals. It also means preventing unauthorized use of both the IMS system and the application programs that access the database. For example, you do not want a program that processes paychecks to be available to everyone who can access the system.

The security mechanisms that IMS provides are signon, terminal, and password security.

Related reading: For an explanation of how to establish these types of security, see *IMS Version 15.3 System Administration*.

Limiting access to specific individuals: signon security

Signon security is available through Resource Access Control Facility (RACF®) or a user-written security exit routine. With signon security, individuals who want to use IMS must be defined to RACF or its equivalent before they are allowed access.

When a person signs on to IMS, RACF or security exits verify that the person is authorized to use IMS before access to IMS-controlled resources is allowed. This signon security is provided by the `/SIGN ON` command. You can also limit the transaction codes and commands that individuals are allowed to enter. You do this by associating an individual's user identification (USERID) with the transaction codes and commands.

LU 6.2 transactions contain the USERID.

Related reading: For more information on security, see *IMS Version 15.3 Communications and Connections*.

Limiting access for specific terminals: terminal security

Use terminal security to limit the entry of a transaction code to a particular terminal or group of terminals in the system. How you do this depends on how many programs you want to protect.

To protect a particular program, you can either authorize a transaction code to be entered from a list of logical terminals, or you can associate each logical terminal with a list of the transaction codes that a user can enter from that logical terminal. For example, you could protect the paycheck application program by defining the transaction code associated with it as valid only when entered from the terminals in the payroll department. If you wanted to restrict access to this application even more, you could associate the paycheck transaction code with only one logical terminal. To enter that transaction code, a user needs to be at a physical terminal that is associated with that logical terminal.

Restriction: If you are using the shared-queues option, static control blocks representing the resources needed for the security check need to be available in the IMS system where the security check is being made. Otherwise, the security check is bypassed.

Related reading: For more information on shared queues, see *IMS Version 15.3 System Administration*.

Limiting access to the program: password security

Another way you can protect the application program is to require a password when a person enters the transaction code that is associated with the application program you want to protect. If you use only password security, the person entering a particular transaction code must also enter the password of the transaction before IMS processes the transaction.

If you use password security with terminal security, you can restrict access to the program even more. In the paycheck example, using password security and terminal security means that you can restrict unauthorized individuals within the payroll department from executing the program.

Restriction: Password security for transactions is only supported if the transactions that are needed for the security check are defined in the IMS system where the security check is being made. Otherwise, the security check is bypassed.

Allowing access to security data: authorization security

RACF has a data set that you can use to store user-unique information. The AUTH call gives application programs access to the RACF data set security data, and a way to control access to application-defined resources. Thus, application programs can obtain the security information about a particular user.

How IMS security relates to Db2 for z/OS security

An important part of Db2 for z/OS security is the authorization ID. The authorization ID that IMS uses for a program or a user at a terminal depends on the kind of security that is used and the kind of program that is running.

For MPPs, IFPs, and transaction-oriented BMPs, the authorization ID depends on the type of IMS security:

- If signon is required, IMS passes the USERID and group name that are signed-on to Db2 for z/OS.
- If signon is not required, Db2 for z/OS uses the name of the originating logical terminal as the authorization ID.

For batch-oriented BMPs, the authorization ID is dependent on the value specified for the BMPUSID= keyword in the DFSDCxxx PROCLIB member:

- If BMPUSID=USERID is specified, the value from the USER= keyword on the JOB statement is used.
- If USER= is not specified on the JOB statement, the program's PSB name is used.
- If BMPUSID=PSBNAME is specified, or if BMPUSID= is not specified at all, the program's PSB name is used. If PSBNAME is not defined to RACF, the Userid of the current Address Space will be used; this will be the Home Dependent Region one, or the Control Region one if LSO=Y or if PARDLI=1 has been specified for the BMP. Userid of the current Address Space will be used also if DFSBSEX0 has returned RC08.

Supplying security information

When you evaluate your application in terms of its security requirements, you need to look at each program individually. When you have done this, you can supply the following information to your security personnel.

- For programs that require signon security:
 - List the individuals who should be able to access IMS.
- For programs that require terminal security:
 - List the transaction codes that must be secured.
 - List the terminals that should be allowed to enter each of these transaction codes. If the terminals you are listing are already installed and being used, identify the terminals by their logical terminal

names. If not, identify them by the department that will use them (for example, the accounting department).

- For programs that require password security:
 - List the transaction codes that require passwords.
- For commands that require security:
 - List the commands that require signon or password security.

Related concepts

[“An overview of application design” on page 15](#)

The application design process varies from place to place and from application to application. The overview that is given in this section and the suggestions about documenting application design and converting existing applications are not the only way that these tasks are performed.

Analyzing screen and message formats

When an application program communicates with a terminal, an editing procedure translates messages from the way they are entered at the terminal to the way the program expects to receive and process them.

The decisions about how IMS will edit your program's messages are based on how your data should be presented to the person at the terminal and to the application program. You need to describe how you want data from the program to appear on the terminal screen, and how you want data from the terminal to appear in the application program's I/O area. (The I/O area contains the segments being processed by the application program.)

To supply information that will be helpful in these decisions, you should be familiar with how IMS edits messages. IMS has two editing procedures:

- **Message Format Service (MFS)** uses control blocks that define what a message should look like to the person at the terminal and to the application program.
- **Basic edit** is available to all IMS application programs. Basic edit removes control characters from input messages and inserts the control characters you specify in output messages to the terminal.

Related reading: For information on defining IMS editing procedures and on other design considerations for IMS networks, see *IMS Version 15.3 Communications and Connections*.

An overview of MFS

MFS uses four kinds of control blocks to format messages between an application program and a terminal. The information you gather about how you want the data formatted when it is passed between the application program and the terminal is contained in these control blocks.

The two control blocks that describe input messages to IMS are:

- The device input format (DIF) describes to IMS what the input message is to look like when it is entered at the terminal.
- The message input descriptor (MID) tells IMS how the application program expects to receive the input message in its I/O area.

By using the DIF and the MID, IMS can translate the input message from the way that it is entered at the terminal to the way it should appear in the program's I/O area.

The two control blocks that describe output messages to IMS are:

- The message output descriptor (MOD) tells IMS what the output message is to look like in the program's I/O area.
- The device output format (DOF) tells IMS how the message should appear on the terminal.

To define the MFS control blocks for an application program, you need to know how you want the data to appear at the terminal and in the application program's I/O area for both input and output.

An overview of basic edit

Basic edit removes the control characters from an input message before the application program receives it, and inserts the control characters you specify when the application program sends a message back to the terminal.

To format output messages at a terminal using basic edit, you need to supply the necessary control characters for the terminal you are using.

If your application will use basic edit, you should describe how you want the data to be presented at the terminal, and what it is to look like in the program's I/O area.

Editing considerations in your application

Before you describe the editing requirements of your application, be sure that you are aware of your standards concerning screen design. Make sure that the requirements that you describe comply with those standards.

Provide the following information about your program's editing requirements:

- How you want the screen to be presented to the person at the terminal for the person to enter the input data. For example, if an airline agent wants to reserve seats on a particular flight, the screen that asks for this information might look like this:

FLIGHT#:

NAME:

NO. IN PARTY:

- What the data should look like when the person at the terminal enters the input message.
- What the input message should look like in the program's I/O area.
- What the data should look like when the program builds the output message in its I/O area.
- How the output message should be formatted at the terminal.
- The length and type of data that your program and the terminal will be exchanging.

The type of data you are processing is only one consideration when you analyze how you want the data presented at the terminal. In addition, you should weigh the needs of the person at the terminal (the human factors aspects in your application) against the effect of the screen design on the efficiency of the application program (the performance factors in the application program). Unfortunately, sometimes a trade-off between human factors and performance factors exists. A screen design that is easily understood and used by the person at the terminal may not be the design that gives the application program its best performance. Your first concern should be that you are following whatever are your established screen standards.

A terminal screen that has been designed with human factors in mind is one that puts the person at the terminal first; it is one that makes it as easy as possible for that person to interact with IMS. Some of the things you can do to make it easy for the person at the terminal to understand and respond to your application program are:

- Display a small amount of data at one time.
- Use a format that is clear and uncluttered.
- Provide clear and simple instructions.
- Display one idea at a time.
- Require short responses from the person at the terminal.
- Provide some means for help and ease of correction for the person at the terminal.

At the same time, you do not want the way in which a screen is designed to have a negative effect on the application program's response time, or on the system's performance. When you design a screen with performance first in mind, you want to reduce the processing that IMS must do with each message. To do this, the person at the terminal should be able to send a lot of data to the application program in one

screen so that IMS does not have to process additional messages. And the program should not require two screens to give the person at the terminal information that it could give on one screen.

When describing how the program should receive the data from the terminal, you need to consider the program logic and the type of data you are working with.

Gathering requirements for conversational processing

When you use *conversational processing*, the person at the terminal enters some information, and an application program processes the information and responds to the terminal. The person at the terminal then enters more information for an application program to process. Each of these interactions between the person at the terminal and the program is called a *step* in the conversation. Only MPPs can be conversational programs; Fast Path programs and BMPs cannot be conversational.

Definition: Conversational processing means that the person at the terminal can communicate with the application program.

What happens in a conversation

A *conversation* is defined as a dialog between a user at a terminal and IMS through a scratchpad area (SPA) and one or more application programs.

During a conversation, the user at the terminal enters a request, receives the information from IMS, and enters another request. Although it is not apparent to the user, a conversation can be processed by several application programs or by one application program.

To continue a conversation, the program must have the necessary information to continue processing. IMS stores data from one step of the conversation to the next in a SPA. When the same program or a different program continues the conversation, IMS gives the program the SPA for the conversation associated with that terminal.

In the preceding airline example, the first program might save the flight number and the names of the people traveling, and then pass control to another application program to reserve seats for those people on that flight. The first program saves this information in the SPA. If the second application program did not have the flight number and names of the people traveling, it would not be able to do its processing.

Designing a conversation

The first part of designing a conversation is to design the flow of the conversation. If the requests from the person at the terminal are to be processed by only one application program, you need only to design that program. If the conversation should be processed by several application programs, you need to decide which steps of the conversation each program is to process, and what each program is to do when it has finished processing its step of the conversation.

When a person at a terminal enters a transaction code that has been defined as conversational, IMS schedules the conversational program (for example, Program A) associated with that transaction code. When Program A issues its first call to the message queue, IMS returns the SPA that is defined for that transaction code to Program A's I/O area. The person at the terminal must enter the transaction code (and password, if one exists) only on the first input screen; the transaction code need not be entered during each step of the conversation. IMS treats data in subsequent screens as a continuation of the conversation started on the first screen.

After the program has retrieved the SPA, Program A can retrieve the input message from the terminal. After it has processed the message, Program A can either continue the conversation, or end it.

To continue the conversation, Program A can do any of the following:

- Reply to the terminal that sent the message.
- Reply to the terminal and pass the conversation to another conversational program, for example Program B. This is called a *deferred program switch*.

Definition: A deferred program switch means that Program A responds to the terminal and then passes control to another conversational program, Program B. After passing control to Program B, Program A is no longer part of the conversation. The next input message that the person at the terminal enters goes to Program B, although the person at the terminal is unaware that this message is being sent to a second program.

Restriction: A deferred program switch is disallowed if the application is involved in an inbound protected conversation. The application will receive an X6 status code if it attempts to perform a deferred program switch in this environment.

- Pass control of the conversation to another conversational program without first responding to the originating terminal. This is called an *immediate program switch*.

Definition: An immediate program switch lets you pass control directly to another conversational program without having to respond to the originating terminal. When you do this, the program that you pass the conversation to must respond to the person at the terminal. To continue the conversation, Program B then has the same choices as Program A did: It can respond to the originating terminal and keep control, or it can pass control in a deferred or immediate program switch.

Restriction: An immediate program switch is disallowed if the application is involved in an inbound protected conversation. The application will be abended with a U711 if it attempts to perform an immediate program switch in this environment.

To end the conversation, Program A can do either of the following:

- Move a blank to the first byte of the transaction code area of the SPA and then return the SPA to IMS.
- Respond to the terminal and pass control to a nonconversational program. This is also called a deferred program switch, but Program A ends the conversation before passing control to another application program. The second application program can be an MPP or a transaction-oriented BMP that processes transactions from the conversational program.

Important points about the scratchpad area (SPA)

When program A passes control of a conversation to program B, program B needs to have the data that program A saved in the SPA in order to continue the conversation. IMS gives the SPA for the transaction to program B when program B issues its first message call.

The SPA is kept with the message. When the truncated data option is on, the size of the retained SPA is the largest SPA of any transaction in the conversation.

For example, if the conversation starts with TRANA (SPA=100), and the program switches to a TRANB (SPA=50), the input message for TRANB will contain a SPA segment of 100 bytes. IMS adjusts the size of the SPA so that TRANB receives only the first 50 bytes.

Recovery considerations in conversations

Because a conversation involves several steps and can involve several application programs, consider the following items.

- One way you can make recovery easier is to design the conversation so that all the database updates are done in the last step of the conversation. This way, if the conversation terminates abnormally, IMS can back out all the updates because they were all made during the same step of the conversation. Updating the database during the last step of the conversation is also a good idea, because the input from each step of the conversation is available.
- Although a conversation can terminate abnormally during any step of the conversation, IMS backs out only the database updates and output messages resulting during the last step of the conversation. IMS does not back out database updates or cancel output messages for previous steps, even though some of that processing might be inaccurate as a result of the abnormal termination.
- Certain IMS system service calls can be helpful if the program determines that some of its processing was invalid. These calls include ROLB, SETS, SETU, and ROLS. The Roll Back call (ROLB) backs out all of the changes that the program has made to the database. ROLB also cancels the output messages

that the program has created (except those sent with an express PCB) since the program's last commit point.

The SETS, or SETU, and ROLS (with a *token*) calls work together to allow the application program to set intermediate backout points within the call processing of the program. The application program can set up to nine intermediate backout points. Your program needs to use the SETS or SETU call to specify a token for each point. A subsequent ROLS call, using the same token, can back out all database changes and discard all nonexpress messages processed since that SETS or SETU call.

Definition: A token is a 4-byte identifier.

- The program can use an express PCB to send a message to the person at the terminal and to the master terminal operator. When the application program inserts messages using an express PCB, IMS waits until it has the complete message, rather than for the occurrence of a commit point, to transmit the message to its destination. (In this context, "insert" refers to a situation in which the application program sends the message and it is received by IMS; "transmit" refers to a situation in which IMS begins sending the message to its destination.) Therefore, when IMS has the complete message, it will be transmitted even if the program abnormally terminates. Messages sent with an express PCB are sent to their final destinations even if the program terminates abnormally or issues a ROLB call.
- To verify the accuracy of the previous processing, and to correct the processing that is determined to be inaccurate, you can use the Conversational Abnormal termination routine, DFSCONE0.

Related reading: For more information on DFSCONE0, see *IMS Version 15.3 Exit Routines*.

- You can write an MPP to examine the SPA, send a message notifying the person at the terminal of the abnormal termination, make any necessary database calls, and use a user-written or system-provided exit routine to schedule it.

Related concepts

[“To other programs and terminals” on page 98](#)

When you want to send an output message to a terminal other than, or in addition to, the terminal that sent the input message, you use an alternate PCB. You can set the alternate PCB for a specific logical terminal when the program's PSB is generated, or you can define the alternate PCB as being modifiable. A program can change the destination of a modifiable alternate PCB while the program is running, so you can send output messages to several alternate destinations.

Identifying output message destinations

An application program can send messages to another application program or to IMS terminals. To send output messages, the program issues a call and references the I/O PCB or an alternate PCB. The I/O PCB and alternate PCBs represent logical terminals and other application programs with which the application program communicates.

Definition: An *alternate PCB* is a data communication program communication block (DCPCB) that you define to describe output message destinations other than the terminal that originated the input message.

Related concepts

[“Batch message processing: transaction-oriented BMPs” on page 44](#)

Transaction-oriented BMPs can access z/OS files, GSAM databases, Db2 for z/OS databases, full-function databases, DEDBs, and MSDBs.

The originating terminal

To send a message to the logical terminal that sent the input message, the program uses an I/O PCB. IMS puts the name of the logical terminal that sent the message in the I/O PCB when the program receives the message.

As a result, the program need not do anything to the I/O PCB before sending the message. If a program receives a message from a batch-oriented BMP or CPI Communications driven program, no logical terminal name is available to put into the I/O PCB. In these cases, the logical terminal name field contains blanks.

Related concepts

[“Identifying output message destinations” on page 97](#)

An application program can send messages to another application program or to IMS terminals. To send output messages, the program issues a call and references the I/O PCB or an alternate PCB. The I/O PCB and alternate PCBs represent logical terminals and other application programs with which the application program communicates.

To other programs and terminals

When you want to send an output message to a terminal other than, or in addition to, the terminal that sent the input message, you use an alternate PCB. You can set the alternate PCB for a specific logical terminal when the program's PSB is generated, or you can define the alternate PCB as being modifiable. A program can change the destination of a modifiable alternate PCB while the program is running, so you can send output messages to several alternate destinations.

The application program might need to respond to the originating terminal before the person at the originating terminal can send any more messages. This might occur when a terminal is in **response mode** or in **conversational mode**:

- **Response mode** can apply to a communication line, a terminal, or a transaction. When response mode is in effect, IMS does not accept any input from the communication line or terminal until the program has sent a response to the previous input message. The originating terminal is unusable (for example, the keyboard locks) until the program has processed the transaction and sent the reply back to the terminal.

If a response-mode transaction is processed, including Fast Path transactions, and the application does not insert a response back to the terminal through either the I/O PCB or alternate I/O PCB, but inserts a message to an alternate PCB (program-to-program switch), the second or subsequent application program must respond to the originating terminal and satisfy the response. IMS will not take the terminal out of response mode.

If an application program terminates normally and does not issue an ISRT call to the I/O PCB, alternate I/O PCB, or alternate PCB, IMS sends system message DFS2082I to the originating terminal to satisfy the response for all response-mode transactions, including Fast Path transactions.

You can define communication lines and terminals as operating in response mode, not operating in response mode, or operating in response mode only if processing a transaction that is been defined as response mode. You specify response mode for communication lines and terminals on the TYPE and TERMINAL macros, respectively, at IMS system definition. You can define any transaction as a response-mode transaction; you do this on the TRANSACT macro at IMS system definition. Response mode is in effect if:

- The communication line has been defined as being in response mode.
- The terminal has been defined as being in response mode.
- The transaction code has been defined as response mode.

- **Conversational mode** applies to a transaction. When a program is processing a conversational transaction, the program must respond to the originating terminal after each input message it receives from the terminal.

In these processing modes, the program must respond to the originating terminal. But sometimes the originating terminal is a physical terminal that is made up of two components—for example, a printer and a display. If the physical terminal is made up of two components, each component has a different logical terminal name. To send an output message to the printer part of the terminal, the program must use a different logical terminal name than the one associated with the input message; it must send the output message to an alternate destination. A special kind of alternate PCB is available to programs in these situations; it is called an *alternate response PCB*.

Definition: An alternate response PCB lets you send messages when exclusive, response, or conversational mode is in effect. See the next section for more information.

Alternate response PCB

The destination of an alternate response PCB must be a logical terminal—you cannot use an alternate response PCB to represent another application program. When you use an alternate response PCB during response mode or conversational mode, the logical terminal represented by the alternate response PCB must represent the same physical terminal as the originating logical terminal.

In these processing modes, after receiving the message, the application program must respond by issuing an ISRT call to one of the following:

- The I/O PCB.
- An alternate response PCB.
- An alternate PCB whose destination is another application program, that is, a program-to-program switch.
- An alternate PCB whose destination is an ISC link. This is allowed only for front-end switch messages.

Related reading: For more information on front-end switch messages, see *IMS Version 15.3 Exit Routines*.

If one of these criteria is not met, message DFS2082I is sent to the terminal.

Express PCB

Consider specifying an alternate PCB as an *express PCB*. The express designation relates to whether a message that the application program inserted is actually transmitted to the destination if the program abnormally terminates or issues a ROLL, ROLB, or ROLS call. For all PCBs, when a program abnormally terminates or issues a ROLL, ROLB, or ROLS call, messages that were inserted but not made available for transmission are cancelled while messages that were made available for transmission are never cancelled.

Definition: An express PCB is an alternate response PCB that allows your program to transmit the message to the destination terminal earlier than when you use a nonexpress PCB.

For a nonexpress PCB, the message is not made available for transmission to its destination until the program reaches a commit point. The commit point occurs when the program terminates, issues a CHKP call, or requests the next input message and when the transaction has been defined with MODE=SNGL.

For an express PCB, when IMS has the complete message, it makes the message available for transmission to the destination. In addition to occurring at a commit point, it also occurs when the application program issues a PURG call using that PCB or when it requests the next input message.

You should provide the answers to the following questions to the data communications administrator to help in meeting your application's message processing requirements:

- Will the program be required to respond to the terminal before the terminal can enter another message?
- Will the program be responding only to the terminal that sends input messages?
- If the program needs to send messages to other terminals or programs as well, is there only one alternate destination?
- What are the other terminals to which the program must send output messages?
- Should the program be able to send an output message before it terminates abnormally?

Related concepts

[“Recovery considerations in conversations” on page 96](#)

Because a conversation involves several steps and can involve several application programs, consider the following items.

[“Identifying output message destinations” on page 97](#)

An application program can send messages to another application program or to IMS terminals. To send output messages, the program issues a call and references the I/O PCB or an alternate PCB. The I/O PCB and alternate PCBs represent logical terminals and other application programs with which the application program communicates.

Chapter 7. Designing an application for APPC

Advanced Program-to-Program Communication (APPC) is IBM's preferred protocol for program-to-program communication. Application programs can be distributed throughout the network and communicate with each other in many hardware architectures and software environments.

Related Reading: For more information on APPC, see:

- *IMS Version 15.3 Communications and Connections*, which includes an overview of APPC for LU 6.2 devices and CPI Communications concepts.

Overview of APPC and LU 6.2

APPC allows application programs using APPC protocols to enter IMS transactions from LU 6.2 devices. The LU 6.2 application program runs on an LU 6.2 device supporting APPC.

APPC creates an environment that allows:

- Remote LU 6.2 devices to enter IMS local and remote transactions
- IMS application programs to insert transaction output to LU 6.2 devices with no coding changes to existing application programs
- New application programs to make full use of LU 6.2 device facilities
- Data integrity provided by IMS and in LU 6.2 environments that do not have a distributed sync-point function

Application program types

APPC/IMS is part of IMS TM that uses the CPI communications interface to communicate with application programs.

APPC/IMS supports the following types of application programs for LU 6.2 processing:

- Standard DL/I
- Modified standard DL/I
- CPI Communications driven

Standard DL/I application program

A standard DL/I application program does not issue any CPI Communications calls or establish any CPI-C conversations. This application program can communicate with LU 6.2 products that replace other LU-type terminals using the IMS API. A standard DL/I application program does not need to be modified, recompiled, or bound, and it executes as it currently does.

Modified standard DL/I application program

A modified standard DL/I application program is a standard DL/I online IMS TM application program that uses both DL/I calls and CPI Communications calls. It can be an MPP, BMP, or IFP that can access full-function databases, DEDBs, MSDBs, and Db2 for z/OS databases.

A modified standard DL/I application program uses CPI Communications (CPI-C) calls to provide support for an LU 6.2 and non-LU 6.2 mixed network. The same application program can be a standard DL/I on one execution, when the CPI Communications ALLOCATE verb is not issued, and a modified standard DL/I on a different execution when the CPI Communications ALLOCATE verb is issued.

A modified standard DL/I application program receives its messages using DL/I GU calls to the I/O PCB and issues output responses using DL/I ISRT calls. CPI Communications calls can also be used to allocate new conversations and to send and receive data for them.

Related Reading: For a list of the CPI Communications calls, see CPI Communications Reference.

Use a modified standard DL/I application program when you want to use an existing standard DL/I application program to establish a conversation with another LU 6.2 device or the same network destination. The standard DL/I application program is optionally modified and uses new functions, new application and transaction definitions, and modified DL/I calls to initiate LU 6.2 application programs. Program calls and parameters are available to use the IMS-provided implicit API and the CPI Communications explicit API.

CPI Communications driven program

A CPI Communications driven application program uses Commit and Backout calls, and CPI Communications interface calls or LU 6.2 verbs for input and output message processing. This application program uses the CPI Communications explicit API, and can access full-function databases, DEDBs, MSDBs, and Db2 for z/OS databases. An LU 6.2 device can activate a CPI Communications driven application program only by allocating a conversation.

Unlike a standard DL/I or modified standard DL/I application program, input and output message processing for a CPI Communications driven program uses APPC/MVS buffers and bypasses IMS message queueing. Because these application programs do not use the IMS message queue, they can control their own execution with the partner LU 6.2 system. An IMS APSB call enables you to allocate a PSB for accessing IMS databases and alternate PCBs.

The application program uses the Common Programming Interface Resource Recovery (CPI-RR) SRRCMIT verb to initiate an IMS sync point and the CPI-RR SRRBACK verb for backout. CPI Communications driven application programs use the CPI-RR calls to initiate IMS sync point processing prior to program termination.

A CPI Communications driven application program is able to:

- Access any type of database
- Receive and send large messages like the standard DL/I and modified standard DL/I application programs
- Control the flow of input and output with CPI Communications calls
- Allocate multiple conversations with partner LU 6.2 devices
- Cause synchronization with conversation partners
- Use the IMS implicit API (for example, IMS queue services)
- Use IMS services (for example, sync point at program termination) regardless of the API that is used

Application objectives

Each application type has a different purpose, and its ease-of-use varies depending on whether the program is a standard DL/I, modified standard DL/I, or a CPI Communications driven application program.

The following table lists the purpose and ease-of-use for each application type (standard DL/I, modified standard DL/I, and PI-C driven). This information must be balanced with IMS resource use.

Table 26. Using application programs in APPC

Purpose of application program	Ease of use		
	Standard DL/I program	Modified standard DL/I program	PI-C driven program
Inquiry	Easy	Neutral	Very Difficult
Data Entry	Easy	Easy	Difficult
Bulk Transfer	Easy	Easy	Neutral
Cooperative	Difficult	Difficult	Desirable

Table 26. Using application programs in APPC (continued)

Purpose of application program	Ease of use		
	Standard DL/I program	Modified standard DL/I program	PI-C driven program
Distributed	Difficult	Neutral	Desirable
High Integrity	Neutral	Neutral	Desirable
Client Server	Easy	Neutral	Very Difficult

Choosing conversation attributes

The LU 6.2 transaction program indicates how the transaction is to be processed by IMS. Two processing modes are available: **synchronous** and **asynchronous**.

Synchronous conversation

A conversation is synchronous if the partner waits for the response on the same conversation used to send the input data.

Synchronous processing is requested by issuing the RECEIVE_AND_WAIT verb after the SEND_DATA verb. Use this mode for IMS response-mode transactions and IMS conversational-mode transactions.

Example:

```
MC_ALLOCATE TPN(MYTXN)
MC_SEND_DATA 'THIS CAN BE A RESPONSE MODE'
MC_SEND_DATA 'OR CONVERSATIONAL MODE'
MC_SEND_DATA 'IMS TRANSACTION'
MC_RECEIVE_AND_WAIT
```

Asynchronous conversation

A conversation is asynchronous if the partner program normally deallocates a conversation after sending the input data. Output is sent to the TP name of DFSASYNC.

Asynchronous processing is requested by issuing the DEALLOCATE verb after the SEND_DATA verb. Use asynchronous processing for IMS commands, message switches, and non-response, non-conversational transactions.

Example:

```
MC_ALLOCATE TPN(OTHERTXN)
MC_SEND_DATA 'THIS MUST BE A MESSAGE SWITCH, IMS COMMAND'
MC_SEND_DATA 'OR A NON-RESP NON-CONV TRANSACTION'
MC_DEALLOCATE
```

Asynchronous output delivery

Asynchronous output is held on the IMS message queue for delivery. When the output is enqueued, IMS attempts to allocate a conversation to send this output. If this fails, IMS holds the output for later delivery. This delivery can be initiated by an operator command (/ALLOC), or by the enqueue of a new message for this LU 6.2 destination.

MSC synchronous and asynchronous conversation

MSC remote application messages from both synchronous and asynchronous APPC conversations can be queued on the multiple systems coupling (MSC) link. These messages can then be sent across the MSC link to a remote IMS for processing.

Related concepts

[“LU 6.2 flow diagrams” on page 111](#)

The following diagrams show the flows for transactions that are sent from an LU 6.2 device.

Conversation type

The APPC conversation type defines how data is passed on and retrieved from APPC verbs.

It is similar in concept to file blocking and affects both ends of the conversation.

APPC supports two types of conversations:

Basic conversation

This low-conversation allows programs to exchange data in a standardized format. This format is a stream of data containing 2-byte length fields (referred to as LLs) that specify the amount of data to follow before the next length field. The typical data pattern is:

```
LL, data, LL, data
```

Each grouping of LL, data is referred to as a logical record. A basic conversation is used to send multiple segments with one verb and to receive maximum data with one verb.

Mapped conversation

This high-conversation allows programs to exchange arbitrary data records in data formats approved by application programmers. One send verb results in one receive verb, and z/OS and VTAM® handle the buffering.

Related Reading: For more information on basic and mapped conversations, see

- *Systems Network Architecture: LU 6.2 Reference: Peer Protocols* and
- *Systems Network Architecture: Transaction Programmer's Reference Manual for LU Type 6.2*

Conversation state

CPI Communications uses conversation state to determine what the next set of actions will be.

Examples of conversation states are:

RESET

The initial state before communications begin.

SEND

The program can send or optionally receive.

RECEIVE

The program must receive or abort.

CONFIRM

The program must respond to a partner.

The basic rules for APPC verbs are:

- The program that initiates the conversation speaks first.
- Only one APPC verb can be outstanding at time.
- Programs take turns sending and receiving.
- The state of the conversation determines the verbs a program can issue.

Synchronization level

The APPC synchronization level defines the protocol that is used when changing conversation states.

APPC and IMS support the following synchronization level values:

SYNCLVL=NONE

Specifies that the programs do not issue calls or recognize returned parameters relating to synchronization.

SYNCLVL=CONFIRM

Specifies that the programs can perform confirmation processing on the conversation.

SYNCLVL=SYNCPT

Specifies that the programs participate in coordinated commit processing on resources that are updated during the conversation under the z/OS Resource Recovery Services (RRS) recovery platform. A conversation with this level is also called a *protected conversation*.

Additionally, either IMS or RRS can be specified as the synchronization point manager.

RRS=Y

If AOS=B, AOS=S, or AOS=X, transactions with SYNCLVL=NONE or CONFIRM are processed with IMS as the synchronization point manager.

If AOS=B or AOS=Y, transactions with SYNCLVL=SYNCPT are processed with RRS as the synchronization point manager.

In a shared message queue environment where the front-end IMS system is also the back-end IMS system, transactions with SYNCLVL=SYNCPT are processed with RRS as the synchronization point manager.

In a non-shared message queue environment, transactions with SYNCLVL=SYNCPT are processed with RRS as the synchronization point manager.

Restriction: The AOS= setting is applicable to shared message queue environment only.

RRS=N

If AOS=B, AOS=S, or AOS=X, transactions with SYNCLVL=NONE or CONFIRM are processed with IMS as the synchronization point manager.

If the back-end IMS system has RRS=N specified, transactions with SYNCLVL=SYNCPT are processed only at the front-end IMS system. However, if the front-end IMS system also has RRS=N specified, transactions with SYNCLVL=SYNCPT are not processed at all.

Allocating a conversation with SYNCLVL=SYNCPT requires the RRS as the synchronization point manager. RRS controls the commitment of protected resources by coordinating the commit or backout request with the participating owners of the updated resources, the resource managers. IMS is the resource manager for DL/I, Fast Path data, and the IMS message queues. The application program decides whether the data is to be committed or aborted and communicates this decision to the synchronization point manager. The synchronization point manager then coordinates the actions in support of this decision among the resource managers.

Related concepts

[Activating protected conversations \(Communications and Connections\)](#)

Introduction to resource recovery

Most customers maintain computer resources that are essential to the survival of their businesses. When these resources are updated in a controlled and synchronized manner, they are said to be **protected resources** or **recoverable resources**. These resources can all reside locally (on the same system) or be distributed (across nodes in the network). The protocols and mechanisms for regulating the updating of multiple protected resources in a consistent manner is provided in z/OS with z/OS Resource Recovery Services (RRS).

Participants in resource recovery

As shown in the following figure, the Resource Recovery environment is composed of three participants:

- Sync-point manager
- Resource managers

- Application program

RRS is the **sync-point manager**, also known as the coordinator. The sync-point manager controls the commitment of protected resources by coordinating the commit request (or backout request) with the **resource managers**, the participating owners of the updated resources. These resource managers are known as participants in the sync-point process. IMS participates as a resource manager for DL/I, Fast Path, and Db2 for z/OS data if this data has been updated in such an environment.

The final participant in this resource recovery protocol is the **application program**, the program accessing and updating protected resources. The application program decides whether the data is to be committed or aborted and relates this decision to the sync-point manager. The sync-point manager then coordinates the actions in support of this decision among the resource managers.




Figure 25. Participants in resource recovery

Two-phase commit protocol

As shown in the following figure, the two-phase commit protocol is a process involving the sync-point manager and the resource manager participants to ensure that of the updates made to a set of resources by a third participant, the application program, either **all updates occur or none**. In simple terms, the application program decides to commit its changes to some resources; this commit is made to the sync-point manager that then polls all of the resource managers as to the feasibility of the commit call. This is the prepare phase, often called phase one. Each resource manager votes yes or no to the commit.

After the sync-point manager has gathered all the votes, phase two begins. If all votes are to commit the changes, then the phase two action is commit. Otherwise, phase two becomes a backout. System failures, communication failures, resource manager failures, or application failures are not barriers to the completion of the two-phase commit process.

The work done by various resource managers is called a *unit of recovery (UOR)* and spans the time from one consistent point of the work to another consistent point, usually from one commit point to another. It is the unit of recovery that is the object of the two-phase commit process.




Figure 26. Two-phase commit process with one resource manager

Notes:

1. The application and IMS make a connection.
2. IMS expresses protected interest in the work started by the application. This tells RRS that IMS will participate in the 2-phase commit process.
3. The application makes a read request to an IMS resource.
4. Control is returned to the application following its read request.
5. The application updates a protected resource.
6. Control is returned to the application following its update request.
7. The application requests that the update be made permanent by way of the SRRCMIT call.
8. RRS calls IMS to do the prepare (phase 1) process.
9. IMS returns to RRS with its vote to commit.
10. RRS calls IMS to do the commit (phase 2) process.
11. IMS informs RRS that it has completed phase 2.
12. Control is returned to the application following its commit request.

Local versus distributed

The residence of the participants involved in the recovery process determines whether that recovery is considered local or distributed. In a **local** recovery scenario, all the participants reside on the same single system. In a **distributed** recovery scenario, the participants are scattered over multiple systems. The following figure shows the communication between Resource Manager participants in a distributed resource recovery. There is no conceptual difference between a local and distributed recovery in the functions provided by RRS. However, to distribute the original sync-point manager's function to

involve remote sync-point managers, a special resource manager is required. The APPC communications resource manager provides this support in the distributed environment.




Figure 27. Distributed resource recovery

Summary of z/OS Resource Recovery Services support

z/OS Resource Recovery Services (RRS) provides a system resource recovery platform so that applications running on z/OS can access local and distributed resources and have system coordinated recovery management of these resources.

The support includes:

- A synchronization point manager to coordinate the two-phase commit process
- Implementation of the SAA Commit and Backout callable services for use by application programs
- A mechanism to associate resources with an application instance
- Services for resource manager registration and participation in the two-phase commit process with RRS
- Services to allow resource managers to express interest in an application instance and be informed of commit and backout requests
- Services to enable resource managers to obtain system data to restore their resources to consistent state

- A communications resource manager (called APPC/PC for APPC/Protected Conversations) so that distributed applications can coordinate their recovery with participating local resource managers

Restrictions:

- Extended Recovery Facility (XRF)

Running protected conversations in an IMS-XRF environment does not guarantee that the alternate system can resume and resolve any unfinished work started by the active system. This process is not guaranteed because a failed resource manager must re-register with its original RRS system if the RRS is still available when the resource manager restarts. Only if the RRS on the active system is not available can an XRF alternate can register with another RRS in the sysplex and obtain the incomplete unit of recovery data of the failing active.

Recommendation: Because IMS retains indoubt units-of-recovery indefinitely until they are resolved, switch back to the original active system as soon as possible to pick up unit-of-recovery information to resolve and complete all the work of the resource managers involved. If this is not possible, the indoubt units-of-recovery can be resolved using commands.

- Batch and non-message-driven BMPs in a DBCTL Environment

Distributed Sync Point does not support the IMS batch environment. In a DBCTL environment, inbound protected conversations are not possible. However, a BMP in a DBCTL environment can allocate an outbound protected conversation, which will be supported by Distributed Sync Point and RRS.

Distributed sync point

The Distributed Sync Point support enables IMS and remote application programs (APPC or OTMA) to participate in protected conversations with coordinated resource updates and recoveries. Before this support, IMS acted as the sync-point manager. In this new scenario, z/OS manages the sync-point process on behalf of the conversation participants: the application program and IMS (now acting as a resource manager).

z/OS implements a system resource recovery platform, the z/OS Resource Recovery Services (RRS). RRS supports the Common Programming Interface - Resource Recovery (CPI-RR), an element of the SAA Common Programming Interface that defines resource recovery and provides for the coordinated management of resource recovery for both local *and* distributed resources. In addition to RRS, a communications resource manager (called APPC/PC for APPC/Protected Conversations) provides distribution of the recovery.

In the APPC environment, a protected conversation is initiated when the application program allocates an APPC conversation with SYNC_LEVEL=SYNCPT. Both IMS and APPC are resource managers in this scenario. In the OTMA environment, some additional code is required because OTMA is not a resource manager. The additional code needed is an OTMA adapter, IBM supplied or equivalent. This adapter indicates to IMS (in the OTMA message prefix) that this message is part of a protected conversation, and thus IMS and the adapter are participants in the coordinated commit process as managed by RRS.

Application programmers can now develop APPC application programs (local and remote) and remote OTMA application programs that use RRS as the sync-point manager, rather than IMS. This enhancement enables resources across multiple platforms to be updated and recovered in a coordinated manner.

Distributed sync point concepts

The Distributed Sync Point support entails:

- Changes in IMS that allow it to function as a resource manager under RRS
- Changes to the application program environment that support using applications in protected conversations
- Changes to some commands that aid the user

Impact on the network

Network traffic will increase as a result of the conversation participants and the sync-point manager communicating with each other.

Application programming interface for LU type 6.2

IMS application programs can use the IMS implicit LU 6.2 API to access LU 6.2 devices. This API provides compatibility with non-LU 6.2 device types so that the same application program can be used from both LU 6.2 and non-LU 6.2 devices.

The API adds to the APPC interface by supplying IMS-provided processing for the application program. You can use the explicit CPI Communications interface for APPC functions and facilities for new or rewritten IMS application programs.

Implicit API

The implicit API accesses an APPC conversation indirectly. This API uses the standard DL/I calls (GU, ISRT, PURG) to send and receive data. It allows application programs that are not specific to LU 6.2 protocols to use LU 6.2 devices.

The API uses new and changed DL/I calls (CHNG, INQY, SETO) to utilize LU 6.2. Using the existing IMS application programming base, you can write specific applications for LU 6.2 using this API and not using the CPI Communications calls. Although the implicit API uses only some of the LU 6.2 capabilities, it can be a useful simplification for many applications. The implicit API also provides function outside of LU 6.2, like message queueing and automatic asynchronous message delivery.

IMS generates all CPI Communications calls under the implicit API. The application interaction is strictly with the IMS message queue.

The remote LU 6.2 system must be able to handle the LU 6.2 flows. APPC/MVS generates these flows from the CPI Communications calls issued by the IMS application program using the implicit API. An IMS application program can use the explicit API to issue the CPI Communications directly. This is useful with remote LU 6.2 systems that have incomplete LU 6.2 implementations, or that are incompatible with the IMS implicit API support.

The existing API is extended so that:

- Asynchronous LU 6.2 output is created by using alternate PCBs that reference LU 6.2 destinations. The DL/I CHNG call can supply parameters to specify an LU 6.2 destination. Default values are used for omitted parameters.
- An application program can retrieve the current conversation attributes such as the conversation type (basic or mapped), the sync_level (NONE, CONFIRM, or SYNCPT), and asynchronous or synchronous conversation.
- A terminal message switch can be used to and from LU 6.2 devices.

Explicit API

The explicit API (the CPI Communications API) can be used by any IMS application program to access an APPC conversation directly.

IMS resources are available to the CPI Communications driven application program only if the application issues the APSB (Allocate PSB) call. The CPI Communications driven application program must use the CPI-RR SRRCMIT and SRRBACK verbs to initiate an IMS sync point or backout, or if SYNCLVL=SYNCPT is specified, to communicate the sync point decision to the z/OS Resource Recovery Services sync point manager.

Related Reading: For a description of the SRRCMIT and SRRBACK verbs, see *SAA CPI Resource Recovery Reference*.

LU 6.2 partner program design

The flow of a transaction that is sent from an LU 6.2 device differs, depending on the conversation attributes and synchronization levels. Different results occur, and the partner system takes actions accordingly.

LU 6.2 flow diagrams

The following diagrams show the flows for transactions that are sent from an LU 6.2 device.

The following figures show:

- The flow between a synchronous or asynchronous LU 6.2 application program and an IMS application program in a single (local) IMS system
- The flow between a synchronous or asynchronous LU 6.2 application program in a single (local) IMS system and an IMS application program in a remote IMS system across a multiple systems coupling (MSC) link
- A backout scenario with SYNC_LEVEL=SYNCPT

Differences in buffering and encapsulation of control data with user data may cause variations in the flows. The control data are the 3 returned fields from the Receive APPC verb: Status_received, Data_received, and Request_to_send_received. Any variations based on these differences will not affect the function or use of the flows.




Figure 28. Flow of a local IMS synchronous transaction when Sync_level=None

Figure 29 on page 113 shows the flow of a local synchronous transaction when Sync_level is Confirm.




Figure 29. Flow of a local IMS synchronous transaction when Sync_level=Confirm

Figure 30 on page 114 shows the flow of a local asynchronous transaction when Sync_level is None.




Figure 30. Flow of a local IMS asynchronous transaction when Sync_level=None

Figure 31 on page 115 shows the flow of a local asynchronous transaction when Sync_level is Confirm.




Figure 31. Flow of a local IMS asynchronous transaction when `Sync_level=Confirm`

The following figure shows the flow of a local conversational transaction When `Sync_level` is None.




Figure 32. Flow of a local IMS conversational transaction when Sync_level=None

The following figure shows the flow of a local IMS command when Sync_level is None.




Figure 33. Flow of a local IMS command when Sync_level=None

The following figure shows the flow of a local asynchronous command when Sync_level is Confirm.




Figure 34. Flow of a local IMS asynchronous command when Sync_level=Confirm

The following figure shows the flow of a message switch When Sync_level is None.




Figure 35. Flow of a message switch when `Sync_level=None`

Synchronous is used to verify that no error has occurred while processing DFSAPPC. If an error occurred, the error message returns before DEALLOCATE.

The following figure shows the flow of a CPI-C driven program when `Sync_level` is None.




Figure 36. Flow of a local CPI communications driven program when Sync_level=None

The following figure shows the flow of a remote synchronous transaction when Sync_level is None.




Figure 37. Flow of a remote IMS synchronous transaction when Sync_level=None

The following figure shows the flow of a remote asynchronous transaction when Sync_level is None.




Figure 38. Flow of a remote IMS asynchronous transaction when Sync_level=None

The following figure shows the flow of a remote asynchronous transaction when Sync_level is Confirm.




Figure 39. Flow of a remote IMS asynchronous transaction when Sync_level=Confirm

The following figure shows the flow of a remote synchronous transaction when Sync_level is Confirm.




Figure 40. Flow of a remote IMS synchronous transaction when Sync_level=Confirm

The scenarios shown in the following figure provide examples of the two-phase process for the supported application program types. The LU 6.2 verbs are used to illustrate supported functions and interfaces between the components. Only parameters pertinent to the examples are included. This does not imply that other parameters are not supported.

The following figure shows a standard DL/I program commit scenario when Sync_Level=Syncpt.




Figure 41. Standard DL/I program commit scenario when Sync_Level=Syncpt

Notes:

- 1 Sync_Level=Syncpt triggers a protected resource update.
- 2 This application program inserts output for the remote application to the IMS message queue.
- 3 The GU initiates the transfer of the output.
- 4 The remote application sends a Confirmed after receiving data (output).
- 5 IMS issues ATRCMIT (equivalent to SRRCMIT) to start the two-phase process.

The following figure shows a CPI-C driven commit scenario when Sync_Level=Syncpt.




Figure 42. CPI-C driven commit scenario when Sync_Level=Syncpt

Notes:

- 1 Sync_Level=Syncpt triggers a protected resource update.
- 2 The programs send and receive data.
- 3 The remote application decides to commit the updates.
- 4 The CPI-C program issues SRRCMIT to commit the changes.
- 5 The commit return code is returned to the remote application.

The following figure shows a standard DL/I program backout scenario when Sync_Level=Syncpt.




Figure 43. Standard DL/I program U119 backout scenario when Sync_Level=Syncpt

Notes:

- 1 Sync_Level=Syncpt triggers a protected-resource update.
- 2 This application program inserts output for the remote application to the IMS message queue.
- 3 The GU initiates the transfer of the output.
- 4 The remote application decides to back out any updates.
- 5 IMS abends the application with a U119 to back out the application.
- 6 The backout return code is returned to the remote application.

The following figure shows a standard DL/I program backout scenario when Sync_Level=Syncpt.




Figure 44. Standard DL/I program U0711 backout scenario when Sync_Level=Syncpt

Notes:

- 1 Sync_Level=Syncpt triggers a protected-resource update.
- 2 This application program inserts output for the remote application to the IMS message queue.
- 3 The GU initiates the transfer of the output.
- 4 The remote application sends a Confirmed after receiving data (output).
- 5 IMS issues ATBRCVW on behalf of the DL/I application to wait for a commit or backout.
- 6 The remote application decides to back out any updates.
- 7 IMS abends the application with U0711 to back out the application.
- 8 The backout return code is returned to the remote application.

The following figure shows a standard DL/I program ROLB scenario when Sync_Level=Syncpt.




Figure 46. Multiple transactions in same commit when `Sync_Level=Syncpt`

Notes:

- 1 An allocate with `Sync_Level=Syncpt` triggers a protected resource update with Conversation 1.
- 2 The first transaction provides the output for Conversation 1.
- 3 An allocate with `Sync_Level=Syncpt` triggers a protected resource update with Conversation 2.
- 4 The second transaction provides the output for Conversation 2.
- 5 The remote application issues `SRRCMIT` to commit both transactions.
- 6 IMS issues `ATRMIT` to start the two-phase process on behalf of each DL/I application.

Related concepts

[“Application objectives” on page 102](#)

Each application type has a different purpose, and its ease-of-use varies depending on whether the program is a standard DL/I, modified standard DL/I, or a CPI Communications driven application program.

Integrity tables

The following tables show the message integrity of conversations, results of processing when integrity is compromised, and how IMS recovers APPC messages.

The following table shows the results, from the viewpoint of the IMS partner system, of normal conversation completion, abnormal conversation completion due to a session failure, and abnormal conversation completion due to non-session failures. These results apply to asynchronous and synchronous conversations and both input and output. This table also shows the outcome of the message, and the action that the partner system takes when it detects the failure. An example of an action, under "LU 6.2 Session Failure," is a programmable work station (PWS) resend.

Table 27. Message integrity of conversations

Conversation attributes	Normal	LU 6.2 session failure ¹	Other failure ²
Synchronous Sync_level=NONE	Input: Reliable Output: Reliable	Input: PWS resend Output: PWS resend	Input: Reliable Output: Reliable
Synchronous Sync_level=CONFIRM	Input: Reliable Output: Reliable	Input: PWS resend Output: Reliable	Input: Reliable Output: Reliable
Synchronous Sync_level=SYNCPT	Input: Reliable Output: Reliable	Input: PWS resend Output: Reliable	Input: Reliable Output: Reliable
Asynchronous Sync_level=NONE	Input: Ambiguous Output: Reliable	Input: Undetectable Output: Reliable	Input: Undetectable Output: Reliable
Asynchronous Sync_level=CONFIRM	Input: Reliable Output: Reliable	Input: PWS resend Output: Reliable	Input: Reliable Output: Reliable
Asynchronous Sync_level=SYNCPT	Input: Reliable Output: Reliable	Input: PWS resend Output: Reliable	Input: Reliable Output: Reliable

Notes:

1. A *session failure* is a network-connectivity breakage.
2. A *non-session failure* is any other kind of failure, such as invalid security authorization.
3. IMS resends asynchronous output if CONFIRM is lost; therefore, the PWS must tolerate duplicate output.

The following table shows the specifics of the processing windows when integrity is compromised (the message is either lost or its state is ambiguous). The table indicates the relative probability of an occurrence of each window and whether output is lost or duplicated.

A Sync_level value of NONE does not apply to asynchronous output, because IMS always uses Sync_level=CONFIRM for such output.

Table 28. Results of processing when integrity is compromised

Conversation attributes	State of window ₁ before accepting transaction	Probability of window state	Possible action while sending response	Probability of action while sending response
Synchronous Sync_level=NONE	ALLOCATE to PREPARE_TO_ RECEIVE return	Medium	Can lose or send duplicate output.	Medium
Synchronous Sync_level=CONFIRM	PREPARE_TO_ RECEIVE to PREPARE_TO_ RECEIVE return	Small	CONFIRM to IMS receipt. Can cause duplicate output.	Small
Synchronous Sync_level=SYNCPT	PREPARE_TO_ RECEIVE to PREPARE_TO_ RECEIVE return	Small	CONFIRM to IMS receipt. Can cause duplicate output.	Small
Asynchronous Sync_level=NONE	Allocate to Deallocate	High	CONFIRMED to IMS receipt. Can cause duplicate output.	Small
Asynchronous Sync_level=CONFIRM	PREPARE_TO_ RECEIVE to PREPARE_TO_ RECEIVE return	Small ²	CONFIRMED to IMS receipt. Can cause duplicate output.	Small
Asynchronous Sync_level=SYNCPT	PREPARE_TO_ RECEIVE to PREPARE_TO_ RECEIVE return	Small ²	CONFIRMED to IMS receipt. Can cause duplicate output.	

Notes:

1. The term *window* refers to a period of time when certain events can occur, such as the consequences described in this table.
2. Can be recoverable.

The following table indicates how IMS recovers APPC transactions across IMS warm starts, XRF takeovers, APPC session failures, and MSC link failures.

Table 29. Recovering APPC messages

Message type	IMS warm start (NRE or ERE)	XRF takeover	APPC (LU 6.2) session fail	MSC LINK failure
Local Recoverable Tran., Non Resp., Non Conversation - APPC Sync. Conv. Mode - APPC Async. Conv. Mode	Discarded (2) Recovered	Discarded (4) Recovered	Discarded (6) Recovered (1)	N/A (9) N/A (9)
Local Recoverable Tran., Conv. or Resp. mode - APPC Sync. Conv. Mode - APPC Async. Conv. Mode	Discarded (2) N/A (8)	Discarded (4) N/A (8)	Discarded (6) N/A (8)	N/A (9) N/A (8,9)
Local Non Recoverable Tran., - APPC Sync. Conv. Mode - APPC Async. Conv. Mode	Discarded (2) Discarded (2)	Discarded (4)	Discarded (6) Recovered (1)	N/A (9) N/A (9)

Table 29. Recovering APPC messages (continued)

Message type	IMS warm start (NRE or ERE)	XRF takeover	APPC (LU 6.2) session fail	MSC LINK failure
Remote Recoverable Tran., Non Resp., Non Conv. - APPC Sync. Conv. Mode - APPC Async. Conv. Mode	Discarded (2,5) Recovered	Discarded (3,5) Recovered	Recovered (1) Recovered (1)	Recovered (7) Recovered (7)
Remote Recoverable Tran., Conv. or Resp. mode - APPC Sync. Conv. Mode - APPC Async. Conv. Mode	Discarded (2,5) N/A (8)	Discarded (3,5) N/A (8)	Recovered (1) N/A (8)	Recovered (7) N/A (8)
Remote Non Recoverable Tran., - APPC Sync. Conv. Mode - APPC Async. Conv. Mode	Discarded (2,5) Discarded (2,5)	Discarded (3,5) Discarded (3,5)	Recovered (1) Recovered (1)	Recovered (7) Recovered (7)

Note:

1. This recovery scenario assumes the message was enqueued before failure; otherwise, the message is discarded.
2. The message is discarded during IMS warm-start processing.
3. The message is discarded when the MSC link is restarted and when the message is taken off the queue (for sending across the link).
4. The message is discarded when the message region is started and when the message is taken off the queue (for processing by the application program).
5. For all remote MSC APPC transactions, if the message has already been sent across the MSC link to the remote system when the failure occurs in the local IMS, the message is processed. After the message is processed by the remote application program and a response message is sent back to the local system, it is enqueued to the DFSASYNC TP name of the LU 6.2 device or program that submitted the original transaction.
6. At sync point, the User Message Control Error exit routine (DFSCMUX0) can prevent the transaction from being aborted and the output message can be rerouted (recovered).
For more information about this exit routine, see *IMS Version 15.3 Exit Routines*.
7. The standard MSC Link recovery protocol recovers all messages that are queued or are in the process of being sent across the MSC link when the link fails.
8. IMS conversational-mode and response-mode transactions cannot be submitted from APPC asynchronous conversation sessions. APPC synchronous conversation-mode must be used.
9. MSC link failures do not affect local transactions.

DFSAPPC message switch

DFSAPPC is an LU 6.2 descriptor that provides an IMS system service.

It allows LU 6.2 application programs to send messages to the following:

- Application programs (transactions)
- IMS-managed local or remote LTERMs (message switches)
- LU name and TP name

Messages sent with the LTERM= option are directed to IMS-managed local or remote LTERMs. Messages sent without the LTERM= option are sent to the appropriate LU 6.2 application or IMS application program.

Because the LTERM can be an LU 6.2 descriptor name, the message is sent to the LU 6.2 application program as if an LU 6.2 device had been explicitly selected.

With DFSAPPC, message delivery is asynchronous. If a message is allocated and the allocate fails, the message is held on the IMS message queue until it can be successfully delivered.

Example: In the LU 6.2 conversation example, an IMS application issues a DFSAPPC message switch to its partner with the LU name FRED and TPN name REPORT. REPI is the user data.

```
DFSAPPC (TPN=REPORT LU=FRED) REP1
```

You can use a 17-byte network-qualified name in the LU= field.

Restriction: LU 6.2 architecture prohibits the use of the ALTRESP PCB on a CHNG call in an LU 6.2 conversation. The LU 6.2 conversation can only be associated with the IOPCB. The application sends a message on the existing LU 6.2 conversation (synchronous) or has IMS create a new conversation (asynchronous) using the IOPCB. Since there is no LTERM associated with an LU 6.2 conversation, only the IOPCB represents the original LU 6.2 conversation.

Related Reading: For more information about DFSAPPC, see *IMS Version 15.3 Communications and Connections*.

Chapter 8. Testing an IMS application program

You should perform a program unit test on your IMS application program to ensure that the program correctly handles its input data, processing, and output data. The amount and type of testing you do depends on the individual program.

Recommendations for testing an IMS program

Before you start testing your program, be aware of your established test procedures.

To start testing, you need the following three items:

- Test JCL.
- A test database. Never test a program using a production database because the program, if faulty, might damage valid data.
- Test input data. The input data that you use need not be current, but it should be valid. You cannot be sure that your output data is valid unless you use valid input data.

The purpose of testing the program is to make sure that the program can correctly handle all the situations that it might encounter. To thoroughly test the program, try to test as many of the paths that the program can take as possible.

Recommendations:

- Test each path in the program by using input data that forces the program to execute each of its branches.
- Be sure that your program tests its error routines. Again, use input data that will force the program to test as many error conditions as possible.
- Test the editing routines your program uses. Give the program as many different data combinations as possible to make sure it correctly edits its input data.

Testing DL/I call sequences (DFSDDLTO) before testing your IMS program

The DL/I test program, DFSDDLTO, is an IMS application program that executes the DL/I calls you specify against any database.

Restriction: DFSDDLTO does not work if you are using a coordinator controller (CCTL).

An advantage of using DFSDDLTO is that you can test the DL/I call sequence you will use prior to coding your program. Testing the DL/I call sequence before you test the program makes debugging easier, because by the time you test the program, you know that the DL/I calls are correct. When you test the program, and it does not execute correctly, you know that the DL/I calls are not part of the problem if you have already tested them using DFSDDLTO.

For each DL/I call that you want to test, you give DFSDDLTO the call and any SSAs that you are using with the call. DFSDDLTO then executes and gives you the results of the call. After each call, DFSDDLTO shows you the contents of the DB PCB mask and the I/O area. This means that for each call, DFSDDLTO checks the access path you have defined for the segment, and the effect of the call. DFSDDLTO is helpful in debugging because it can display IMS application control blocks.

To indicate to DFSDDLTO the call you want executed, you use four types of control statements:

Status statements establish print options for DFSDDLTO's output and select the DB PCB to use for the calls you specify.

Comment statements let you choose whether you want to supply comments.

Call statements indicate to DFSDDLTO the call you want to execute, any SSAs you want used with the call, and how many times you want the call executed.

Compare statements tell DFSDDLTO that you want it to compare its results after executing the call with the results you supply.

In addition to testing call sequences to see if they work, you can also use DFSDDLTO to check the performance of call sequences.

Using BTS to test your IMS program

IMS Batch Terminal Simulator for z/OS (BTS) is a valuable tool for testing programs because you can use it to test call sequences. The documentation that BTS produces is helpful in debugging. You can also test online application programs without actually running them online.

Restriction: BTS does not work if you are using a CCTL or running under DBCTL.

Related reading: For information about how to use BTS, see *IMS Batch Terminal Simulator for z/OS User's Guide*.

Tracing DL/I calls with image capture for your IMS program

The DL/I image capture program (DFSDLTRO) is a trace program that can trace and record DL/I calls issued by all types of IMS application programs.

Restriction: The image capture program does not trace calls to Fast Path databases.

You can run the image capture program in a DB/DC or a batch environment to:

- **Test your program**

If the image capture program detects an error in a call it traces, it reproduces as much of the call as possible, although it cannot document where the error occurred, and cannot always reproduce the full SSA.

- **Produce input for DFSDDLTO**

You can use the output produced by the image capture program as input to DFSDDLTO. The image capture program produces status statements, comment statements, call statements, and compare statements for DFSDDLTO.

- **Debug your program**

When your program terminates abnormally, you can rerun the program using the image capture program, which can then reproduce and document the conditions that led to the program failure. You can use the information in the report produced by the image capture program to find and fix the problem.

Using image capture with DFSDDLTO

The image capture program produces the following control statements that you can use as input to DFSDDLTO.

- **Status statements**

When you invoke the image capture program, it produces the status statement. The status statement it produces:

- Sets print options so that DFSDDLTO prints all call trace comments, all DL/I calls, and the results of all comparisons.
- Determines the new relative PCB number each time a PCB change occurs while the application program is executing.

- **Comments statement**

The image capture program also produces a comments statement when you invoke it. The comments statements give:

- The time and date IMS started the trace

- The name of the PSB being traced

The image capture program also produces a comments statement preceding any call in which IMS finds an error.

- **Call statements**

The image capture program produces a call statement for each DL/I call the application program issues. It also generates a CHKP call when it starts the trace and after each commit point or CHKP request.

- **Compare statements**

The image capture program produces data and PCB comparison statements if you specify COMP on the TRACE command (if you run the image capture program online), or on the DLITRACE control statement (if you run the image capture program as a batch job).

Restrictions on using image capture output

The status statement of the image capture call is based on relative PCB position.

When the PCB parameter LIST=NO has been specified, the status statement may need to be changed to select the PCB as follows:


- If all PCBs have the parameter LIST=YES, the status statement does not need to be changed.
- If all PCBs have the parameter LIST=NO, the status statement needs to be changed from the relative PCB number to the correct PCB name.
- If some PCBs have the parameter LIST=YES and some have the parameter LIST=NO, the status statement needs to be changed as follows:
 - The PCB relative position is based on all PCBs as if LIST=YES.
 - For PCBs that have a PCB name, the status statement can be changed to use the PCB name based on a relative PCB number.
 - For PCBs that have LIST=YES and no PCB name, change the relative PCB number to refer to the relative PCB number in the user list by looking at the PCB list using LIST=YES and LIST=NO.

Running image capture online

When you run the image capture program online, the trace output goes to the IMS log data set. To run the image capture program online, you issue the IMS TRACE command from the IMS master terminal.

If you trace a BMP or an MPP and you want to use the trace results with DFSDDLTO, the BMP or MPP must have exclusive write access to the databases it processes. If the application program does not have exclusive access, the results of DFSDDLTO may differ from the results of the application program. When you trace a BMP that accesses GSAM databases, you must include an //IMSERR DD statement to get a formatted dump of the GSAM control blocks.

The following diagram shows the TRACE command format:



SET ON|OFF

Turns the trace on or off.

PSB *psbname*

Specifies the name of the PSB you want to trace. You can trace more than one PSB at the same time by issuing a separate TRACE command for each PSB.

COMP|NOCOMP

Specifies whether you want the image capture program to produce data and PCB compare statements to be used as input to DFSDDLTO.

Running image capture as a batch job

To run the image capture program as a batch job, you use the DLITRACE control statement in the DFSVSAMP DD data set.

In the DLITRACE control statement, you specify:

- Whether you want to trace all of the DL/I calls the program issues or trace only a certain group of calls.
- Whether you want the trace output to go to:

A sequential data set that you specify

The IMS log data set

Both sequential and IMS log data sets

If the program being traced issues CHKP and XRST calls, the checkpoint and restart information may not be directly reproducible when you use the trace output with DFSDDLTO.

When you run DFSDDLTO in an IMS DL/I or DBB batch region with trace output, the results are the same as the application program's results, but only if the database has not been altered.

For information on the format of the DLITRACE control statement in the DFSVSAMP DD data set, see the topic "Defining DL/I call image trace" in *IMS Version 15.3 System Definition*.

Retrieving image capture data from the log data set

If the trace output is sent to the IMS log data set, you can retrieve it by using utility DFSERA10 and a DL/I call trace exit routine, DFSERA50. DFSERA50 deblocks, formats, and numbers the image capture program records that are to be retrieved.

To use DFSERA50, you must insert a DD statement defining a sequential output data set in the DFSERA10 input stream. The default ddname for this DD statement is TRCPUNCH. The statement must specify BLKSIZE=80.

For example, you can use the following examples of DFSERA10 input control statements in the SYSIN data set to retrieve the image capture program data from the log data set:

- **Print all image capture program records:**

Column 1	Column 10
OPTION	PRINT OFFSET=5,VALUE=5F,FLDTYP=X

- **Print selected image capture program records by PSB name:**

Column 1	Column 10
OPTION	PRINT OFFSET=5,VALUE=5F,COND=M
OPTION	PRINT OFFSET=25,FLDTYP=C,FLDLLEN=8,VALUE=psbname, COND=E

- **Format image capture program records (in a format that can be used as input to DFSDDLTO):**

Column 1	Column 10
OPTION	PRINT OFFSET=5,VALUE=5F,COND=M
OPTION	PRINT EXITR=DFSERA50,OFFSET=25,FLDTYP=C
	VALUE=psbname,FLDLLEN=8,DDNAME=OUTDDN,COND=E

Remember: The DDNAME= parameter names the DD statement to be used by DFSERA50. The data set that is defined on the OUTDDN DD statement is used instead of the default TRCPUNCH DD statement. For this example, the DD is:

```
//OUTDDN DD ...,DCB=(BLKSIZE=80),...
```

Requests for monitoring and debugging your IMS program

You can use the STAT and LOG requests to help you in debugging your program.

- The Statistics (STAT) call retrieves database statistics.
- The Log (LOG) call makes it possible for the application program to write a record on the system log.

The enhanced OSAM and VSAM STAT calls provide additional information for monitoring performance and fine tuning of the system for specific needs.

When the enhanced STAT call is issued, the following information is returned:

- OSAM statistics for each defined subpool
- VSAM statistics that also include hyperspace statistics
- OSAM and VSAM count fields that have been expanded to 10 digits

Retrieving database statistics: the STAT call

The STAT call is helpful in debugging a program because it retrieves IMS database statistics. It is also helpful in monitoring and fine tuning for performance. The STAT call retrieves OSAM database buffer pool statistics and VSAM database buffer subpool statistics.

This topic contains Product-sensitive Programming Interface information.

When you issue the STAT call, you indicate:

- An I/O area into which the statistics are to be returned.
- A statistics function, which is the name of a 9-byte area whose contents describe the type and format of the statistics you want returned. The contents of the area are defined as follows:
 - The first 4 bytes define the type of statistics desired (OSAM or VSAM).
 - The 5th byte defines the format to be returned (formatted, unformatted, or summary).
 - The remaining 4 bytes are defined as follows:
 - The normal or enhanced STAT call contains 4 bytes of blanks.
 - The extended STAT call contains the 4-byte parameter ' E1 ' (a 1-byte blank, followed by a 2-byte character string, and then another 1-byte blank).

Related reference

[STAT call \(Application Programming APIs\)](#)

Format of OSAM buffer pool statistics

For OSAM buffer pool statistics, the values are possible for the stat-function parameter and for the format of the data that is returned to the application program. If no OSAM buffer pool is present, a GE status code is returned to the program.

DBASF

This function value provides the full OSAM database buffer pool statistics in a formatted form. The application program I/O area must be at least 360 bytes. Three 120-byte records (formatted for printing) are provided as two heading lines and one line of statistics. The following diagram shows the data format.

BLOCK REQ nnnnnnnn	FOUND IN POOL nnnnnnnn	READS ISSUED nnnnnn	BUFF ALTS nnnnnnnn	OSAM WRITES nnnnnnnn	BLOCKS WRITTEN nnnnnnnn	NEW BLOCKS nnnnnn	CHAIN WRITES nnnnnn
WRITTEN AS NEW nnnnnnnn	LOGICAL CYL FORMAT nnnnnnnn	PURGE REQ nnnnnnnn	RELEASE REQ nnnnnnnn	ERRORS nn/nn			

BLOCK REQ

Number of block requests received.

FOUND IN POOL

Number of times the block requested was found in the buffer pool.

READS ISSUED

Number of OSAM reads issued.

BUFF ALTS

Number of buffers altered in the pool.

OSAM WRITES

Number of OSAM writes issued.

BLOCKS WRITTEN

Number of blocks written from the pool.

NEW BLOCKS

Number of new blocks created in the pool.

CHAIN WRITES

Number of chained OSAM writes issued.

WRITTEN AS NEW

Number of blocks created.

LOGICAL CYL FORMAT

Number of format logical cylinder requests issued.

PURGE REQ

Number of purge user requests.

RELEASE REQ

Number of release ownership requests.

ERRORS

Number of write error buffers currently in the pool or the largest number of errors in the pool during this execution.

DBASU

This function value provides the full OSAM database buffer pool statistics in an unformatted form. The application program I/O area must be at least 72 bytes. Eighteen fullwords of binary data are provided:

Word**Contents****1**

A count of the number of words that follow.

2-18

The statistic values in the same sequence as presented by the DBASF function value.

DBASS

This function value provides a summary of the OSAM database buffer pool statistics in a formatted form. The application program I/O area must be at least 180 bytes. Three 60-byte records (formatted for printing) are provided. The following diagram shows the data format.

```
DATA BASE BUFFER POOL:  SIZE nnnnnnn
                        REQ1 nnnnnn REQ2 nnnnnn READ nnnnnn WRITES nnnnnn LCYL nnnnn
                        PURG nnnnnn OWNRR nnnnnn ERRORS nn/nn
```

SIZE

Buffer pool size.

REQ1

Number of block requests.

Number of block requests satisfied in the pool plus new blocks created.

Number of read requests issued.

Number of OSAM writes issued.

Number of format logical cylinder requests.

Number of purge user requests.

Number of release ownership requests.

Number of permanent errors now in the pool or the largest number of permanent errors during this execution.

Format of VSAM buffer subpool statistics

Because there might be several buffer subpools for VSAM databases, the STAT call is iterative when requesting these statistics. If more than one VSAM local shared resource pool is defined, statistics are retrieved for all VSAM local shared resource pools in the order in which they are defined. For each local shared resource pool, statistics are retrieved for each subpool according to buffer size.

The first time the call is issued, the statistics for the subpool with the smallest buffer size are provided. For each succeeding call (without intervening use of the PCB), the statistics for the subpool with the next-larger buffer size are provided.

If index subpools exist within the local shared resource pool, the index subpool statistics always follow statistics of the data subpools. Index subpool statistics are also retrieved in ascending order based on the buffer size.

The final call for the series returns a GA status code in the PCB. The statistics returned are totals for all subpools in all local shared resource pools. If no VSAM buffer subpools are present, a GE status code is returned to the program.

VBASF

This function value provides the full VSAM database subpool statistics in a formatted form. The application program I/O area must be at least 360 bytes. Three 120-byte records (formatted for printing) are provided as two heading lines and one line of statistics. Each successive call returns the statistics for the next data subpool. If present, statistics for index subpools follow the statistics for data subpools.

The following diagram shows the data format.

[illegible]

ID of the local shared resource pool.

Size of the buffers in this VSAM subpool. In the final call, this field is set to ALL.

Number of buffers in this subpool. In the final call, this is the number of buffers in all subpools.

RET RBA

Number of retrieve-by-RBA calls received by the buffer handler.

RET KEY

Number of retrieve-by-key calls received by the buffer handler.

ISRT ES

Number of logical records inserted into ESDSs.

ISRT KS

Number of logical records inserted into KSDSs.

BFR ALT

Number of logical records altered in this subpool. Delete calls that result in erasing records from a KSDS are not counted.

BGWRT

Number of times the background-write function was executed by the buffer handler.

SYN PTS

Number of Synchronization calls received by the buffer handler.

GETS

Number of VSAM GET calls issued by the buffer handler.

SCHBFR

Number of VSAM SCHBFR calls issued by the buffer handler.

FOUND

Number of times VSAM found the control interval already in the subpool.

READS

Number of times VSAM read a control interval from external storage.

USR WTS

Number of VSAM writes initiated by IMS.

NUR WTS

Number of VSAM writes initiated to make space in the subpool.

ERRORS

Number of write error buffers currently in the subpool or the largest number of write errors in the subpool during this execution.

VBASU

This function value provides the full VSAM database subpool statistics in a unformatted form. The application program I/O area must be at least 72 bytes. Eighteen fullwords of binary data are provided for each subpool:

Word**Contents****1**

A count of the number of words that follow.

2-18

The statistic values in the same sequence as presented by the VBASF function value, except for POOLID, which is not included in this unformatted form.

VBASS

This function value provides a summary of the VSAM database subpool statistics in a formatted form. The application program I/O area must be at least 180 bytes. Three 60-byte records (formatted for printing) are provided.

The following diagram shows the data format.

```
DATA BASE BUFFER POOL:  BSIZE nnnnnnn  POOLID xxxx  Type x
RRBA nnnnn  RKEY nnnnn  BFALT nnnnn  NREC nnnnn  SYN PTS nnnnn
NMBUFS nnn  VRDS nnnnn  FOUND nnnnn  VWTS nnnnn  ERRORS nn/nn
```

BSIZE

Size of the buffers in this VSAM subpool.

POOLID

ID of the local shared resource pool.

TYPE

Indicates a data (D) subpool or an index (I) subpool.

RRBA

Number of retrieve-by-RBA requests.

RKEY

Number of retrieve-by-key requests.

BFALT

Number of logical records altered.

NREC

Number of new VSAM logical records created.

SYN PTS

Number of sync point requests.

NMBUFS

Number of buffers in this VSAM subpool.

VRDS

Number of VSAM control interval reads.

FOUND

Number of times VSAM found the requested control interval already in the subpool.

VWTS

Number of VSAM control interval writes.

ERRORS

Number of permanent write errors now in the subpool or the largest number of errors in this execution.

Format of enhanced/extended OSAM buffer subpool statistics

The enhanced OSAM buffer pool statistics provide additional information generated for each defined subpool. Because there might be several buffer subpools for OSAM databases, the enhanced STAT call repeatedly requests these statistics. The first time the call is issued, the statistics for the subpool with the smallest buffer size is provided. For each succeeding call (without intervening use of the PCB), the statistics for the subpool with the next-larger buffer size is provided.

The final call for the series returns a GA status code in the PCB. The statistics returned are the totals for all subpools. If no OSAM buffer subpools are present, a GE status code is returned.

Extended OSAM buffer pool statistics can be retrieved by including the 4-byte parameter 'bE1b' following the enhanced call function. The extended STAT call returns all of the statistics returned with the enhanced call, plus the statistics on the coupling facility buffer invalidates, OSAM caching, and sequential buffering IMMED/SYNC read counts.

Restriction: The extended format parameter is supported by the DBESO, DBESU, and DBESF functions only.

DBESF

This function value provides the full OSAM subpool statistics in a formatted form. The application program I/O area must be at least 600 characters. For OSAM subpools, five 120-byte records (formatted for printing) are provided. Three of the records are heading lines and two of the records are lines of subpool statistics.

The following example shows the enhanced stat call format:

B U F F E R H A N D L E R O S A M S T A T I S T I C S										
BSIZ	NBUFS	LOCATE-REQ	NEW-BLOCKS	ALTER- REQ	PURGE- REQ	FND-IN-POOL	BUFRS-SRCH	READ- REQ	REQS	BUFSTL-WRT
nn1K	nnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn

The following example shows the extended stat call format:

B U F F E R H A N D L E R O S A M S T A T I S T I C S										
BSIZ	NBUFS	LOCATE-REQ	NEW-BLOCKS	ALTER- REQ	PURGE- REQ	STG CLS=	FND-IN-POOL	BUFRS-SRCH	READ- REQ	REQS
nn1K	nnnnnnnn5	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0
		nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0	nnnnnnnnnn0
		CF-READS	EXPCTD-NF	CFWRT-PRI	CFWRT-CHG	STGCLS-FULL	XI-CNT	VECTR-XI	SB-SEQRD	SB-ANTICIP
		nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn	nnnnnnnnnn

FIXOPT

Fixed options for this subpool. Y or N indicates whether the data buffer prefix and data buffers are fixed.

POOLID

ID of the local shared resource pool.

BSIZ

Size of the buffers in this subpool. Set to ALL for total line. For the summary totals (BSIZ=ALL), the FIXOPT and POOLID fields are replaced by an OSM= field. This field is the total size of the OSAM subpool.

NBUFS

Number of buffers in this subpool. This is the total number of buffers in the pool for the ALL line.

LOCATE-REQ

Number of LOCATE-type calls.

NEW-BLOCKS

Number of requests to create new blocks.

ALTER-REQ

Number of buffer alter calls. This count includes NEW BLOCK and BYTALT calls.

PURGE-REQ

Number of PURGE calls.

FND-IN-POOL

Number of LOCATE-type calls for this subpool where data is already in the OSAM pool.

BUFRS-SRCH

Number of buffers searched by all LOCATE-type calls.

READ-REQS

Number of READ I/O requests.

BUFSTL-WRT

Number of single block writes initiated by buffer steal routine.

PURGE-WRTS

Number of blocks for this subpool written by purge.

WT-BUSY-ID

Number of LOCATE calls that waited due to busy ID.

WT-BUSY-WR

Number of LOCATE calls that waited due to buffer busy writing.

WT-BUSY-RD

Number of LOCATE calls that waited due to buffer busy reading.

WT-RLSEOWN

Number of buffer steal or purge requests that waited for ownership to be released.

WT-NO-BFRS

Number of buffer steal requests that waited because no buffers are available to be stolen.

ERRORS

Total number of I/O errors for this subpool or the number of buffers locked in pool due to write errors.

CF-READS

Number of blocks read from CF.

EXPCTD-NF

Number of blocks expected but not read.

CFWRT-PRI

Number of blocks written to CF (prime).

CFWRT-CHG

Number of blocks written to CF (changed).

STGCLS-FULL

Number of blocks not written (STG CLS full).

XI-CNTL

Number of XI buffer invalidate calls.

VECTR-XI

Number of buffers found invalidated by XI on VECTOR call.

SB-SEQRD

Number of immediate (SYNC) sequential reads (SB stat).

SB-ANTICIP

Number of anticipatory reads (SB stat).

DBESU

This function value provides full OSAM statistics in an unformatted form. The application program I/O area must be at least 84 bytes. Twenty-one fullwords of binary data are provided for each subpool:

Word**Contents****1**

A count of the number of words that follow.

2-19

The statistics provided in the same sequence as presented by the DBESF function value.

20

The POOLID provided at subpool definition time.

21

The second byte contains the following fix options for this subpool:

- X'04' = DATA BUFFER PREFIX fixed
- X'02' = DATA BUFFERS fixed

The summary totals (word 2=ALL), for word 21, contain the total size of the OSAM pool.

22-30

Extended stat data in same sequence as on DBESF call.

DBESS

This function value provides a summary of the OSAM database buffer pool statistics in a formatted form. The application program I/O area must be at least 360 bytes. Six 60-byte records (formatted for printing) are provided. This STAT call is a restructured DBASF STAT call that allows for 10-digit count fields. In addition, the subpool header blocks give a total of the number of OSAM buffers in the pool.

The following shows the data format:

```
DATA BASE BUFFER POOL:  NSUBPL 00000000  NBUFS 0000000000
BLKREQ 000000000000  INPOOL 000000000000  READS 000000000000
BUFALT 000000000000  WRITES 000000000000  BLKWRT 000000000000
NEWBLK 000000000000  CHNWRT 000000000000  WRTNEW 000000000000
LCYLFM 000000000000  PURGRQ 000000000000  RLSERQ 000000000000
FRCWRT 000000000000  ERRORS 000000000/000000000
```

NSUBPL

Number of subpools defined for the OSAM buffer pool.

NBUFS

Total number of buffers defined in the OSAM buffer pool.

BLKREQ

Number of block requests received.

INPOOL

Number of times the block requested is found in the buffer pool.

READS

Number of OSAM reads issued.

BUFALT

Number of buffers altered in the pool.

WRITES

Number of OSAM writes issued.

BLKWRT

Number of blocks written from the pool.

NEWBLK

Number of blocks created in the pool.

CHNWRT

Number of chained OSAM writes issued.

WRTNEW

Number of blocks created.

LCYLFM

Number of format logical cylinder requests issued.

PURGRQ

Number of purge user requests.

RLSERQ

Number of release ownership requests.

FRCWRT

Number of forced write calls.

ERRORS

Number of write error buffers currently in the pool or the largest number of errors in the pool during this execution.

DBESO

This function value provides the full OSAM database subpool statistics in a formatted form for online statistics that are returned as a result of a /DIS POOL command. This call can also be a user-application

STAT call. When issued as an application DL/I STAT call, the program I/O area must be at least 360 bytes. Six 60-byte records (formatted for printing) are provided.

Example: The following shows the enhanced stat call format:

```
OSAM DB BUFFER POOL:ID xxxx BSIZE nnnnnK NBUFnnnnnnn FX=X/X
LCTREQ nnnnnnnnnn NEWBLK nnnnnnnnnn ALTREQ nnnnnnnnnn
PURGRQ nnnnnnnnnn FNDIPL nnnnnnnnnn BFSRCH nnnnnnnnnn
RDREQ nnnnnnnnnn BFSTLW nnnnnnnnnn PURGWR nnnnnnnnnn
WBSYID nnnnnnnnnn WBSYWR nnnnnnnnnn WBSYRD nnnnnnnnnn
WRLSEO nnnnnnnnnn WNOBFR nnnnnnnnnn ERRORS nnnnn/nnnnn
```

Example: The following shows the extended stat call format:

```
OSAM DB BUFFER POOL:ID xxxx BSIZE nnnnnK NBUFnnnnnnn FX=X/X
LCTREQ nnnnnnnnnn NEWBLK nnnnnnnnnn ALTREQ nnnnnnnnnn
PURGRQ nnnnnnnnnn FNDIPL nnnnnnnnnn BFSRCH nnnnnnnnnn
RDREQ nnnnnnnnnn BFSTLW nnnnnnnnnn PURGWR nnnnnnnnnn
WBSYID nnnnnnnnnn WBSYWR nnnnnnnnnn WBSYRD nnnnnnnnnn
WRLSEO nnnnnnnnnn WNOBFR nnnnnnnnnn ERRORS nnnnn/nnnnn
CFREAD nnnnnnnnnn CFEXPC nnnnnnnnnn CFWRPR nnnnn/nnnnn
CFWRCH nnnnnnnnnn STGCLF nnnnnnnnnn XIINV nnnnn/nnnnn
XICLCT nnnnnnnnnn SBSEQR nnnnnnnnnn SBANTR nnnnn/nnnnn
```

POOLID

ID of the local shared resource pool.

BSIZE

Size of the buffers in this subpool. Set to ALL for summary total line. For the summary totals (BSIZE=ALL), the FX= field is replaced by the OSAM= field. This field is the total size of the OSAM buffer pool. The POOLID is not shown. For the summary totals (BSIZE=ALL), the FX= field is replaced by the OSAM= field. This field is the total size of the OSAM buffer pool. The POOLID is not shown.

NBUF

Number of buffers in this subpool. Total number of buffers in the pool for the ALL line.

FX=

Fixed options for this subpool. Y or N indicates whether the data buffer prefix and data buffers are fixed.

LCTREQ

Number of LOCATE-type calls.

NEWBLK

Number of requests to create new blocks.

ALTREQ

Number of buffer alter calls. This count includes NEW BLOCK and BYTALT calls.

PURGRQ

Number of PURGE calls.

FNDIPL

Number of LOCATE-type calls for this subpool where data is already in the OSAM pool.

BFSRCH

Number of buffers searched by all LOCATE-type calls.

RDREQ

Number of READ I/O requests.

BFSTLW

Number of single-block writes initiated by buffer-steal routine.

PURGWR

Number of buffers written by purge.

WBSYID

Number of LOCATE calls that waited due to busy ID.

WBSYWR

Number of LOCATE calls that waited due to buffer busy writing.

WBSYRD

Number of LOCATE calls that waited due to buffer busy reading.

WRLSEO

Number of buffer steal or purge requests that waited for ownership to be released.

WNOBRF

Number of buffer steal requests that waited because no buffers are available to be stolen.

ERRORS

Total number of I/O errors for this subpool or the number of buffers locked in pool due to write errors.

CFREAD

Number of blocks read from CF.

CFEXPC

Number of blocks expected but not read.

CFWRPR

Number of blocks written to CF (prime).

CFWRCH

Number of blocks written to CF (changed).

STGCLF

Number of blocks not written (STG CLS full).

XIINV

Number of XI buffer invalidate calls.

XICLCT

Number of buffers found invalidated by XI on VECTOR call.

SBSEQR

Number of immediate (SYNC) sequential reads (SB stat).

SBANTR

Number of anticipatory reads (SB stat).

Format of enhanced VSAM buffer subpool statistics

The enhanced VSAM buffer subpool statistics provide information on the total size of VSAM subpools in virtual storage and in hiperspace. All count fields are 10 digits.

Because there might be several buffer subpools for VSAM databases, the enhanced STAT call repeatedly requests these statistics. If more than one VSAM local shared resource pool is defined, statistics are retrieved for all VSAM local shared resource pools in the order in which they are defined. For each local shared resource pool, statistics are retrieved for each subpool according to buffer size.

The first time the call is issued, the statistics for the subpool with the smallest buffer size are provided. For each succeeding call (without intervening use of the PCB), the statistics for the subpool with the next-larger buffer size are provided.

If index subpools exist within the local shared resource pool, the index subpool statistics always follow the data subpools statistics. Index subpool statistics are also retrieved in ascending order based on the buffer size.

The final call for the series returns a GA status code in the PCB. The statistics returned are totals for all subpools in all local shared resource pools. If no VSAM buffer subpools are present, a GE status code is returned to the program.

VBESF

This function value provides the full VSAM database subpool statistics in a formatted form. The application program I/O area must be at least 600 bytes. For each shared resource pool ID, the first call returns five 120-byte records (formatted for printing). Three of the records are heading lines and two of the records are lines of subpool statistics.


```

    BUFFER HANDLER STATISTICS / VSAM STATISTICS FIXOPT=X/X/X POOLID: xxxx
BSIZ NBUFFRS HS-NBUF RETURN-RBA RETURN-KEY ESDS-INSRT KSDS-INSRT BUFFRS-ALT BKGRND-WRT SYNC-POINT ERRORS
      VSAM-GETS SCHED-BUFR VSAM-FOUND VSAM-READS USER-WRITS VSAM-WRITS HSRDS-SUCC HSWRT-SUCC HSR/W-FAIL
nn1K nnnnnn nnnnnnnn nnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn
      nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn nnnnnnnnnnnn

```

Fixed options for this subpool. Y or N indicates whether the data buffer prefix, the index buffers, and the data buffers are fixed.

Size of the buffers in this subpool. Set to ALL for total line. For the summary totals (BSIZ=ALL), the FIXOPT and POOLID fields are replaced by a VS= field and a HS= field. The VS= field is the total size of the VSAM subpool in virtual storage. The HS= field is the total size of the VSAM subpool in hiperspace.

Number of successful VSAM reads from hiperspace buffers.

HSWRT-SUCC

Number of successful VSAM writes from hiperspace buffers.

HSR/W-FAIL

Number of failed VSAM reads from hiperspace buffers/number of failed VSAM writes to hiperspace buffers. This indicates the number of times a VSAM READ/WRITE request from or to hiperspace resulted in DASD I/O.

VBESU

This function value provides full VSAM statistics in an unformatted form. The application program I/O area must be at least 104 bytes. Twenty-five fullwords of binary data are provided for each subpool.

Word

Contents

1

A count of the number of words that follow.

2-23

The statistics provided in the same sequence as presented by the VBESF function value.

24

The POOLID provided at the time the subpool is defined.

25

The first byte contains the subpool type, and the third byte contains the following fixed options for this subpool:

- X'08' = INDEX BUFFERS fixed
- X'04' = DATA BUFFER PREFIX fixed
- X'02' = DATA BUFFERS fixed

The summary totals (word 2=ALL) for word 25 and word 26 contain the virtual and hiperspace pool sizes.

VBESS

This function value provides a summary of the VSAM database subpool statistics in a formatted form. The application program I/O area must be at least 360 bytes. For each shared resource pool ID, the first call provides six 60-byte records (formatted for printing).

The following shows the data format:

VSAM DB	BUFFER POOL: ID	xxxx	BSIZE	nnnnnnK	TYPE	x	FX=X/X/X
RRBA	nnnnnnnnnn	RKEY	nnnnnnnnnn	BFALT	nnnnnnnnnn		
NREC	nnnnnnnnnn	SYNC PT	nnnnnnnnnn	NBUFS	nnnnnnnnnn		
VRDS	nnnnnnnnnn	FOUND	nnnnnnnnnn	VWTS	nnnnnnnnnn		
HSR-S	nnnnnnnnnn	HSW-S	nnnnnnnnnn	HS NBUFS	nnnnnnnnnn		
HS-R/W-FAIL	nnnnn/nnnnn	ERRORS	nnnnnn/nnnnnn				

POOLID

ID of the local shared resource pool.

BSIZE

Size of the buffers in this VSAM subpool.

TYPE

Indicates a data (D) subpool or an index (I) subpool.

FX

Fixed options for this subpool. Y or N indicates whether the data buffer prefix, the index buffers, and the data buffers are fixed.

RRBA

Number of retrieve-by-RBA calls received by the buffer handler.

RKEY

Number of retrieve-by-key calls received by the buffer handler.

BFALT

Number of logical records altered.

NREC

Number of new VSAM logical records created.

SYNC PT

Number of sync point requests.

NBUFS

Number of buffers in this VSAM subpool.

VRDS

Number of VSAM control interval reads.

FOUND

Number of times VSAM found the requested control interval already in the subpool.

VWTS

Number of VSAM control interval writes.

HSR-S

Number of successful VSAM reads from hiperspace buffers.

HSW-S

Number of successful VSAM writes to hiperspace buffers.

HS NBUFS

Number of VSAM hiperspace buffers defined for this subpool.

HS-R/W-FAIL

Number of failed VSAM reads from hiperspace buffers and number of failed VSAM writes to hiperspace buffers. This indicates the number of times a VSAM READ/WRITE request to or from hiperspace resulted in DASD I/O.

ERRORS

Number of permanent write errors now in the subpool or the largest number of errors in this execution.

Writing Information to the system log: the LOG request

An application program can write a record to the system log by issuing the LOG call.

When you issue the LOG request, you specify the I/O area that contains the record you want written to the system log. You can write any information to the log that you want, and you can use different log codes to distinguish between different types of information.

Related Reading: For information about coding the LOG request, see the appropriate application programming reference information.

What to do when your IMS program terminates abnormally

When your program terminates abnormally, you can take the following actions to simplify the task of finding and fixing the problem.

- Record as much information as possible about the circumstances under which the program terminated abnormally.
- Check for certain initialization and execution errors.

Recommended actions after an abnormal termination of an IMS program

Many places have guidelines on what you should do if your program terminates abnormally. The suggestions given here are common guidelines:

- Document the error situation to help in investigating and correcting it. The following information can be helpful:
 - The program's PSB name
 - The transaction code that the program was processing (online programs only)
 - The text of the input message being processed (online programs only)
 - The call function
 - The name of the originating logical terminal (online programs only)
 - The contents of the PCB that was referenced in the call that was executing
 - The contents of the I/O area when the problem occurred
 - If a database call was executing, the SSAs, if any, that the call used
 - The date and time of day
- When your program encounters an error, it can pass all the required error information to a standard error routine. You should not use STAE or ESTAE routines in your program; IMS uses STAE or ESTAE routines to notify the control region of any abnormal termination of the application program. If you call your own STAE or ESTAE routines, IMS may not get control if an abnormal termination occurs.
- Online programs might want to send a message to the originating logical terminal to inform the person at the terminal that an error has occurred. Unless you are using a CCTL, your program can get the logical terminal name from the I/O PCB, place it in an express PCB, and issue one or more ISRT calls to send the message.
- An online program might also want to send a message to the master terminal operator giving information about the program's termination. To do this, the program places the logical terminal name of the master terminal in an express PCB and issues one or more ISRT calls. (This is not applicable if you are using a CCTL.)
- You might also want to send a message to a printer so that you will have a hard-copy record of the error.
- You can send a message to the system log by issuing a LOG request.
- Some places run a BMP at the end of the day to list all the errors that have occurred during the day. If your shop does this, you can send a message using an express PCB that has its destination set for that BMP. (This is not applicable if you are using a CCTL.)

Diagnosing an abnormal termination of an IMS program

If your program does not run correctly when you are testing it or when it is executing, you need to isolate the problem. The problem might be anything from a programming error (for example, an error in the way you coded one of your requests) to a system problem. This section gives some guidelines about the steps that you, as the application programmer, can take when your program fails to run, terminates abnormally, or gives incorrect results.

IMS program initialization errors

Before your program receives control, IMS must have correctly loaded and initialized the PSB and DBDs used by your application program. Often, when the problem is in this area, you need a system programmer or DBA (or your equivalent specialist) to fix the problem. One thing you can do is to find out if there have been any recent changes to the DBDs, PSB, and the control blocks that they generate.

IMS program execution errors

If you do not have any initialization errors, check:

1. The output from the compiler. Make sure that all error messages have been resolved.
2. The output from the binder:
 - Are all external references resolved?
 - Have all necessary modules been included?

- Was the language interface module correctly included?
- Is the correct entry point specified?

3. Your JCL:

- Is the information that described the files that contain the databases correct? If not, check with your DBA.
- Have you included the DL/I parameter statement in the correct format?
- Have you included the region size parameter in the EXEC statement? Does it specify a region or partition large enough for the storage required for IMS and your program?
- Have you declared the fields in the PCB masks correctly?
- If your program is an assembler language program, have you saved and restored registers correctly? Did you save the list of PCB addresses at entry? Does register 1 point to a parameter list of fullwords before issuing any DL/I calls?
- For COBOL for z/OS and PL/I for MVS and VM, are the literals you are using for arguments in DL/I calls producing the results you expect? For example, in PL/I for MVS and VM, is the parameter count being generated as a half-word instead of a fullword, and is the function code producing the required 4-byte field?
- Use the PCB as much as possible to determine what in your program is producing incorrect results.

Related concepts

[“Use of STAE or ESTAE and SPIE in IMS programs” on page 53](#)

IMS uses STAE or ESTAE routines in the control region, the dependent (MPP, IFP, BMP) regions, and the batch regions. In the control region, STAE or ESTAE routines ensure that database logging and various resource cleanup functions are complete.

Chapter 9. Testing a CICS application program

You should perform a program unit test on your CICS application program to ensure that the program correctly handles its input data, processing, and output data. The amount and type of testing you do depends on the individual program.

Recommendations for testing a CICS program

When you are ready to test your program, be aware of your established test procedures before you start.

To start testing, you need the following three items:

- Test JCL.
- A test database. When you are testing a program, do not execute it against a production database because the program, if faulty, might damage valid data.
- Test input data. The input data that you use need not be current, but it should be valid data. You cannot be sure that your output data is valid unless you use valid input data.

The purpose of testing the program is to make sure that the program can correctly handle all the situations that it might encounter.

To thoroughly test the program, try to test as many of the paths that the program can take as possible. For example:

- Test each path in the program by using input data that forces the program to execute each of its branches.
- Be sure that your program tests its error routines. Again, use input data that will force the program to test as many error conditions as possible.
- Test the editing routines your program uses. Give the program as many different data combinations as possible to make sure it correctly edits its input data.

Testing your CICS program

You can use different tools to test a CICS program, depending on the type of program.

The following table summarizes the tools that are available for online DBCTL, batch, and BMP programs.

Table 30. Tools you can use for testing your program

Tool	Online (DBCTL)	Batch	BMP
Execution Diagnostic Facility (EDF)	Yes ¹	No	No
CICS dump control	Yes	No	No
CICS trace control	Yes	Yes	No
DFSDDLTO	No	Yes ²	Yes ²
DL/I image capture program	Yes	Yes	Yes

Notes:

1. For online, command-level programs only.
2. For call-level programs only. (For a command-level batch program, you can use DL/I image capture program first, to produce calls for DFSDDLTO.)

Using the Execution Diagnostic Facility (command-level only)

You can use the Execution Diagnostic Facility (EDF) to test command-level programs online. EDF can display EXEC CICS and EXEC DLI commands in online programs; it cannot intercept DL/I calls.

With EDF you can:

- Display and modify working storage; you can change values in the DIB.
- Display and modify a command before it is executed. You can modify the value of any argument, and then execute the command.
- Modify the return codes after the execution of the command. After the command has been executed, but before control is returned to the application program, the command is intercepted to show the response and any argument values set by CICS.

You can run EDF on the same terminal as the program you are testing.

Related Reading: For more information about using EDF, see "Execution (Command-Level) Diagnostic Facility" in CICS Transaction Server for z/OS CICS Application Programming Reference.

Using CICS dump control

You can use the CICS dump control facility to dump virtual storage areas, CICS tables, and task-related storage areas. For more information about using the CICS dump control facility, see the CICS application programming reference manual that applies to your version of CICS.

Using CICS trace control

You can use the trace control facility to help debug and monitor your online programs in the DBCTL environment. You can use trace control requests to record entries in a trace table. The trace table can be located either in virtual storage or on auxiliary storage. If it is in virtual storage, you can gain access to it by investigating a dump; if it is on auxiliary storage, you can print the trace table. For more information about the control statements you can use to produce trace entries, see the information about trace control in the application programming reference manual that applies to your version of CICS.

Tracing DL/I calls with image capture

DL/I image capture program (DFSDLTRO) is a trace program that can trace and record DL/I calls issued by batch, BMP, and online (DBCTL environment) programs. You can also use the image capture program with command-level programs, and you can produce calls for use as input to DFSDDLTO.

You can use the image capture program to:

- **Test your program**

If the image capture program detects an error in a call it traces, it reproduces as much of the call as possible, although it cannot document where the error occurred, and cannot always reproduce the full SSA.

- **Produce input for DFSDDLTO (DL/I test program)**

You can use the output produced by the image capture program as input to DFSDDLTO. The image capture program produces status statements, comment statements, call statements, and compare statements for DFSDDLTO. For example, you can use the image capture program with a command-level program, to produce calls for DFSDDLTO.

- **Debug your program**

When your program terminates abnormally, you can rerun the program using the image capture program. The image capture program can then reproduce and document the conditions that led to the program failure. You can use the information in the report produced by the image capture program to find and fix the problem.

Using image capture with DFSDDLTO

The image capture program produces the following control statements that you can use as input to DFSDDLTO:

- **Status statements**

When you invoke the image capture program, it produces the status statement. The status statement it produces:

- Sets print options so that DFSDDLTO prints all call trace comments, all DL/I calls, and the results of all comparisons.
- Determines the new relative PCB number each time a PCB change occurs while the application program is executing.

- **Comments statement**

The image capture program also produces a comments statement when you invoke it. The comments statements give:

- The time and date IMS started the trace
- The name of the PSB being traced

The image capture program also produces a comments statement preceding any call in which IMS finds an error.

- **Call statements**

The image capture program produces a call statement for each DL/I call or EXEC DLI command the application program issues. It also generates a CHKP call when it starts the trace and after each commit point or CHKP request.

- **Compare statements**


If you specify COMP on the DLITRACE control statement, the image capture program produces data and PCB comparison statements.

Running image capture online

When you run the image capture program online, the trace output goes to the IMS log data set. To run the image capture program online, you issue the IMS TRACE command from the z/OS console.

If you trace a BMP and you want to use the trace results with DFSDDLTO, the BMP must have exclusive write access to the databases it processes. If the application program does not have exclusive access, the results of DFSDDLTO may differ from the results of the application program.

The following diagram shows TRACE command format:



SET ON|OFF

Turns the trace on or off.

PSB psbname

Specifies the name of the PSB you want to trace. You can trace more than one PSB at the same time, by issuing a separate TRACE command for each PSB.

COMP|NOCOMP

Specifies whether you want the image capture program to produce data and PCB compare statements to be used with DFSDDLTO.

Running image capture as a batch job

To run the image capture program as a batch job, you use the DLITRACE control statement in the DFSVSAMP DD data set.

In the DLITRACE control statement, you specify:

- Whether you want to trace all of the DL/I calls the program issues or trace only a certain group of calls.
- Whether you want the trace output to go to:

- A sequential data set that you specify
- The IMS log data set
- Both sequential and IMS log data sets

If the program being traced issues CHKP and XRST calls, the checkpoint and restart information may not be directly reproducible when you use the trace output with DFSDDLTO.

When you run DFSDDLTO in an IMS DL/I or DBB batch region with trace output, the results are the same as the application program's results, but only if the database has not been altered.

For information about the format of the DLITRACE control statement in the DFSVSAMP DD data set, see [Defining DL/I call image trace \(System Definition\)](#).

Example of DLITRACE

This example shows a DLITRACE control statement that traces the first 14 DL/I calls or commands that the program issues, sends the output to the IMS log data set, and produces data and PCB comparison statements for DFSDDLTO.

```
//DFSVSAMP DD *  
DLITRACE LOG=YES,STOP=14,COMP  
/*
```

Special JCL requirements

The following are special JCL requirements:

//IEFRDER DD

If you want log data set output, this DD statement is required to define the IMS log data set.

//DFSTROUT DD|anyname

If you want sequential data set output, this DD statement is required to define that data set. If you want to specify an alternate DDNAME (anyname), it must be specified using the DDNAME parameter on the DLITRACE control statement.

The DCB parameters on the JCL statement are not required. The data set characteristics are:

- RECFM=FB
- LRECL=80

Notes on using image capture

- If the program being traced issues CHKP and XRST calls, the checkpoint and restart information may not be directly reproducible when you use the trace output with the DFSDDLTO.
- When you run DFSDDLTO in an IMS DL/I or DBB batch region with trace output, the results are the same as the results of the application program provided the database was not altered.

Retrieving image capture data from the log data set

If the trace output is sent to the IMS log data set, you can retrieve it by using utility DFSERA10 and a DL/I call trace exit routine, DFSERA50. DFSERA50 deblocks, formats, and numbers the image capture program records to be retrieved. To use DFSERA50, you must insert a DD statement defining a sequential output data set in the DFSERA10 input stream. The default ddname for this DD statement is TRCPUNCH. The

card must specify BLKSIZE=A, where A can be any multiple of 80 ranging from 80 to 32720. For example, 27920.

For example, you can use the following examples of DFSERA10 input control statements in the SYSIN data set to retrieve the image capture program data from the log data set:

- **Print all image capture program records:**

Column 1 OPTION	Column 10 PRINT OFFSET=5,VALUE=5F,FLDTYP=X

- **Print selected image capture program records by PSB name:**

Column 1 OPTION	Column 10 PRINT OFFSET=5,VALUE=5F,COND=M
OPTION	PRINT OFFSET=25,FLDTYP=C,FLDLEN=8, VALUE=psbname,COND=E

- **Format image capture program records (in a format that can be used as input to DFSDDLTO):**

Column 1 OPTION	Column 10 PRINT OFFSET=5,VALUE=5F,COND=M
OPTION	PRINT EXITR=DFSERA50,OFFSET=25,FLDTYP=C VALUE=psbname,FLDLEN=8,DDNAME=OUTDDN,COND=E

The DDNAME= parameter is used to name the DD statement used by DFSERA50. The data set defined on the OUTDDN DD statement is used instead of the default TRCPUNCH DD statement. For this example, the DD appears as:

```
//OUTDDN DD ...,DCB=(BLKSIZE=27920),...
```

Requests for monitoring and debugging your CICS program

You can use the STAT and LOG requests to help you in debugging your program.

- The statistics (STAT) request retrieves database statistics. STAT can be issued from both call- and command-level programs.
- The log (LOG) request makes it possible for the application program to write a record on the system log. You can issue LOG as a command or call in a batch program; in this case, the record is written to the IMS log. You can issue LOG as a call or command in an online program in the DBCTL environment; in this case, the record is written to the DBCTL log.

What to do when your CICS program terminates abnormally

Whenever your program terminates abnormally, you can take some actions to simplify the task of finding and fixing the problem.

First, you can record as much information as possible about the circumstances under which the program terminated abnormally; and second, you can check for certain initialization and execution errors.

Recommended actions after an abnormal termination of CICS

Many places have guidelines on what you should do if your program terminates abnormally. The suggestions given here are some common guidelines:

- Document the error situation to help in investigating and correcting it. Some of the information that can be helpful is:
 - The program's PSB name
 - The transaction code that the program was processing (online programs only)
 - The text of the input screen being processed (online programs only)
 - The call function

- The terminal ID (online programs only)
- The contents of the PCB or the DIB
- The contents of the I/O area when the problem occurred
- If a database request was executing, the SSAs or SEGMENT and WHERE options, if any, the request used
- The date and time of day
- When your program encounters an error, it can pass all the required error information to a standard error routine.
- An online program might also want to send a message to the master terminal destination (CSMT) and application terminal operator, giving information about the program's termination.
- You can send a message to the system log by issuing a LOG request.

Diagnosing an abnormal termination of CICS

If your program does not run correctly when you are testing it or when it is executing, you need to isolate the problem. The problem might be anything from a programming error (for example, an error in the way you coded one of your requests) to a system problem. This section gives some guidelines about the steps that you, as the application programmer, can take when your program fails to run, terminates abnormally, or gives incorrect results.

CICS initialization errors

Before your program receives control, IMS must have correctly loaded and initialized the PSB and DBDs used by your application program. Often, when the problem is in this area, you need a system programmer or DBA (or your equivalent specialist) to fix the problem. One thing you can do is to find out if there have been any recent changes to the DBDs, PSB, and the control blocks that they generate.

CICS execution errors

If you do not have any initialization errors, check the following in your program:

1. The output from the compiler. Make sure that all error messages have been resolved.
2. The output from the binder:
 - Are all external references resolved?
 - Have all necessary modules been included?
 - Was the language interface module correctly included?
 - Is the correct entry point specified (for batch programs only)?
3. Your JCL:
 - Is the information that described the files that contain the databases correct? If not, check with your DBA.
 - Have you included the DL/I parameter statement in the correct format (for batch programs only)?
 - Have you included the region size parameter in the EXEC statement? Does it specify a region or partition large enough for the storage required for IMS and your program (for batch programs only)?
4. Your call-level program:
 - Have you declared the fields in the PCB masks correctly?
 - If your program is an assembler language program, have you saved and restored registers correctly? Did you save the list of PCB addresses at entry? Does register 1 point to a parameter list of full words before issuing any DL/I calls?
 - For COBOL for z/OS and PL/I for MVS and VM, are the literals you are using for arguments in DL/I calls producing the results you expect? For example, in PL/I for MVS and VM, is the parameter count

being generated as a half word instead of a fullword, and is the function code producing the required 4-byte field?

- Use the PCB as much as possible to determine what in your program is producing incorrect results.

5. Your command-level program:

- Did you use the FROM option with your ISRT or REPL command? If not, data will not be transferred to the database.
- Check translator messages for errors.

Chapter 10. Documenting your application program

Many places establish standards for program documentation; make sure you are aware of your established standards.

Documentation for other programmers

Documenting a program is not something you do at the end of the project; your documentation will be much more complete, and more useful to others, if you record information about the program as you structure and code it. Include any information that might be useful to someone else who must work with your program.

The reason you record this information is so that people who maintain your program know why you chose certain commands, options, call structures, and command codes. For example, if the DBA were considering reorganizing the database in some way, information about why your program accesses the data the way it does would be helpful.

Information you can include for other programmers includes:

- Flowcharts and pseudocode for the program
- Comments about the program from code inspections
- A written description of the program flow
- Information about why you chose the call sequence you did, such as:
 - Did you test the call sequence using DFSDDLTO?
 - In cases where more than one combination of calls would have had the same results, why did you choose the sequence you did?
 - What was the other sequence? Did you test it using DFSDDLTO?
- Any problems you encountered in structuring or coding the program
- Any problems you had when you tested the program
- Warnings about what should not be changed in the program

All this information relates to structuring and coding the program. In addition, you should include the documentation for end users with the documentation for programmers.

Ultimately, you must determine the level of detail necessary and the most suitable form for documenting the program. These documentation guidelines are provided as suggestions.

Documentation for end users

In addition to documenting the design of the application, you should record information about how the program is used.

The amount of information that users need and how much of it you should supply depends upon whom the users of the program are and what type of program it is.

At a minimum, include the following information for those who use your program:

- What one needs in order to use the program, for example:
 - For online programs, is there a password?
 - For batch programs, what is the required JCL?
- The input that one needs to supply to the program, for example:
 - For an MPP, what is the MOD name that must be entered to initially format the screen?
 - For a CICS online program, what is the CICS transaction code that must be entered? What terminal input is expected?

- For a batch program, is the input in the form of a tape, or a disk data set? Is the input originally output from a previous job?
- The content and form of the program's output, for example:
 - If it is a report, show the format or include a sample listing.
 - For an online application program, show what the screen will look like.
- For online programs, if decisions must be made, explain what is involved in each decision. Present the choices and the defaults.

If the people that will be using your program are unfamiliar with terminals, they will need a user's guide also. This guide should give explicit instructions on how to use the terminal and what a user can expect from the program. The guide should contain discussions of what should be done if the task or program abends, whether the program should be restarted, or if the database requires recovery. Although you may not be responsible for providing this kind of information, you should provide any information that is unique to your application to whomever is responsible for this kind of information.

Part 2. Application programming for IMS DB

IMS provides support for writing application programs to access the IMS database.

Chapter 11. Writing your application programs for IMS DB

You can write application programs in High Level Assembler language, C language, COBOL, Java, Pascal, and PL/I to access data in the IMS DB.

Related concepts

[“IMS solutions for Java development overview” on page 553](#)

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

Programming guidelines

The number, type, and sequence of the IMS requests your program issues affects the efficiency of your program. A program that is poorly designed can still run if it is coded correctly. IMS will not find design errors for you. The suggestions that follow will help you develop the most efficient design possible for your application program.

When you have a general sequence of calls mapped out for your program, look over the guidelines on sequence to see if you can improve it. An efficient sequence of requests results in efficient internal IMS processing. As you write your program, keep in mind the guidelines explained in this section. The following list offers programming guidelines that will help you write efficient and error-free programs.

- Use the most simple call. Qualify your requests to narrow the search for IMS.
- Use the request or sequence of requests that will give IMS the shortest path to the segment you want.
- Use as few requests as possible. Each DL/I call your program issues uses system time and resources. You may be able to eliminate unnecessary calls by:
 - Using path requests when you are replacing, retrieving, or inserting more than one segment in the same path. If you are using more than one request to do this, you are issuing unnecessary requests.
 - Changing the sequence so that your program saves the segment in a separate I/O area, and then gets it from that I/O area the subsequent times it needs the segment. If your program retrieves the same segment more than once during program execution, you are issuing unnecessary requests.
 - Anticipating and eliminating needless and nonproductive requests, such as requests that result in GB, GE, and II status codes. For example, if you are issuing GN calls for a particular segment type, and you know how many occurrences of that segment type exist, do not issue the GN that results in a GE status code. Keep track of the number of occurrences your program retrieves, and then continue with other processing when you know you have retrieved all the occurrences of that segment type.
 - Issuing an insert request with a qualification for each parent, rather than issuing Get requests for the parents to make sure that they exist. If IMS returns a GE status code, at least one of the parents does not exist. When you are inserting segments, you cannot insert dependent segments unless the parent segments exist.
- Commit your updates regularly. To ensure optimal performance and system stability, it is recommended that updates to full-function databases be committed regularly.

IMS has removed the restriction that previously limited the number of databases with uncommitted updates to 300 at a time. However, to minimize the risk of contention or potential resource issues, it is advised to maintain the number of databases with uncommitted updates within the 300-database threshold whenever possible.

- Keep the main section of the program logic together. For example, branch to conditional routines, such as error and print routines in other parts of the program, instead of branching around them to continue normal processing.
- Use call sequences that make good use of the physical placement of the data. Access segments in hierarchic sequence as often as possible, and avoid moving backward in the hierarchy.

- Process database records in order of the key field of the root segments. (For HDAM and PHDAM databases, this order depends on the randomizing routine that is used. Check with your DBA for this information.)
- Avoid constructing the logic of the program and the structure of commands or calls in a way that depends heavily on the database structure. Depending on the current structure of the hierarchy reduces the program's flexibility.
- Minimize the number of segments your program locks. You may need to take checkpoints to release the locks on updated segments and the lock on the current database record for each PCB your program uses. Each PCB used by your program has the current database record locked at share or update level. If this lock is no longer required, issuing the GU call, qualified at the root level with a greater-than operator for a key of X'FF' (high values), releases the current lock without acquiring a new lock.

Do not use the minimization technique if you use a randomizer that puts high values at the end of the database and you use secondary indexes. If there is another root beyond the supposed high value key, IMS returns a GE to allow the application to determine the next step. A secondary index might not work because the hierarchical structure is inverted, and although the key is past the last root in the index, it might not be past the last root in the database.

Using PCBs with a processing option of get (G) results in locks for the PCB at share level. This allows other programs that use the get processing option to concurrently access the same database record. Using a PCB with a processing option that allows updates (I, R, or D) results in locks for the PCB at update level. This does not allow any other program to concurrently access the same database record.

Related concepts

[“Reserving segments for the exclusive use of your program” on page 270](#)

You may want to reserve a segment and prohibit other programs from updating the segment while you are using it. To some extent, IMS does this for you through resource lock management. The Q command code lets you reserve segments in a different way.

Segment search arguments (SSAs)

Segment search arguments (SSAs) specify information for IMS to use in processing a DL/I call. Regardless of the datatype for the field specified in a SSA, the SSA treats the field as a binary type and does a binary comparison.

A DL/I call with one or more SSAs is a *qualified call*, and a DL/I call without SSAs is an *unqualified call*.

Unqualified SSAs

Contains only a segment name.

Qualified SSAs

Includes one or more qualification statements that name a segment occurrence. The C command and a segment occurrence's concatenated key can be substituted for a qualification statement.

You can use SSA to select segments by name and to specify search criteria for specific segments. Specific segments are described by adding qualification statements to the DL/I call. You can further qualify your calls by using command codes.

Unqualified SSAs

An unqualified SSA gives the name of the segment type that you want to access. In an unqualified SSA, the segment name field is 8 bytes and must be followed by a 1-byte blank. If the actual segment name is fewer than 8 bytes long, it must be padded to the right with blanks. An example of an unqualified SSA follows:

```
PATIENTbb
```

Qualified SSAs

To qualify an SSA, you can use either a field or the sequence field of a virtual child. A qualified SSA describes the segment occurrence that you want to access. This description is called a qualification statement and has three parts. The following table shows the structure of a qualified SSA.

Table 31. Qualified SSA structure

SSA Component	Field Length
Segment name	8
(	1
Field name	8
Relative operator	2
Field value	Variable
)	1

Using a qualification statement enables you to give IMS information about the particular segment occurrence that you are looking for. You do this by giving IMS the name of a field within the segment and the value of the field you are looking for. The field and the value are connected by a relational operator (R.O. in the previous table) which tells IMS how you want the two compared. For example, to access the PATIENT segment with the value 10460 in the PATNO field, you could use this SSA:

```
PATIENTb(PATNObb=b10460)
```

Alternatively, if the DL/I call uses command code O, you can use a 4-byte starting offset position and 4-byte data length instead of an 8-byte field name. The starting offset is relative to the physical segment definition and starts with 1. The maximum length that can be retrieved is the maximum segment size for the database type, and the minimum length is 1. The two fields are specified in the following format: 'ooooLLLL'. oooo is the offset position and LLLL is the length of the data that you want to retrieve. You can use this approach to search for and retrieve data without a field definition.

The qualification statement is enclosed in parentheses. The first field contains the name of the field (Fld Name in the previous table) that you want IMS to use in searching for the segment. The second field contains a relational operator. The relational operator can be any one of the following:

- Equal, represented as

```
=b  
b=  
EQ
```

- Greater than, represented as

```
>b  
b>  
GT
```

- Less than, represented as

```
<b  
b<  
LT
```

- Greater than or equal to, represented as

```
>=  
=>  
GE
```

- Less than or equal to, represented as

<=
=<
LE

- Not equal to, represented as

¬=
=¬
NE

The third field (Fld Value in the previous table) contains the value that you want IMS to use as the comparative value. The length of Fld Value must be the same length as the field specified by Fld Name.

You can use more than one qualification statement in an SSA. Special cases exist, such as in a virtual logical child segment when the sequence field consists of multiple fields.

Sequence fields of a virtual logical child

As a general rule, a segment can have only one sequence field. However, in the case of the virtual logical-child segment type, multiple FIELD statements can be used to define a noncontiguous sequence field.

When specifying the sequence field for a virtual logical child segment, if the field is not contiguous, the length of the field named in the SSA is the concatenated length of the specified field plus all succeeding sequence fields. The following figure shows a segment with a noncontiguous sequence field.



Figure 47. Segment with a noncontiguous sequence field

If the first sequence field is not included in a "scattered" sequence field in an SSA, IMS treats the argument as a data field specification, rather than as a sequence field.

Related reading: For more information on the virtual logical child segment, refer to *IMS Version 15.3 Database Administration*.

Related concepts

[“Specifying segment search arguments using the SSAList interface” on page 646](#)

The SSAList interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

SSA guidelines

Using SSAs can simplify your programming, because the more information you can give IMS to do the searching for you, the less program logic you need to analyze and compare segments in your program.

Using SSAs does not necessarily reduce system overhead, such as internal logic and I/Os, required to obtain a specific segment. To locate a particular segment without using SSAs, you can issue DL/I calls and include program logic to examine key fields until you find the segment you want. By using SSAs in your DL/I calls, you can reduce the number of DL/I calls that are issued and the program logic needed to examine key fields. When you use SSAs, IMS does this work for you.

Recommendations:

- Use qualified calls with qualified SSAs whenever possible. SSAs act as filters, returning only the segments your program requires. This reduces the number of calls your program makes, which provides better performance. It also provides better documentation of your program. Qualified SSAs

are particularly useful when adding segments with insert calls. They ensure that the segments are inserted where you want them to go.

- For the root segment, specify the key field and an equal relational operator, if possible. Using a key field with an equal-to, equal-to-or-greater-than, or greater-than operator lets IMS go directly to the root segment.
- For dependent segments, it is desirable to use the key field in the SSA, although it is not as important as at the root level. Using the key field and an equal-to operator lets IMS stop the search at that level when a higher key value is encountered. Otherwise IMS must search through all occurrences of the segment type under its established parent in order to determine whether a particular segment exists.
- If you often must search for a segment using a field other than the key field, consider putting a secondary index on the field.

For example, suppose you want to find the record for a patient by the name of "Ellen Carter". As a reminder, the patient segment in the examples contains three fields: the patient number, which is the key field; the patient name; and the patient address. The fact that patient number is the key field means that IMS stores the patient segments in order of their patient numbers. The best way to get the record for "Ellen Carter" is to supply her patient number in the SSA. If her number is 09000, your program uses this call and SSA:

```
GUbbbbbbPATIENTb (PATN0bbb=b09000)
```

If your program supplies an invalid number, or if someone has deleted Ellen Carter's record from the database, IMS does not need to search through all the PATIENT occurrences to determine that the segment does not exist.

However, if your program does not have the number and must give the name instead, IMS must search through all the patient segments and read each patient name field until it finds "Ellen Carter" or until it reaches the end of the patient segments.

Related concepts

[“Secondary indexing and logical relationships” on page 271](#)

Secondary indexing and logical relationships are techniques that can change your application program's view of the data. The DBA makes the decision about whether to use these options.

Multiple qualification statements

When you use a qualification statement, you can do more than give IMS a field value with which to compare the fields of segments in the database. You can give several field values to establish limits for the fields you want IMS to compare.

You can use a maximum of 1024 qualification statements on a call.

Connect the qualification statements with one of the Boolean operators. You can indicate to IMS that you are looking for a value that, for example, is greater than A **and** less than B, or you can indicate that you are looking for a value that is equal to A **or** greater than B. The Boolean operators are:

Logical AND

For a segment to satisfy this request, the segment must satisfy both qualification statements that are connected with the logical AND (coded * or &).

Logical OR

For a segment to satisfy this request, the segment can satisfy either of the qualification statements that are connected with the logical OR (coded + or |).

One more Boolean operator exists and is called the independent AND. Use it only with secondary indexes.

For a segment to satisfy multiple qualification statements, the segment must satisfy a set of qualification statements. A set is a number of qualification statements that are joined by an AND. To satisfy a set, a segment must satisfy each of the qualification statements within that set. Each OR starts a new set of qualification statements. When processing multiple qualification statements, IMS reads them left to right and processes them in that order.

When you include multiple qualification statements for a root segment, the fields you name in the qualification statements affect the range of roots that IMS examines to satisfy the call. DL/I examines the qualification statements to determine the minimum acceptable key value.

If one or more of the sets do not include at least one statement that is qualified on the key field with an operator of equal-to, greater-than, or equal-to-or-greater-than, IMS starts at the first root of the database and searches for a root that meets the qualification.

If each set contains at least one statement that is qualified on the key field with an equal-to, greater-than, or equal-to-or-greater-than operator, IMS uses the lowest of these keys as the starting place for its search. After establishing the starting position for the search, IMS processes the call by searching forward sequentially in the database, similar to the way it processes GN calls. IMS examines each root it encounters to determine whether the root satisfies a set of qualification statements. IMS also examines the qualification statements to determine the maximum acceptable key value.

If one or more of the sets do not include at least one statement that is qualified on the key field with an operator of equal-to, less-than-or-equal-to, or less-than, IMS determines that no maximum key value exists. If each set contains at least one statement that is qualified on the key field with an equal-to, less-than, or equal-to-or-less-than operator, IMS uses the maximum of these keys to determine when the search stops.

IMS continues the search until it satisfies the call, encounters the end of the database, or finds a key value that exceeds the maximum. If no maximum key value is found, the search continues until IMS satisfies the call or encounters the end of the database.

Examples: Shown below are cases of SSAs used at the root level:

```
ROOTKEYb
=b10&FIELDBb
b=XYZ+ROOTKEYb
=10&FIELDBb
b
=ABC
```

In this case, the minimum and maximum key is 10. This means that IMS starts searching with key 10 and stops when it encounters the first key greater than 10. To satisfy the SSA, the ROOTKEY field must be equal to 10, and FIELDB must be equal to either ABC or XYZ.

```
ROOTKEYb
=>10&ROOTKEYb
<20
```

In this case, the minimum key is 10 and the maximum key is 20. Keys in the range of 10 to 20 satisfy the SSA. IMS stops the search when it encounters the first key greater than 20.

```
ROOTKEYb
10&ROOTKEYb
=<20+ROOTKEYb
=>110&ROOTKEYb
=<120
```

In this case, the minimum key is 10 and the maximum key is 120. Keys in the range of 10 to 20 and 110 to 120 satisfy the call. IMS stops the search when it encounters the first key greater than 120. IMS does not scan from 20 to 110 but skips forward (using the index for HIDAM or PHIDAM) from 20 to 110. Because of this, you can use ranges for more efficient program operation.

When you use multiple qualification statement segments that are part of logical relationships, additional considerations exist.

Related concepts

[“Multiple qualification statements with secondary indexes” on page 272](#)

When you qualify a call using the name of an indexed field, you can include multiple qualification statements.

Example of how to use multiple qualification statements

The following example shows how you can use multiple qualification statements.

Given the sample Medical database, we want to answer the following question:

Did we see patient number 04120 during 1992?

To find the answer to this question, you need to give IMS more than the patient's name; you want IMS to search through the ILLNESS segments for that patient, read each one, and return any that have a date in 1992. The call you would issue to do this is:

```
GU  PATIENTb(PATN0bbbEQ04120)
    ILLNESSb(ILLDATEb>=19920101&ILLDATEb<=19921231)
```

In other words, you want IMS to return any ILLNESS segment occurrences under patient number 04120 that have a date on or after January 1, 1992, and on or before December 31, 1992, joined with an AND connector. Suppose you wanted to answer the following request:

Did we see Judy Jennison during January of 1992 or during July of 1992? Her patient number is 05682.

You could issue a GU call with the following SSAs:

```
GU  PATIENTb(PATN0bEQ05682)
    ILLNESSb(ILLDATEb>=19920101&ILLDATEb<=19920131|ILLDATEb>=19920701&ILLDATEb<=19920731)
```

To satisfy this request, the value for ILLDATE must satisfy either of the two sets. IMS returns any ILLNESS segment occurrences for the month of January 1992, or for the month of July 1992.

Multiple qualification statements for HDAM, PHDAM, or DEDB

For HDAM (Hierarchical Direct Access Method), PHDAM (partitioned HDAM), or data entry database (DEDB) organizations, a randomizing exit routine usually does not store the root keys in ascending key sequence. For these organizations, IMS determines the minimum and maximum key values. The minimum key value is passed to the randomizing exit routine, which determines the starting anchor point.

The first root off this anchor is the starting point for the search. When IMS encounters a key that exceeds the maximum key value, IMS terminates the search with a GE status code. If the randomizing routine randomized so that the keys are stored in ascending key sequence, a call for a range of keys will return all of the keys in the range. However, if the randomizing routine did not randomize into key sequence, the call does not return all keys in the requested range. Therefore, use calls for a range of key values only when the keys are in ascending sequence (when the organization is HDAM, PHDAM, or DEDB).

Recommendations:

- When the organization is HDAM, PHDAM, or DEDB, use calls for a range of key values only when the keys are in ascending sequence.
- When the organization is HDAM, PHDAM or DEDB, do not use calls that allow a range of values at the root level.

While not recommended, a sequential search of the database can be accomplished with the use of command codes A and G when making GN/GHN database calls. Command code A will clear positioning and cause the call to start at the beginning of the database. Command code G will prevent randomization and cause a sequential search of the database when used with SSAs that specify a range of values at the root level. The returned segments may not be in sequential order depending on how they were randomized.

To search the database sequentially, you can use the use the following segment search argument (SSA) together with SSAs that specify a range of values at the root level.

key field > hex zeros & key field < all f's key

The returned segments may not be in sequential order depending on how they were randomized.

For more details about HDAM or PHDAM databases, see *IMS Version 15.3 Database Administration*.

SSAs and command codes

SSAs can also include one or more command codes, which can change and extend the functions of DL/I calls.

For information on DL/I command codes, see [Command code reference \(Application Programming APIs\)](#).

IMS always returns the lowest segment in the path to your I/O area. If your program codes a D command code in an SSA, IMS also returns the segment described by that SSA. A call that uses the D command code is called a *path call*.

For example, suppose your program codes a D command code on a GU call that retrieves segment F and all segments in the path to F in the hierarchy shown in the following figure.




Figure 48. D command code example

The call function and the SSAs for the call look like this:

```
GU  Abbbbbbb  
    *D  
    Cbbbbbbb  
    *D  
    Ebbbbbbb  
    Fbbbbbbb
```

A command code consists of one letter. Code the command codes in the SSA after the segment name field. Separate the segment name field and the command code with an asterisk, as shown in the following table.

Table 32. Unqualified SSA with command code

SSA Component	Field Length
Seg Name	8
*	1
Cmd Code	Variable
b	1

Your program can use command codes in both qualified and unqualified SSAs. However, command codes cannot be used by MSDB calls. If the command codes are not followed by qualification statements, they

must each be followed by a 1-byte blank. If the command codes are followed by qualification statements, do not use the blank. The left parenthesis of the qualification statement follows the command code instead, as indicated in the following table.

Table 33. Qualified SSA with command code

SSA Component	Field Length
Seg Name	8
*	1
Cmd Code	Variable
(	1
Field name or, if the O cmd code is specified, either the field name or the field position and length	8
Relational Operator (R.O.)	2
Field Value	Variable
)	1

By giving IMS the field position within the segment and the value of the field you are looking for, the field position and the value are connected by a relational operator which tells IMS how you want the two to be compared. The field position can be either a searchable field name as defined in the DBD or a position and length when using command code O.

If your program uses command codes to manage subset pointers in a DEDB, enter the number of the subset pointer immediately after the command code. Subset pointers are a means of dividing a chain of segment occurrences under the same parent into two or more groups or subsets. Your program can define as many as eight subset pointers for any segment type. Using an application program, your program can then manage these subset pointers.

Related concepts

“Processing Fast Path DEDBs with subset pointer command codes” on page 304

Subset pointers and the command codes you use with them are optimization tools that significantly improve the efficiency of your program when you need to process long segment chains.

Considerations for coding DL/I calls and data areas

If you have made all the design decisions about your program, coding the program is a matter of implementing the decisions that you have made. In addition to knowing the design and processing logic for your program, you need to know about the data that your program is processing, the PCBs it references, and the segment formats in the hierarchies your program processes.

You can use the following list as a checklist to make sure you are not missing any information. If you are missing information about data, IMS options being used in the application program, or segment layouts and the application program's data structures, obtain this information from the DBA or the equivalent specialist at your installation. Be aware of the programming standards and conventions that have been established at your installation.

Program design considerations:

- The sequence of calls for your program.
- The format of each call:
 - Does the call include any SSAs?
 - If so, are they qualified or unqualified?
 - Does the call contain any command codes?

- The processing logic for the program.
- The routine the program uses to check the status code after each call.
- The error routine the program uses.

Checkpoint considerations:

- The type of checkpoint call to use (basic or symbolic).
- The identification to assign to each checkpoint call, regardless of whether the Checkpoint call is basic or symbolic.
- If you are going to use the symbolic checkpoint call, which areas of your program to checkpoint.

Segment considerations:

- Whether the segment is fixed length or variable length.
- The length of the segment (the maximum length, if the segment is variable length).
- The names of the fields that each segment contains.
- Whether the segment has a key field. If it does, is the key field unique or non-unique? If it does not, what sequencing rule has been defined for it? (A segment's key field is defined in the SEQ keyword of the FIELD statement in the DBD. The sequencing rule is defined in the RULES keyword of the SEGM statement in the DBD.)
- The segment's field layouts:
 - The byte location of each field.
 - The length of each field.
 - The format of each field.

Data structure considerations:

- Each data structure your program processes has been defined in a DB PCB. All of the PCBs your program references are part of a PSB for your application program. You need to know the order in which the PCBs are defined in the PSB.
- The layout of each of the data structures your program processes.
- Whether multiple or single positioning has been specified for each data structure. This is specified in the POS keyword of the PCB statement during PSB generation.
- Whether any data structures use multiple DB PCBs.

Preparing to run your CICS DL/I call program

You must perform several steps before you run your CICS DL/I call program.

Refer to the appropriate CICS reference information:

- For information on translating, compiling, and binding your CICS online program, see the description of installing application programs in *CICS Transaction Server for z/OS CICS System Definition Guide*.
- For information on which compiler options should be used for a CICS online program, as well as for CICS considerations when converting a CICS online COBOL program with DL/I calls to Enterprise COBOL, see *CICS Transaction Server for z/OS CICS Application Programming Guide*.

Examples of how to code DL/I calls and data areas

You can code DL/I calls and data areas in assembler language, C, COBOL, Pascal, Java, and PL/I.

Coding a batch program in assembler language

The following code example shows how to write an IMS program to access the IMS database in assembler language.

The numbers to the right of the program refer to the notes that follow the program. This kind of program can run as a batch program or as a batch-oriented BMP.

Sample assembler language program

			NOTES
PGMSTART	CSECT		1
*	EQUATE	REGISTERS	
*	USEAGE	OF REGISTERS	
R1	EQU	1 ORIGINAL PCBLIST ADDRESS	
R2	EQU	2 PCBLIST ADDRESS1	
R5	EQU	5 PCB ADDRESSS	
R12	EQU	12 BASE ADDRESS	
R13	EQU	13 SAVE AREA ADDRESS	
R14	EQU	14	
R15	EQU	15	
*			
	USING	PGMSTART,R12 BASE REGISTER ESTABLISHED	2
	SAVE	(14,12) SAVE REGISTERS	
	LR	12,15 LOAD REGISTERS	
	ST	R13,SAVEAREA+4 SAVE AREA CHAINING	
	LA	R13,SAVEAREA NEW SAVE AREA	
	USING	PCBLIST,R2 MAP INPUT PARAMETER LIST	
	USING	PCBNAME,R5 MAP DB PCB	
	LR	R2,R1 SAVE INPUT PCB LIST IN REG 2	
	L	R5,PCBDETA LOAD DETAIL PCB ADDRESS	
	LA	R5,0(R5) REMOVE HIGH ORDER END OF LIST FLAG	3
	CALL	ASMTDLI,(GU,(R5),DETSEGIO,SSANAME),VL	4
*			
*			
	L	R5,PCBMSTA LOAD MASTER PCB ADDRESS	
	CALL	ASMTDLI,(GHU,(R5),MSTSEGIO,SSAU),VL	5
*			
*			
	CALL	ASMTDLI,(GHN,(R5),MSTSEGIO),VL	6
*			
*			
	CALL	ASMTDLI,(REPL,(R5),MSTSEGIO),VL	
*			
*			
	L	R13,4(R13) RESTORE SAVE AREA	
	RETURN	(14,12) RETURN BACK	7
*			
*	FUNCTION	CODES USED	
*			
GU	DC	CL4'GU'	
GHU	DC	CL4'GHU'	
GHN	DC	CL4'GHN'	
REPL	DC	CL4'REPL'	8
*			
*	SSAS		
*			
SSANAME	DS	0C	
	DC	CL8'ROOTDET'	
	DC	CL1'('	
	DC	CL8'KEYDET'	9
	DC	CL2'='	
NAME	DC	CL5' '	
	DC	C')'	
*			
SSAU	DC	CL9'ROOTMST'*	
MSTSEGIO	DC	CL100' '	
DETSEGIO	DC	CL100' '	
SAVEAREA	DC	18F'0'	
*			10
PCBLIST	DSECT		
PCBIO	DS	A ADDRESS OF I/O PCB	
PCBMSTA	DS	A ADDRESS OF MASTER PCB	

PCBDETA	DS	A	ADDRESS OF DETAIL PCB	11
*				
PCBNAME	DSECT			
DBPCBDBD	DS	CL8	DBD NAME	
DBPCBLEV	DS	CL2	LEVEL FEEDBACK	
DBPCBSTC	DS	CL2	STATUS CODES	
DBPCBPRO	DS	CL4	PROC OPTIONS	
DBPCBRVS	DS	F	RESERVED	
DBPCBSFD	DS	CL8	SEGMENT NAME FEEDBACK	
DBPCBMKL	DS	F	LENGTH OF KEY FEEDBACK	
DBPCBNSS	DS	F	NUMBER OF SENSITIVE SEGMENTS IN PCB	
DBPCBKFD	DS	C	KEY FEEDBACK AREA	
	END	PGMSTART		

Note:

1. The entry point to an assembler language program can have any name. Also, you can substitute CBLTDLI for ASMTDLI in any of the calls.
2. When IMS passes control to the application program, register 1 contains the address of a variable-length fullword parameter list. Each word in this list contains the address of a PCB that the application program must save. The high-order byte of the last word in the parameter list has the 0 bit set to a value of 1 which indicates the end of the list. The application program subsequently uses these addresses when it executes DL/I calls.
3. The program loads the address of the DETAIL DB PCB.
4. The program issues a GU call to the DETAIL database using a qualified SSA (SSANAME).
5. The program loads the address of the HALDB master PCB.
6. The next three calls that the program issues are to the HALDB master. The first is a GHU call that uses an unqualified SSA. The second is an unqualified GHN call. The REPL call replaces the segment retrieved using the GHN call with the segment in the MSTSEGIO area.

You can use the *parmcount* parameter in DL/I calls in assembler language instead of the VL parameter, except for in the call to the sample status-code error routine.
7. The RETURN statement loads IMS registers and returns control to IMS.
8. The call functions are defined as four-character constants.
9. The program defines each part of the SSA separately so that it can modify the SSA's fields.
10. The program must define an I/O area that is large enough to contain the largest segment it is to retrieve or insert (or the largest path of segments if the program uses the D command code). This program's I/O areas are 100 bytes each.
11. A fullword must be defined for each PCB. The assembler language program can access status codes after a DL/I call by using the DB PCB base addresses.

This example assumes that an I/O PCB was passed to the application program. If the program is a batch program, CMPAT=YES must be specified on the PSBGEN statement of PSBGEN so that the I/O PCB is included. Because the I/O PCB is required for a batch program to make system service calls, CMPAT=YES should always be specified.

Restriction: The IMS language interface module (DFSLI000) must be bound to the compiled assembler language program.

Coding a CICS online program in assembler language

The following code example in assembler language shows how you define and establish addressability to the UIB.

The numbers to the right of the program refer to the notes that follow the program. This program can run in a CICS environment using DBCTL.

Sample call-level assembler language program (CICS online)

PGMSTART	DSECT			NOTES
UIBPTR	DS	F		

IOAREA	DS	0CL40	1
AREA1	DS	CL3	
AREA2	DS	CL37	
	DLIUIB		
	USING	UIB,8	2
PCBPTRS	DSECT		
*	PSB	ADDRESS LIST	
PCB1PTR	DS	F	
PCB1	DSECT		
	USING	PCB1,6	3
DBPC1DBD	DS	CL8	
DBPC1LEV	DS	CL2	
DBPC1STC	DS	CL2	
DBPC1PRO	DS	CL4	
DBPC1RSV	DS	F	
DBPC1SFD	DS	CL8	
DBPC1MKL	DS	F	
DBPC1NSS	DS	F	
DBPC1KFD	DS	0CL256	
DBPC1NM	DS	0CL12	
DBPC1NMA	DS	0CL14	
DBPC1NMP	DS	CL17	
ASMUIB	CSECT		
	B	SKIP	
PSBNAME	DC	CL8'ASMP SB'	
PCBFUN	DC	CL4'PCB'	
REPLFUN	DC	CL4'REPL'	
TERMFUN	DC	CL4'TERM'	
GHUFUN	DC	CL4'GHU'	
SSA1	DC	CL9'AAAA4444'	
GOODRC	DC	XL1'00'	
GOODSC	DC	CL2' '	
SKIP	DS	0H	4
*		SCHEDULE PSB AND OBTAIN PCB ADDRESSES	
	CALLDLI	ASMTDLI,(PCBFUN,PSBNAME,UIBPTR)	
	L	8,UIBPTR	5
	CLC	UIBFCTR,X'00'	
	BNE	ERROR1	
*		GET PSB ADDRESS LIST	
	L	4,UIBPCBAL	
	USING	PCBPTRS,4	
*		GET ADDRESS OF FIRST PCB IN LIST	
	L	6,PCB1PTR	
*		ISSUE DL/I CALL: GET A UNIQUE SEGMENT	
	CALLDLI	ASMTDLI,(GHUFUN,PCB1,IOAREA,SSA1)	6
	CLC	UIBFCTR,GOODRC	
	BNE	ERROR2	
	CLC	DBPC1STC,GOODSC	
	BNE	ERROR3	7
*		PERFORM SEGMENT UPDATE ACTIVITY	
	MVC	AREA1,.....	
	MVC	AREA2,.....	
*		ISSUE DL/I CALL: REPLACE SEGMENT AT CURRENT POSITION	
	CALLDLI	ASMTDLI,(REPLFUN,PCB1,IOAREA,SSA1)	8
	CLC	UIBFCTR,GOODRC	
	BNE	ERROR4	
	CLC	DBPC1STC,GOODSC	
	B	TERM	
ERROR1	DS	0H	
*		INSERT ERROR DIAGNOSTIC CODE	
	B	TERM	
ERROR2	DS	0H	
*		INSERT ERROR DIAGNOSTIC CODE	
	B	TERM	
ERROR3	DS	0H	
*		INSERT ERROR DIAGNOSTIC CODE	
	B	TERM	
ERROR4	DS	0H	
*		INSERT ERROR DIAGNOSTIC CODE	
ERROR5	DS	0H	
*		INSERT ERROR DIAGNOSTIC CODE	
	B	TERM	
TERM	DS	0H	
*		RELEASE THE PSB	
	CALLDLI	ASMDLI,(TERMFUN)	
	EXEC	CICS RETURN	
	END	ASMUIB	9,10

Note:

1. The program must define an I/O area that is large enough to contain the largest segment it is to retrieve or insert (or the largest path of segments if the program uses the D command code).
2. The DLIUIB statement copies the UIB DSECT.
3. A fullword must be defined for each DB PCB. The assembler language program can access status codes after a DL/I call by using the DB PCB base addresses.
4. This is an unqualified SSA. For qualified SSA, define each part of the SSA separately so that the program can modify the fields of the SSA.
5. This call schedules the PSB and obtains the PSB address.
6. This call retrieves a segment from the database.

CICS online assembler language programs use the CALLDLI macro, instead of the call statement, to access DL/I databases. This macro is similar to the call statement. It looks like this:

```
CALLDLI ASMTDLI,(function,PCB-name,ioarea, SSA1,...SSAn),VL
```

7. CICS online programs must check the return code in the UIB before checking the status code in the DB PCB.
8. The REPL call replaces the data in the segment that was retrieved by the most recent Get Hold call. The data is replaced by the contents of the I/O area referenced in the call.
9. This call releases the PSB.
10. The RETURN statement loads IMS registers and returns control to IMS.

Related reading: For more information on installing CICS application programs, see *CICS Transaction Server for z/OS CICS Application Programming Reference*.

Related reference

[“Specifying the UIB \(CICS online programs only\)” on page 219](#)

The interface between your CICS online program and DL/I passes additional information to your program in a user interface block (UIB). The UIB contains the address of the PCB list and any return codes your program must examine before checking the status code in the DB PCB.

Coding a batch program in C language

The following code example shows how to write an IMS program to access the IMS database in C language.

The numbers to the right of the program refer to the notes that follow the program.

Sample C language program

```
#pragma runopts(env(IMS),plist(IMS))
#include <ims.h>
#include <stdio.h>
main() {
    /*
    /*      descriptive statements
    /*
    IO_PCB_TYPE *IO_PCB = (IO_PCB_TYPE*)PCBLIST[0];
    struct {PCB_STRUCT(10)} *mast_PCB = __pcblist[1];
    struct {PCB_STRUCT(20)} *detail_PCB = __pcblist[2];
    const static char func_GU[4] = "GU ";
    const static char func_GN[4] = "GN ";
    const static char func_GHU[4] = "GHU ";
    const static char func_GHN[4] = "GHN ";
    const static char func_GNP[4] = "GNP ";
    const static char func_GHNP[4] = "GHNP";
    const static char func_ISRT[4] = "ISRT";
    const static char func_REPL[4] = "REPL";
    const static char func_DLET[4] = "DLET";
    char qual_ssa[8+1+8+2+6+1+1]; /* initialized by sprintf
    /*below. See the */
    /*explanation for */
    /*sprintf in note 7 for the */
    /*meanings of 8,1,8,2,6,1 —*/
    /*the final 1 is for the */
}
```

NOTES

1
2

3

4

5


```

static const char unqual_ssa[] = /*trailing '\0' of string */
                                "NAME ";
                                /* 12345678_ */
struct {
    ____
    ____
    ____
} mast_seg_io_area;

struct {
    ____
    ____
    ____
} det_seg_io_area;

/*
/*      Initialize the qualifier
/*
/*
/*
    sprintf(qual_ssa,
            "8.8s(8.8s6.6s)",
            "ROOT", "KEY", "=", "vvvvv");
/*
/*      Main part of C batch program
/*
/*
    ctdli(func_GU, detail_PCB,
          &det_seg_io_area, qual_ssa);

    ctdli(func_GHU, mast_PCB,
          &mast_seg_io_area, qual_ssa);

    ctdli(func_GHN, mast_PCB,
          &mast_seg_io_area);

    ctdli(func_REPL, mast_PCB,
          &mast_seg_io_area);
}

```

6

7

8

9

10

11

12

Note:

1. The **env(IMS)** establishes the correct operating environment and the **plist(IMS)** establishes the correct parameter list when invoked under IMS. The **ims.h** header file contains declarations for PCB layouts, **__pcblist**, and the **ctdli** routine. The PCB layouts define masks for the PCBs that the program uses as structures. These definitions make it possible for the program to check fields in the PCBs.

The **stdio.h** header file contains declarations for **sprintf** (used to build up the SSA).

2. After IMS has loaded the application program's PSB, IMS gives control to the application program through this entry point.
3. The C run-time sets up the **__pcblist** values. The order in which you refer to the PCBs must be the same order in which they have been defined in the PSB. (Values other than "10" and "20" can be used, according to the actual key lengths needed.) These declarations can be done using macros, such as:

```

#define IO_PCB (IO_PCB_TYPE *) (__pcblist[0])
#define mast_PCB (__pcblist[1])
#define detail_PCB (__pcblist[2])

```

This example assumes that an I/O PCB was passed to the application program. When the program is a batch program, **CMPAT=YES** must be specified on the **PSBGEN** statement of **PSBGEN** so that the I/O PCB is included. Because the I/O PCB is required for a batch program to make system service calls, **CMPAT=YES** should always be specified for batch programs.

4. Each of these areas defines one of the call functions used by the batch program. Each character string is defined as four alphanumeric characters, with a value assigned for each function. (If the **[4]s** had been left out, 5 bytes would have been reserved for each constant.) You can define other constants in the same way. Also, you can store standard definitions in a source library and include them by using a **#include** directive.

Instead, you can define these by macros, although each string would have a trailing null (**'\0'**).

5. The SSA is put into a string (see note 7). You can define a structure, as in COBOL, PL/I, or Pascal, but using **sprintf** is more convenient. (Remember that C strings have trailing nulls that cannot be passed

to IMS.) Note that the string is 1 byte longer than required by IMS to contain the trailing null, which is ignored by IMS. Note also that the numbers in brackets assume that six fields in the SSA are equal to these lengths.

6. The I/O areas that will be used to pass segments to and from the database are defined as structures.
7. The **sprintf** function is used to fill in the SSA. The "**%-8.8s**" format means "a left-justified string of exactly eight positions". The "**%2.2s**" format means "a right-justified string of exactly two positions".

Because the **ROOT** and **KEY** parts do not change, this can also be coded:

```
sprintf(qual_ssa,
        "ROOT      (KEY      =%-6.6s)", "vvvvvv");
/* 12345678 12345678 */
```

8. This call retrieves data from the database. It contains a qualified SSA. Before you can issue a call that uses a qualified SSA, initialize the data field of the SSA. Before you can issue a call that uses an unqualified SSA, initialize the segment name field. Unlike the COBOL, PL/I, and Pascal interface routines, **ctdli** also returns the status code as its result. (Blank is translated to 0.) So, you can code:

```
switch (ctdli(...)) {
    case 0: ... /* everything ok */

        break;
    case 'AB': ....

        break;
    case 'IX': ...

        break;
    default:

}
}
```

You can pass only the PCB pointer for DL/I calls in a C program.

9. This is another call with a qualified SSA.
10. This call is an unqualified call that retrieves data from the database. Because it is a Get Hold call, it can be followed by REPL or DLET.
11. The REPL call replaces the data in the segment that was retrieved by the most recent Get Hold call. The data is replaced by the contents of the I/O area that is referenced in the call.
12. The end of the **main** routine (which can be done by a **return** statement or **exit** call) returns control to IMS.

Restriction: IMS provides a language interface module (DFSLI000) that is an interface between IMS and the C language. This module must be made available to the application program at bind time.

Coding a batch program in COBOL

The following code example shows how to write an IMS program to access the IMS database in COBOL.

The numbers to the right of the program refer to the notes that follow the program. This kind of program can run as a batch program or as a batch-oriented BMP.

Sample COBOL program

```
Identification Division.
Program-ID. BATCOBOL.
Environment Division.
Data Division.
Working-Storage Section.
01 Func-Codes.
   05 Func-GU          Picture XXXX Value 'GU  '.
   05 Func-GHU         Picture XXXX Value 'GHU  '.
   05 Func-GN          Picture XXXX Value 'GHN  '.
   05 Func-GHN         Picture XXXX Value 'GHN  '.
   05 Func-GNP         Picture XXXX Value 'GNP  '.
   05 Func-GHNP        Picture XXXX Value 'GHNP '.
   05 Func-REPL        Picture XXXX Value 'REPL '.
   05 Func-ISRT        Picture XXXX Value 'ISRT '.
```

```

05 Func-DLET          Picture XXXX Value 'DLET'.
05 Parmcount          Picture S9(5) Value +4 Comp-5.
01 Unqual-SSA.
05 Seg-Name           Picture X(08) Value ' '.
05 Filler              Picture X          Value ' '.
01 Qual-SSA-Mast.
05 Seg-Name-M          Picture X(08) Value 'ROOTMast'.
05 Begin-Paren-M       Picture X          Value '('.
05 Key-Name-M          Picture X(08) Value 'KeyMast '.
05 Key-Oper-M          Picture X(05) Value '='.
05 Key-Value-M         Picture X(06) Value 'VVVVVV'.
05 End-Paren-M         Picture X          Value ')'.
01 Qual-SSA-Det.
05 Seg-Name-D          Picture X(08) Value 'ROOTDET '.
05 Begin-Paren-D       Picture X          Value '('.
05 Key-Name-D          Picture X(08) Value 'KEYDET '.
05 Rel-Oper-D          Picture X(05) Value '='.
05 Key-Value-D         Picture X(06) Value 'VVVVVV'.
05 End-Paren-D         Picture X          Value ')'.
01 Det-Seg-In.
05 Data1               Picture X.
05 Data2               Picture X.
01 Mast-Seg-In.
05 Data1               Picture X.
05 Data2               Picture X.
linkage section.
01 IO-PCB.
05 Filler              Picture X(10).
05 IO-Status-Code      Picture XX.
05 Filler              Picture X(20).
01 DB-PCB-Mast.
05 Mast-Dbd-Name       Picture X(8).
05 Mast-Seg-Level      Picture XX.
05 Mast-Status-Code    Picture XX.
05 Mast-Proc-Opt       Picture XXXX.
05 Filler              Picture S9(5) Comp-5.
05 Mast-Seg-Name        Picture X(8).
05 Mast-Len-KFB         Picture S9(5) Comp-5.
05 Mast-Nu-Senseg       Picture S9(5) Comp-5.
05 Mast-Key-FB          Picture X(256).
01 DB-PCB-Detail.
05 Det-Dbd-Name        Picture X(8).
05 Det-Seg-Level       Picture XX.
05 Det-Status-Code     Picture XX.
05 Det-Proc-Opt        Picture XXXX.
05 Filler              Picture S9(5) Comp-5.
05 Det-Seg-Name        Picture X(8).
05 Det-Len-KFB         Picture S9(5) Comp-5.
05 Det-Nu-Senseg       Picture S9(5) Comp-5.
05 Det-Key-FB          Picture X(256).

Procedure Division using IO-PCB DB-PCB-Mast DB-PCB-Detail.
    Call 'CBLTDLI' using Func-GU DB-PCB-Detail
        Det-seg-in Qual-SSA-Det.
    .
    Call 'CBLTDLI' using Parmcount Func-ghu DB-PCB-Mast
        Mast-seg-in Qual-SSA-Mast.
    .
    Call 'CBLTDLI' using Func-GHN DB-PCB-Mast
        Mast-seg-in.
    .
    Call 'CBLTDLI' using Func-REPL DB-PCB-Mast
        Mast-seg-in.
    .
    Goback.

```

Note:

1. You define each of the DL/I call functions the program uses with a 77-level or 01-level working storage entry. Each picture clause is defined as four alphanumeric characters and has a value assigned for each function. If you want to include the optional *parmcount* field, you can initialize count values for each type of call. You can also use a COBOL COPY statement to include these standard descriptions in the program.

2. A 9-byte area is set up for an unqualified SSA. Before the program issues a call that requires an unqualified SSA, it moves the segment name to this area. If a call requires two or more SSAs, you may need to define additional areas.
3. A 01-level working storage entry defines each qualified SSA that the application program uses. Qualified SSAs must be defined separately, because the values of the fields are different.
4. A 01-level working storage entry defines I/O areas that are used for passing segments to and from the database. You can further define I/O areas with sub-entries under the 01-level. You can use separate I/O areas for each segment type, or you can define one I/O area that you use for all segments.
5. A 01-level linkage section entry defines a mask for each of the PCBs that the program requires. The DB PCBs represent both input and output databases. After issuing each DL/I call, the program checks the status code through this linkage. You define each field in the DB PCB so that you can reference it in the program.
6. This is the standard procedure division statement of a batch program. After IMS has loaded the PSB for the program, IMS passes control to the application program. The PSB contains all the PCBs that are defined in the PSB. The coding of USING on the procedure division statement references each of the PCBs by the names that the program has used to define the PCB masks in the linkage section. The PCBs must be listed in the order in which they are defined in the PSB.

The previous code example assumes that an I/O PCB was passed to the application program. When the program is a batch program, CMPAT=YES must be specified on the PSBGEN statement of PSBGEN so that the I/O PCB is included. Because the I/O PCB is required for a batch program to make system service calls, CMPAT=YES should always be specified for batch programs.

The entry DLITCBL statement is only used in the main program. Do not use it in called programs.

7. This call retrieves data from the database by using a qualified SSA. Before issuing the call, the program must initialize the key or data value of the SSA so that it specifies the particular segment to be retrieved. The program should test the status code in the DB PCB that was referenced in the call immediately after issuing the call. You can include the *parmcount* parameter in DL/I calls in COBOL programs, except in the call to the sample status-code error routine. It is never required in COBOL.
8. This is another retrieval call that contains a qualified SSA.
9. This is an unqualified retrieval call.
10. The REPL call replaces the segment that was retrieved in the most recent Get Hold call. The segment is replaced with the contents of the I/O area that is referenced in the call (MAST-SEG-IN).
11. The program issues the GOBACK statement when it has finished processing.

Related reading: For information on how to use these procedures, see *IMS Version 15.3 System Definition*.

Binding COBOL code to the IMS language interface module

IMS supplies a language interface module (DFSLI000). This module must be bound to the batch program after the program has been compiled. It gives a common interface to IMS.

If you use the IMS-supplied procedures (IMSCOBOL or IMSCOBGO), IMS binds the language interface with the application program. IMSCOBOL is a two-step procedure that compiles and binds your program. IMSCOBGO is a three-step procedure that compiles, binds, and executes your program in an IMS batch region.

Coding a CICS online program in COBOL

The following code examples are skeleton online programs in Enterprise COBOL. They show examples of how to define and set up addressability to the UIB.

The numbers to the right of the programs refer to the notes that follow them. This kind of program can run in a CICS environment using DBCTL.

Sample COBOL program that can run in CICS

```
Identification Division.
Program-ID. CBLUIB.
Environment Division.
Data Division.
Working-Storage Section.
  01 Func-Codes.
    05 Psb-Name          Picture X(8) Value 'CBLPSB '.
    05 Func-PCB          Picture X(4) Value 'PCB '.
    05 Func-TERM         Picture X(4) Value 'TERM'.
    05 Func-GHU          Picture X(4) Value 'GHU '.
    05 Func-REPL         Picture X(4) Value 'REPL'.
    05 SSA1              Picture X(9) Value 'AAAA4444 '.
    05 Success-Message   Picture X(40).
    05 Good-Status-Code  Picture XX Value ' '.
    05 good-return-code  Picture X Value low-Value.
  01 Message0.
    05 Message1          Picture X(38).
    05 Message2          Picture XX.
  01 Dli-IO-Area.
    05 Area1              Picture X(3).
    05 Area2              Picture X(37).
  Procedure Division.
  * Schedule the psb and address the uib
    Call 'CBLTDli' using Func-PCB Psb-Name
      address of Dliuib.
    If Uibfctr is not equal low-Values then
  * Insert error diagnostic code
    Exec CICS return end-exec
  End-if.
  Set address of pcb-addresses to pcbaddr.
  * Issue DL/I Call: get a unique segment
    Set address of pcb1 to pcb-address-list(1).
    Call 'CBLTDli' using Func-GHU Pcb1
      Dli-io-area ssa1.
    If uibfctr is not equal good-return-code then
  * Insert error diagnostic code
    Exec CICS return end-Exec
  End-if.
  If pcb1-status-code is not equal good-status-code then
  * Insert error diagnostic code
    Exec CICS return end-Exec
  End-if.
  Perform segment update activity
  Move 'aaa' to area1.
  Move 'bbb' to area2.
  * Issue DL/I Call: replace segment at current position
    Call 'CBLTDli' using Func-REPL Pcb1
      Dli-io-area ssa1
    If uibfctr is not equal good-return-code then
  * Insert error diagnostic code
    Exec CICS return end-Exec
  End-if.
  If pcb1-status-code is not equal good-status-code then
  * Insert error diagnostic code
    Exec CICS return end-Exec
  End-if.
  * Release the psb
    Call 'CBLTDli' using Func-TERM.
  * Other application Function
    Exec CICS return end-Exec.
  Goback.
```

Note:

1. You define each of the DL/I call functions the program uses with a 77-level or 01-level working storage entry. Each picture clause is defined as four alphanumeric characters and has a value assigned for each function. If you want to include the optional *parmcount* field, initialize count values for each type of call. You can also use the COBOL COPY statement to include these standard descriptions in the program.
2. A 9-byte area is set up for an unqualified SSA. Before the program issues a call that requires an unqualified SSA, it can either initialize this area with the segment name or move the segment name to this area. If a call requires two or more SSAs, you may need to define additional areas.

3. An 01-level working storage entry defines I/O areas that are used for passing segments to and from the database. You can further define I/O areas with sub-entries under the 01-level. You can use separate I/O areas for each segment type, or you can define one I/O area that you use for all segments.
4. One PCB layout is defined in the linkage section. The PCB-ADDRESS-LIST occurs *n* times, where *n* is greater than or equal to the number of PCBs in the PSB.
5. The PCB call schedules a PSB for your program to use. The address of the DLIUIB parameter returns the address of DLIUIB.
6. This unqualified GHU call retrieves a segment from the database and places it in the I/O area that is referenced by the call. Before issuing the call, the program must initialize the key or data value of the SSA so that it specifies the particular segment to be retrieved.
7. CICS online programs should test the return code in the UIB before testing the status code in the DB PCB.
8. The REPL call replaces the segment that was retrieved in the most recent Get Hold call with the data that the program has placed in the I/O area.
9. The TERM call terminates the PSB the program scheduled earlier. This call is optional and is only issued if a sync point is desired prior to continued processing. The program issues the EXEC CICS RETURN statement when it has finished its processing. If this is a RETURN from the highest-level CICS program, a TERM call and sync point are internally generated by CICS.

Sample call-level OS/VS COBOL program for CICS online (obsolete with Enterprise COBOL)

Identification Division.	NOTES
Program-ID. CBLUIB.	
Environment Division.	
Data Division.	
Working-Storage Section.	
01 Func-Codes.	
05 Psb-Name Picture X(8) Value 'CBLPSB '.	1
05 Func-PCB Picture X(4) Value 'PCB '.	
05 Func-TERM Picture X(4) Value 'TERM'.	
05 Func-GHU Picture X(4) Value 'GHU '.	
05 Func-REPL Picture X(4) Value 'REPL'.	
05 SSA1 Picture X(9) Value 'AAAA4444 '.	2
05 Success-Message Picture X(40).	
05 Good-Status-Code Picture XX Value ' '.	
05 Good-Return-Code Picture X Value low-Value.	
01 Message0.	
05 Message1 Picture X(38).	
05 Message2 Picture XX.	
01 Dli-IO-Area.	3
05 Area1 Picture X(3).	
05 Area2 Picture X(37).	
Linkage Section.	4
01 BllCells.	
05 Filler Picture S9(8) Comp-5.	
05 Uib-Ptr Picture S9(8) Comp-5.	
05 B-Pcb-Ptrs Picture S9(8) Comp-5.	
05 Pcb1-Ptr Picture S9(8) Comp-5.	
Copy DliUib.	5,6
01 Overlay-Dliuib Redefines Dliuib.	
05 Pcbaddr usage is pointer.	
05 Filler Picture XX.	
01 Pcb-Ptrs.	
05 B-Pcb1-Ptr Picture 9(8) Comp-5.	
01 Pcb1.	7
05 Pcb1-Dbd-Name Picture X(8).	
05 Pcb1-Seg-Level Picture XX.	
05 Pcb1-Status-Code Picture XX.	
05 Pcb1-PROC-OPT Picture XXXX.	
05 Filler Picture S9(5) Comp-5.	
05 Pcb1-Seg-Name Picture X(8).	
05 Pcb1-Len-KFB Picture S9(5) Comp-5.	
05 Pcb1-NU-ENSeg Picture S9(5) Comp-5.	
05 Pcb1-KEY-FB Picture X(256).	
Procedure Division.	8
Call 'CBLTDLI' using Func-PCB Psb-Name Uib-ptr.	
If Uibfctr is not equal low-values then	
* Insert error diagnostic Code	

```

      Exec CICS Return end-Exec
End-if.
Move Uibpcbal to B-Pcb-Ptrs.
Move B-Pcb1-Ptr to Pcb1-Ptr.

*      Issue DL/I Call: get a unique segment
      Call 'CBLTDLI' using Func-GHU Pcb1
      Dli-io-area ssa1.
      Service reload Uib-ptr
      If Uibfctr is not equal Good-Return-Code then
*      Insert error diagnostic Code
      Exec CICS Return end-Exec
      End-if.

      If Pcb1-Status-Code is not equal Good-Status-Code then
*      Insert error diagnostic Code
      Exec CICS Return end-Exec
      End-if.

*      Perform segment update activity
      Move 'aaa' to area1.
      Move 'bbb' to area2.
*      Issue DL/I Call: replace segment at current position
      Call 'CBLTDLI' using Func-REPL Pcb1
      Dli-io-area ssa1.
      If Uibfctr is not equal Good-Return-Code then
*      Insert error diagnostic Code
      Exec CICS Return end-Exec
      End-if.

      If Pcb1-Status-Code is not equal Good-Status-Code then
*      Insert error diagnostic Code
      Exec CICS Return end-Exec
      End-if.

*      Release the PSB
      Call 'CBLTDLI' using Func-TERM.
      Exec CICS Return end-Exec.

```

Note:

1. You define each of the DL/I call functions the program uses with a 77-level or 01-level working storage entry. Each picture clause is defined as four alphanumeric characters and has a value assigned for each function. If you want to include the optional *parmcount* field, you can initialize count values for each type of call. You can also use the COBOL COPY statement to include these standard descriptions in the program.
2. A 9-byte area is set up for an unqualified SSA. Before the program issues a call that requires an unqualified SSA, it can either initialize this area with the segment name or move the segment name to this area. If a call requires two or more SSAs, you may need to define additional areas.
3. An 01-level working storage entry defines I/O areas that are used for passing segments to and from the database. You can further define I/O areas with 02-level entries. You can use separate I/O areas for each segment type, or you can define one I/O area to use for all segments.
4. The linkage section must start with a definition of this type to provide addressability to a parameter list that will contain the addresses of storage that is outside the working storage of the application program. The first 02-level definition is used by CICS to provide addressability to the other fields in the list. A one-to-one correspondence exists between the other 02-level names and the 01-level data definitions in the linkage section.
5. The COPY DLIUIB statement will be expanded.
6. The UIB returns the address of an area that contains the PCB addresses. The definition of PCB pointers is necessary to obtain the actual PCB addresses. Do not alter the addresses in the area.
7. The PCBs are defined in the linkage section.
8. The PCB call schedules a PSB for your program to use.
9. This unqualified GHU call retrieves a segment from the database and places it in the I/O area that is referenced by the call. Before issuing the call, the program must initialize the key or data value of the SSA so that it specifies the particular segment to be retrieved.

10. CICS online programs should test the return code in the UIB before testing the status code in the DB PCB.
11. The REPL call replaces the segment that was retrieved in the most recent Get Hold call with the data that the program has placed in the I/O area.
12. The TERM call terminates the PSB that the program scheduled earlier. This call is optional and is only issued if a sync point is desired prior to continued processing.
13. The program issues the EXEC CICS RETURN statement when it has finished its processing. If this is a return from the highest-level CICS program, a TERM call and sync point are internally generated by CICS.

Related reading: For more information about installing application programs, see *CICS Transaction Server for z/OS CICS Application Programming Guide*.

Related reference

[“Specifying the UIB \(CICS online programs only\)” on page 219](#)

The interface between your CICS online program and DL/I passes additional information to your program in a user interface block (UIB). The UIB contains the address of the PCB list and any return codes your program must examine before checking the status code in the DB PCB.

Coding a program in Java

IMS provides support for developing applications using the Java programming language.

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

Related concepts

[“IMS solutions for Java development overview” on page 553](#)

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

Coding a batch program in Pascal

The following code sample is a skeleton batch program in Pascal. It shows you how the parts of an IMS program that is written in Pascal fit together. The numbers to the right of the program refer to the notes that follow the program.

Restriction: Pascal is not supported by CICS.

<pre> segment PASCIMS; type CHAR2 = packed array [1..2] of CHAR; CHAR4 = packed array [1..4] of CHAR; CHAR6 = packed array [1..6] of CHAR; CHARn = packed array [1..n] of CHAR; DB_PCB_TYPE = record DB_NAME : ALFA; DB_SEG_LEVEL : CHAR2; DB_STAT_CODE : CHAR2; DB_PROC_OPT : CHAR4; FILLER : INTEGER; DB_SEG_NAME : ALFA; DB_LEN_KFB : INTEGER; DB_NO_SENSESEG : INTEGER; DB_KEY_FB : CHARn; end; procedure PASCIMS (var SAVE: INTEGER; var DB_PCB_MAST: DB_PCB_TYPE; var DB_PCB_DETAIL : DB_PCB_TYPE); REENTRANT; procedure PASCIMS; type QUAL_SSA_TYPE = record SEQ_QUAL : CHAR; SEG_KEY_NAME : ALFA; SEG_OPR : CHAR2; SEG_NAME : ALFA; end; </pre>	NOTES 1 2 3 4 5


```

        SEG_KEY_VALUE: CHAR6;
        SEG_END_CHAR : CHAR;
    end;
    MAST_SEG_IO_AREA_TYPE = record
        (* Field declarations *)
    end;
    DET_SEG_IO_AREA_TYPE = record
        (* Field declarations *)
    end;
var
    MAST_SEG_IO_AREA : MAST_SEG_IO_AREA_TYPE;
    DET_SEG_IO_AREA : DET_SEG_IO_AREA_TYPE;
const
    GU   = 'GU   ';
    GN   = 'GN   ';
    GHU  = 'GHU  ';
    GHN  = 'GHN  ';
    GHNP = 'GHNP';
    ISRT = 'ISRT';
    REPL = 'REPL';
    DLET = 'DLET';
    QUAL_SSA = QUAL_SSA_TYPE('ROOT', '(', 'KEY', ' ', '=',
                             'vvvvv', ')');
    UNQUAL_SSA = 'NAME   ';
procedure PASTDLI; GENERIC;
begin
    PASTDLI(const GU,
            var DB_PCB_DETAIL;
            var DET_SEG_IO_AREA;
            const QUAL_SSA);
    PASTDLI(const GHU,
            var DB_PCB_MAST,
            var MAST_SEG_IO_AREA,
            const QUAL_SSA);
    PASTDLI(const GHN,
            var DB_PCB_MAST,
            var MAST_SEG_IO_AREA);
    PASTDLI(const REPL,
            var DB_PCB_MAST,
            var MAST_SEG_IO_AREA);
end;

```

Note:

1. Define the name of the Pascal compile unit.
2. Define the data types that are needed for the PCBs used in your program.
3. Define the PCB data type that is used in your program.
4. Declare the procedure heading for the REENTRANT procedure that is called by IMS. The first word in the parameter list should be an INTEGER, which is reserved for VS Pascal's usage. The rest of the parameters are the addresses of the PCBs that are received from IMS.
5. Define the data types that are needed for the SSAs and I/O areas.
6. Declare the variables used for the I/O areas.
7. Define the constants, such as function codes and SSAs that are used in the PASTDLI DL/I calls.
8. Declare the IMS interface routine by using the GENERIC directive. GENERIC identifies external routines that allow multiple parameter list formats. A GENERIC routine's parameters are "declared" only when the routine is called.
9. This call retrieves data from the database. It contains a qualified SSA. Before you can issue a call that uses a qualified SSA, you must initialize the data field of the SSA. Before you can issue a call that uses an unqualified SSA, you must initialize the segment name field.
10. This is another call that has a qualified SSA.
11. This call is an unqualified call that retrieves data from the database. Because it is a Get Hold call, it can be followed by a REPL or DLET call.
12. The REPL call replaces the data in the segment that was retrieved by the most recent Get Hold call; the data is replaced by the contents of the I/O area that is referenced in the call.

13. You return control to IMS by exiting from the PASCIMS procedure. You can also code a RETURN statement to exit at another point.

Restriction: You must bind your program to the IMS language interface module (DFSLI000) after compiling your program.

Coding a batch program in PL/I

The following code example is a skeleton batch program in PL/I. It shows you how the parts of an IMS program that is written in PL/I fit together.

The numbers to the right of the program refer to the notes that follow. This kind of program can run as a batch program or as a batch-oriented BMP.

Restriction: IMS application programs cannot use PL/I multitasking. This is because all tasks operate as subtasks of a PL/I control task when you use multitasking.

Sample PL/I program

		NOTES
/*	*/	
/*	ENTRY POINT	*/
/*	*/	
DLITPLI: PROCEDURE (IO_PTR_PCB,DB_PTR_MAST,DB_PTR_DETAIL)		1
OPTIONS (MAIN);		
/*	*/	
/*	DESCRIPTIVE STATEMENTS	*/
/*	*/	
DCL IO_PTR_PCB POINTER;		
DCL DB_PTR_MAST POINTER;		
DCL DB_PTR_DETAIL POINTER;		
DCL FUNC_GU CHAR(4) INIT('GU ');		2
DCL FUNC_GN CHAR(4) INIT('GN ');		
DCL FUNC_GHU CHAR(4) INIT('GHU ');		
DCL FUNC_GHN CHAR(4) INIT('GHN ');		
DCL FUNC_GNP CHAR(4) INIT('GNP ');		
DCL FUNC_GHNP CHAR(4) INIT('GHNP');		
DCL FUNC_ISRT CHAR(4) INIT('ISRT');		
DCL FUNC_REPL CHAR(4) INIT('REPL');		
DCL FUNC_DLET CHAR(4) INIT('DLET');		
DCL 1 QUAL_SSA	STATIC UNALIGNED,	3
2 SEG_NAME	CHAR(8) INIT('ROOT '),	
2 SEG_QUAL	CHAR(1) INIT('('),	
2 SEG_KEY_NAME	CHAR(8) INIT('KEY '),	
2 SEG_OPR	CHAR(2) INIT('= '),	
2 SEG_KEY_VALUE	CHAR(6) INIT('vvvvv'),	
2 SEG_END_CHAR	CHAR(1) INIT(')');	
DCL 1 UNQUAL_SSA	STATIC UNALIGNED,	
2 SEG_NAME_U	CHAR(8) INIT('NAME '),	
2 BLANK	CHAR(1) INIT(' ');	4
DCL 1 MAST_SEG_IO_AREA,		
2 —		
2 —		
2 —		
DCL 1 DET_SEG_IO_AREA,		
2 —		
2 —		
2 —		
DCL 1 IO_PCB	BASED (IO_PTR_PCB),	5
2 FILLER	CHAR(10),	
2 STAT	CHAR(2);	
DCL 1 DB_PCB_MAST	BASED (DB_PTR_MAST),	
2 MAST_DB_NAME	CHAR(8),	
2 MAST_SEG_LEVEL	CHAR(2),	
2 MAST_STAT_CODE	CHAR(2),	
2 MAST_PROC_OPT	CHAR(4),	
2 FILLER	FIXED BINARY (31,0),	
2 MAST_SEG_NAME	CHAR(8),	
2 MAST_LEN_KFB	FIXED BINARY (31,0),	
2 MAST_NO_SENSESEG	FIXED BINARY (31,0),	
2 MAST_KEY_FB	CHAR(*);	
DCL 1 DB_PCB_DETAIL	BASE (DB_PTR_DETAIL),	
2 DET_DB_NAME	CHAR(8),	
2 DET_SEG_LEVEL	CHAR(2),	
2 DET_STAT_CODE	CHAR(2),	
2 DET_PROC_OPT	CHAR(4),	
2 FILLER	FIXED BINARY (31,0),	

```

2 DET_SEG_NAME      CHAR(8),
2 DET_LEN_KFB       FIXED BINARY (31,0),
2 DET_NO_SENSESEG   FIXED BINARY (31,0),
2 DET_KEY_FB        CHAR(*);
DCL THREE          FIXED BINARY (31,0) INITIAL(3);
DCL FOUR           FIXED BINARY (31,0) INITIAL(4);
DCL FIVE           FIXED BINARY (31,0) INITIAL(5);
DCL SIX            FIXED BINARY (31,0) INITIAL(6);
/*
/*      MAIN PART OF PL/I BATCH PROGRAM      */
/*
CALL PLITDLI (FOUR, FUNC_GU, DB_PCB_DETAIL, DET_SEG_IO_AREA, QUAL_SSA);
IF DET_STAT_CODE = GOOD_STATUS_CODE THEN DO;
    CALL PLITDLI (FOUR, FUNC_GHU, DB_PCB_MAST, MAST_SEG_IO_AREA, QUAL_SSA);
    IF MAST_STAT_CODE = GOOD_STATUS_CODE THEN DO;
        CALL PLITDLI (THREE, FUNC_GHN, DB_PCB_MAST, MAST_SEG_IO_AREA);
        IF MAST_STAT_CODE = GOOD_STATUS_CODE THEN DO;
            CALL PLITDLI (THREE, FUNC_REPL, DB_PCB_MAST, MAST_SEG_IO_AREA);
            IF MAST_STAT_CODE ^= GOOD_STATUS_CODE THEN DO;
                /* INSERT REPLACE DIAGNOSTIC MESSAGE */
            END;
        END;
    ELSE DO;
        /* INSERT GHN DIAGNOSTIC MESSAGE */
    END;
END;
ELSE DO;
    /* INSERT GHU DIAGNOSTIC MESSAGE */
END;
ELSE DO;
    /* INSERT GU DIAGNOSTIC MESSAGE */
END;
RETURN;
END DLITPLI;

```

Note:

1. After IMS has loaded the PSB of the application program, IMS gives control to the application program through this entry point. PL/I programs must pass the pointers to the PCBs, not the names, in the entry statement. The entry statement lists the PCBs that the program uses by the names that it has assigned to the definitions for the PCB masks. The order in which you refer to the PCBs in the entry statement must be the same order in which they have been defined in the PSB.

The code example assumes that an I/O PCB was passed to the application program. When the program is a batch program, CMPAT=YES must be specified on the PSBGEN statement of PSBGEN so that the I/O PCB is included. Because the I/O PCB is required for a batch program to make system service calls, CMPAT=YES should always be specified for batch programs.

2. Each of these areas defines one of the call functions used by the batch program. Each character string is defined as four alphanumeric characters, with a value assigned for each function. You can define other constants in the same way. Also, you can store standard definitions in a source library and include them by using a %INCLUDE statement.
3. A structure definition defines each SSA the program uses. The unaligned attribute is required for SSAs. The SSA character string must reside contiguously in storage. You should define a separate structure for each qualified SSA, because the value of the data field for each SSA is different.
4. The I/O areas that are used to pass segments to and from the database are defined as structures.
5. Level-01 declaratives define masks for the PCBs that the program uses as structures. These definitions make it possible for the program to check fields in the PCBs.
6. This statement defines the *parmcount* that is required in DL/I calls that are issued from PL/I programs (except for the call to the sample status-code error routine, where it is not allowed). The *parmcount* is the address of a 4-byte field that contains the number of subsequent parameters in the call. The *parmcount* is required only in PL/I programs. It is optional in the other languages. The value in *parmcount* is binary. This example shows how you can code the *parmcount* parameter when three parameters follow in the call:

```
DCL THREE          FIXED BINARY (31,0) INITIAL(3);
```

7. This call retrieves data from the database. It contains a qualified SSA. Before you can issue a call that uses a qualified SSA, initialize the data field of the SSA. Before you can issue a call that uses an unqualified SSA, initialize the segment name field. Check the status code after each DL/I call that you issue.

Although you must declare the PCB parameters that are listed in the entry statement to a PL/I program as POINTER data types, you can pass either the PCB name or the PCB pointer in DL/I calls in a PL/I program.

8. This is another call that has a qualified SSA.
9. This is an unqualified call that retrieves data from the database. Because it is a Get Hold call, it can be followed by REPL or DLET.
10. The REPL call replaces the data in the segment that was retrieved by the most recent Get Hold call; the data is replaced by the contents of the I/O area referenced in the call.
11. The RETURN statement returns control to IMS.

Binding PL/I code to the IMS language interface module

IMS provides a language interface module (DFSII000) which gives a common interface to IMS. This module must be bound to the program.

If you use the IMS-supplied procedures (IMSPLI or IMSPLIGO), IMS binds the language interface module to the application program. IMSPLI is a two-step procedure that compiles and binds your program. IMSPLIGO is a three-step procedure that compiles, binds, and executes your program in a DL/I batch region. For information on how to use these procedures, see *IMS Version 15.3 System Definition*.

Coding a CICS online program in PL/I

The following code example is a skeleton CICS online program in PL/I. It shows you how to define and establish addressability to the UIB.

The numbers to the right of the program refer to the notes that follow. This kind of program can run in a CICS environment using DBCTL.

Sample call-level PL/I program (CICS online)

PLIUIB: PROC OPTIONS(MAIN);	NOTES
DCL PSB_NAME CHAR(8) STATIC INIT('PLIPSB ');	1
DCL PCB_FUNCTION CHAR(4) STATIC INIT('PCB ');	
DCL TERM_FUNCTION CHAR(4) STATIC INIT('TERM');	
DCL GHU_FUNCTION CHAR(4) STATIC INIT('GHU ');	
DCL REPL_FUNCTION CHAR(4) STATIC INIT('REPL');	
DCL SSA1 CHAR(9) STATIC INIT('AAAA4444 ');	2
DCL PARM_CT_1 FIXED BIN(31) STATIC INIT(1);	
DCL PARM_CT_3 FIXED BIN(31) STATIC INIT(3);	
DCL PARM_CT_4 FIXED BIN(31) STATIC INIT(4);	
DCL GOOD_RETURN_CODE BIT(8) STATIC INIT('0'B);	
DCL GOOD_STATUS_CODE CHAR(2) STATIC INIT(' ');	
%INCLUDE DLIUIB;	3
DCL 1 PCB_POINTERS BASED(UIBPCBAL),	4
2 PCB1_PTR POINTER;	
DCL 1 DLI_IO_AREA,	5
2 AREA1 CHAR(3),	
2 AREA2 CHAR(37);	
DCL 1 PCB1 BASED(PCB1_PTR),	6
2 PCB1_DBD_NAME CHAR(8),	
2 PCB1_SEG_LEVEL CHAR(2),	
2 PCB1_STATUS_CODE CHAR(2),	
2 PCB1_PROC_OPTIONS CHAR(4),	
2 PCB1_RESERVE_DLI FIXED BIN (31,0),	
2 PCB1_SEGNAME_FB CHAR(8),	
2 PCB1_LENGTH_FB_KEY FIXED BIN(31,0),	
2 PCB1_NUMB_SENS_SEGS FIXED BIN(31,0),	
2 PCB1_KEY_FB_AREA CHAR(17);	
/* SCHEDULE PSB AND OBTAIN PCB ADDRESSES */	
CALL PLITDLI (PARM_CT_3,PCB_FUNCTION,PSB_NAME,UIBPTR);	7
IF UIBFCTR = GOOD_RETURN_CODE THEN DO;	
/* ISSUE DL/I CALL: GET A UNIQUE SEGMENT */	
CALL PLITDLI (PARM_CT_4,GHU_FUNCTION,PCB1,DLI_IO_AREA,SSA1);	8

```

IF UIBFCTR = GOOD_RETURN_CODE& PCB1_STATUS_CODE = GOOD_STATUS_CODE THEN DO;    9
  /* PERFORM SEGMENT UPDATE ACTIVITY */
  AREA1 = .....;
  AREA2 = .....;
  /* ISSUE DL/I: REPLACE SEGMENT AT CURRENT POSITION */
  PLITDLI (PARM_CT_3,REPL_FUNCTION,PCB1,DLI_IO_AREA);    10
  IF UIBFCTR ^= GOOD_RETURN_CODE
    | PCB1_STATUS_CODE ^= GOOD_STATUS_CODE THEN DO;
    /* INSERT REPL ERROR DIAGNOSTIC CODE */
  END;
END;
ELSE DO;
  /* INSERT GHU ERROR DIAGNOSTIC CODE */
END;
END;
ELSE DO;
  /* ANALYZE UIB PROBLEM */
  /* ISSUE UIB DIAGNOSTIC MESSAGE */
END;
/* RELEASE THE PSB */
CALL PLITDLI(PARM_CT_1,TERM_FUNCTION);    11
EXEC CICS RETURN;    12
END PLIUIB;

```

Note:

1. Each of these areas defines the DL/I call functions the program uses. Each character string is defined as four alphanumeric characters and has a value assigned for each function. You can define other constants in the same way. You can store standard definitions in a source library and include them by using a %INCLUDE statement.
2. A structure definition defines each SSA the program uses. The unaligned attribute is required for SSA. The SSA character string must reside contiguously in storage. If a call requires two or more SSA, you may need to define additional areas.
3. The %INCLUDE DLIUIB statement will be expanded.
4. The UIB returns the address of an area containing the PCB addresses. The definition of PCB pointers is necessary to obtain the actual PCB addresses. Do not alter the addresses in the area.
5. The I/O areas that are used to pass segments to and from the database are defined as structures.
6. The PCBs are defined based on the addresses that are passed in the UIB.
7. The PCB call schedules a PSB for your program to use.
8. This unqualified GHU call retrieves a segment from the database. The segment is placed in the I/O area that is referenced in the call. Before issuing the call, the program must initialize the key or data value of the SSA so that it specifies the particular segment to be retrieved.
9. CICS online programs must test the return code in the UIB before testing the status code in the DB PCB.
10. The REPL call replaces the segment that was retrieved in the most recent Get Hold call. The I/O area that is referenced in the call contains the segment to be replaced.
11. The TERM call terminates the PSB that the program scheduled earlier.
12. The program issues the EXEC CICS RETURN statement when it has finished processing.

Related reading: For more information about installing application programs, see *CICS Transaction Server for z/OS CICS Application Programming Guide*.

Related reference

[“Specifying the UIB \(CICS online programs only\)” on page 219](#)

The interface between your CICS online program and DL/I passes additional information to your program in a user interface block (UIB). The UIB contains the address of the PCB list and any return codes your program must examine before checking the status code in the DB PCB.

Chapter 12. Defining application program elements for IMS DB

Use these specific parameters and formats for making DL/I calls through the language interfaces for your applications program written in assembler language, C language, COBOL, Pascal, and PL/I.

Formatting DL/I calls for language interfaces

When you use DL/I calls in assembler language, C language, COBOL, Pascal, or PL/I, you must call the DL/I language interface to initiate the functions specified with the DL/I calls.

IMS offers several interfaces for DL/I calls:

- A language-independent interface for any programs that are Language Environment® conforming (CEETDLI)
- Language-specific interfaces for all supported languages (xxxTDLI)
- A non-language-specific interface for all supported languages (AIBTDLI)

Java makes use of the all three DL/I language interfaces, but the usage is internal and no calls are necessary to initiate the functions specified with the DL/I calls.

Related concepts

[“IMS solutions for Java development overview” on page 553](#)


You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.


Assembler language application programming

Application programs in assembly language use the following format, parameters, and DL/I calls to communicate with IMS databases.


In assembler language programs, all DL/I call parameters that are passed as addresses can be passed in a register, which, if used, must be enclosed in parentheses.

Format






B



C



Notes:

Parameters

parmcount

Specifies the address of a 4-byte field in user-defined storage that contains the number of parameters in the parameter list that follows *parmcount*. Assembler language application programs must use either *parmcount* or VL.

function

Specifies the address of a 4-byte field in user-defined storage that contains the call function. The call function must be left-justified and padded with blanks (such as GUbB).

db pcb

Specifies the address of the database PCB to be used for the call. The PCB address must be one of the PCB addresses passed on entry to the application program in the PCB list.

tp pcb

Specifies the address of the I/O PCB or alternate PCB to be used for the call. The PCB address must be one of the PCB addresses passed on entry to the application program in the PCB list.

aib

Specifies the address of the application interface block (AIB) in user-defined storage.

i/o area

Specifies the address of the I/O area in user-defined storage that is used for the call. The I/O area must be large enough to contain the returned data.

i/o area length

Specifies the address of a 4-byte field in user-defined storage that contains the I/O area length (specified in binary).

area length

Specifies the address of a 4-byte field in user-defined storage that contains the length (specified in binary) of the area immediately following it in the parameter list. Up to seven area lengths or area pairs can be specified.

area

Specifies the address of the area in user-defined storage to be checkpointed. Up to seven area lengths or area pairs can be specified.

token

Specifies the address of a 4-byte field in user-defined storage that contains a user token.

stat function

Specifies the address of a 9-byte field in user-defined storage that contains the stat function to be performed.

ssa

Specifies the address in user-defined storage that contains the SSAs to be used for the call. Up to 15 SSAs can be specified, one of which is *rootssa*.

rootssa

Specifies the address of a root segment search argument in user-defined storage.

rsa

Specifies the address of the area in user-defined storage that contains the record search argument.

psb name

Specifies the address in user-defined storage of an 8-byte PSB name to be used for the call.

uibptr

Specifies the address in user-defined storage of the user interface block (UIB).

sysserve

Specifies the address of an 8-byte field in user-defined storage to be used for the call.

VL

Signifies the end of the parameter list. Assembler language programs must use either *parmcoun*t or VL.

Example of a DL/I call format

Using the DL/I AIBTDLI interface:

```
CALL    AIBTDLI,(function,aib,i/o area,ssa1),VL
```

Using the DL/I language-specific interface:

```
CALL    ASMTDLI,(function,db pcb,i/o area,ssa1),VL
```

Related concepts

[“AIBTDLI interface” on page 226](#)

Use AIBTDLI as the interface between your application program and IMS.

Related reference


[DL/I calls for database management \(Application Programming APIs\)](#)

[DL/I calls for IMS DB system services \(Application Programming APIs\)](#)


C language application programming

Application programs in C use the following format, parameters, and DL/I calls to communicate with IMS databases.


Format




A



B



C



Notes:

¹ For AIBTDLI, `parmcount` is required for C applications.

Parameters

rc

This parameter receives the DL/I status or return code. It is a two-character field shifted into the 2 low-order bytes of an integer variable (int). If the status code is two blanks, 0 is placed in the field. You can test the **rc** parameter with an **if** statement. For example, **if (rc == 'IX')**. You can also use **rc** in a **switch** statement. You can choose to ignore the value placed in **rc** and use the status code returned in the PCB instead.

parmcount

Specifies the name of a fixed binary (31) variable in user-defined storage that contains the number of parameters in the parameter list that follows **parmcount**.

function

Specifies the name of a character (4) variable, left justified in user-defined storage, that contains the call function to be used. The call function must be left-justified and padded with blanks (such as **GUbb**)

db pcb

Specifies the name of a pointer variable that contains the address of the database to be used for the call. The PCB address must be one of the PCB addresses passed on entry to the application program in the PCB list.

tp pcb

Specifies the name of a pointer variable that contains the address of the I/O PCB or alternate PCB to be used for the call. The PCB address must be one of the PCB addresses passed on entry to the application program in the PCB list.

aib

Specifies the name of the pointer variable that contains the address of the structure that defines the application interface block (AIB) in user-defined storage.

i/o area

Specifies the name of a pointer variable to a major structure, array, or character string that defines the I/O area in user-defined storage used for the call. The I/O area must be large enough to contain all of the returned data.

i/o area length

Specifies the name of a fixed binary (31) variable in user-defined storage that contains the I/O area length.

area length

Specifies the name of a fixed binary (31) variable in user-defined storage that contains the length of the area immediately following it in the parameter list. Up to seven area lengths or area pairs can be specified.

area

Specifies the name of the pointer variable that contains the address of the structure that defines the user-defined storage to be checkpointed. Up to seven area lengths or area pairs can be specified.

token

Specifies the name of a character (4) variable in user-defined storage that contains a user token.

stat function

Specifies the name of a character (9) variable in user-defined storage that contains the stat function to be performed.

ssa

Specifies the name of a character variable in user-defined storage that contains the SSAs to be used for the call. Up to 15 SSAs can be specified, one of which is *rootssa*.

rootssa

Specifies the name of a character variable that defines the root segment search argument in user-defined storage.

rsa

Specifies the name of a character variable that contains the record search argument for a GU call or where IMS should return the *rsa* for an ISRT or GN call.

psb name

Specifies the name of a character (8) variable containing the PSB name to be used for the call.

uibptr

Specifies the name of a pointer variable that contains the address of the structure that defines the user interface block (UIB) that is used in user-defined storage.

sysserve

Specifies the name of a character (8) variable string in user-defined storage to be used for the call.

I/O area

In C, the I/O area can be of any type, including structures or arrays. The **ctdli** declarations in **ims.h** do not have any prototype information, so no type checking of the parameters is done. The area may be **auto**, **static**, or allocated (with **malloc** or **calloc**). You need to give special consideration to C-strings because DL/I does not recognize the C convention of terminating strings with nulls (' \0 '). Instead of the usual **strcpy** and **strcmp** functions, you may want to use **memcpy** and **memcmp**.

Example of a DL/I call format**Using the DL/I CEETDLI interface:**

```
#include <leawi.h>
...
CEETDLI (function,db pcb,i/o area,ssa1);
```

Using the DL/I AIBTDLI interface:

```
int rc;
...
rc=AIBTDLI (parmcount,function,aib,i/o area,ssa1);
```

Using the DL/I language-specific interface:

```
#include <ims.h>
int rc;
...
rc=CTDLI (function,db pcb,i/o area,ssa1);
```

Related concepts

[“AIBTDLI interface” on page 226](#)

Use AIBTDLI as the interface between your application program and IMS.

Related reference

[DL/I calls for database management \(Application Programming APIs\)](#)


[DL/I calls for IMS DB system services \(Application Programming APIs\)](#)

COBOL application programming


Application programs in COBOL use the following format, parameters, and DL/I calls to communicate with IMS databases.

Format


➤ CALL ➔




A



B



C



Note: All apostrophes (') can be replaced by quotation marks (") and can be done regardless of the APOST/QUOTE compiler (or CICS translator) option.

Parameters

parmcount

Specifies the identifier of a usage binary (4) byte data item in user-defined storage that contains the number of parameters in the parameter list that follows *parmcount*. If you define this field as COMP-5 rather than COMP, COMP-4, or BINARY, then it can contain the maximum possible values regardless of the COBOL TRUNC compiler option setting.

function

Specifies the identifier of a usage display (4) byte data item, left justified in user-defined storage that contains the call function to be used. The call function must be left-justified and padded with blanks (such as GUBb).

db pcb

Specifies the identifier of the database PCB group item from the PCB list that is passed to the application program on entry. This identifier will be used for the call.

tp pcb

Specifies the identifier of the I/O PCB or alternate PCB group item from the PCB list that is passed to the application program on entry. This identifier will be used for the call.

aib

Specifies the identifier of the group item that defines the application interface block (AIB) in user-defined storage.

i/o area

Specifies the identifier of a major group item, table, or usage display data item that defines the I/O area length in user-defined storage used for the call. The I/O area must be large enough to contain all of the returned data.

i/o area length

Specifies the identifier of a usage binary (4) byte data item in user-defined storage that contains the I/O area length (specified in binary). If you define this field as COMP-5 rather than COMP, COMP-4, or BINARY, then it can contain the maximum possible values regardless of the COBOL TRUNC compiler option setting.

area length

Specifies the identifier of a usage binary (4) byte data item in user-defined storage that contains the length (specified in binary) of the area immediately following it in the parameter list. Up to seven area lengths or area pairs can be specified. If you define this field as COMP-5 rather than COMP, COMP-4, or BINARY, then it can contain the maximum possible values regardless of the COBOL TRUNC compiler option setting.

area

Specifies the identifier of the group item that defines the user-defined storage to be checkpointed. Up to seven area lengths or area pairs can be specified.

token

Specifies the identifier of a usage display (4) byte data item in user-defined storage that contains a user token.

stat function

Specifies the identifier of a usage display (9) byte data item in user-defined storage that contains the stat function to be performed.

ssa

Specifies the identifier of a usage display data item in user-defined storage that contains the SSAs to be used for the call. Up to 15 SSAs can be specified, one of which is *rootssa*.

rootssa

Specifies the identifier of a usage display data item that defines the root segment search argument in user-defined storage.

rsa

Specifies the identifier of a usage display data item that contains the record search argument.

psb name

Specifies the identifier of a usage display (8) byte data item containing the PSB name to be used for the call.

uibptr

Specifies the identifier of the group item that defines the user interface block (UIB) that is used in user-defined storage.

sysserve

Specifies the identifier of a usage display (8) byte data item in user-defined storage to be used for the call.

Example of a DL/I call format**Using the DL/I CEETDLI interface:**

```
CALL      'CEETDLI' USING function,db pcb,i/o area,ssa1.
```

Using the DL/I AIBTDLI interface:

```
CALL      'AIBTDLI' USING function,aib,i/o area,ssa1.
```

Using the DL/I language-specific interface:

```
CALL      'CBLTDLI' USING function,db pcb,i/o area,ssa1.
```

Related reference

[DL/I calls for database management \(Application Programming APIs\)](#)

[DL/I calls for IMS DB system services \(Application Programming APIs\)](#)

Java application programming for IMS

IMS provides support for developing applications using the Java programming language.

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

Related concepts


[“IMS solutions for Java development overview” on page 553](#)

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.


Pascal application programming

Application programs in Pascal use the following format, parameters, and DL/I calls to communicate with IMS databases.


Format




A




B



C



D



Parameters

parmcount

Specifies the name of a fixed binary (31) variable in user-defined storage that contains the number of parameters in the parameter list that follows parmcount.

function

Specifies the name of a character (4) variable, left justified in user-defined storage, that contains the call function to be used. The call function must be left-justified and padded with blanks (such as GUBb).

db pcb

Specifies the name of a pointer variable that contains the address of the database PCB defined in the call procedure statement.

tp pcb

Specifies the name of a pointer variable that contains the address of the I/O PCB or alternate PCB defined in the call procedure statement.

aib

Specifies the name of the pointer variable that contains the address of the structure that defines the application interface block (AIB) in user-defined storage.

i/o area

Specifies the name of a pointer variable to a major structure, array, or character string that defines the I/O area in user-defined storage used for the call. The I/O area must be large enough to contain all of the returned data.

i/o area length

Specifies the name of a fixed binary (31) variable in user-defined storage that contains the I/O area length.

area length

Specifies the name of a fixed binary (31) variable in user-defined storage that contains the length of the area immediately following it in the parameter list. Up to seven area lengths or area pairs can be specified.

area

Specifies the name of the pointer variable that contains the address of the structure that defines the user-defined storage to be checkpointed. Up to seven area lengths or area pairs can be specified.

token

Specifies the name of a character (4) variable in user-defined storage that contains a user token.

stat function

Specifies the name of a character (9) variable in user-defined storage that contains the stat function to be performed.

ssa

Specifies the name of a character variable in user-defined storage that contains the SSAs to be used for the call. Up to 15 SSAs can be specified, one of which is *rootssa*.

rootssa

Specifies the name of a character variable that defines the root segment search argument in user-defined storage.

rsa

Specifies the name of a character variable that contains the record search argument.

psb name

Specifies the name of a character (8) variable containing the PSB name to be used for the call.

uibptr

Specifies the name of a pointer variable that contains the address of the structure that defines the user interface block (UIB) that is used in user-defined storage.

sysserve

Specifies the name of a character (8) variable string in user-defined storage to be used for the call.

Example of a DL/I call format

Using the DL/I AIBTDLI interface:

```
AIBTDLI(CONST function,  
        VAR aib,  
        VAR i/o area,  
        VAR ssa1);
```

Using the DL/I language-specific interface:

```
PASTDLI(CONST function,  
        VAR db pcb,  
        VAR i/o area,  
        VAR ssa1);
```

Related reference

[DL/I calls for database management \(Application Programming APIs\)](#)


[DL/I calls for IMS DB system services \(Application Programming APIs\)](#)

Application programming for PL/I


Application programs in PL/I use the following format, parameters, and DL/I calls to communicate with IMS databases.

Restriction: For the PLITDLI interface, all parameters except *parmcount* are indirect pointers; for the AIBTDLI interface, all parameters are direct pointers.


Format




A



B



C



Parameters

parmcount

Specifies the name of a fixed binary (31-bit) variable that contains the number of arguments that follow *parmcount*.

function

Specifies the name of a fixed-character (4-byte) variable left-justified, blank padded character string containing the call function to be used (such as GUbB).

db pcb

Specifies the structure associated with the database PCB to be used for the call. This structure is based on a PCB address that must be one of the PCB addresses passed on entry to the application program.

tp pcb

Specifies the structure associated with the I/O PCB or alternate PCB to be used for the call.

aib

Specifies the name of the structure that defines the AIB in your application program.

i/o area

Specifies the name of the I/O area used for the call. The I/O area must be large enough to contain all the returned data.

i/o area length

Specifies the name of a fixed binary (31) variable that contains the I/O area length.

area length

Specifies the name of a fixed binary (31) variable that contains the length of the area immediately following it in the parameter list. Up to seven area lengths or area pairs can be specified.

area

Specifies the name of the area to be checkpointed. Up to seven area lengths or area pairs can be specified.

token

Specifies the name of a character (4) variable that contains a user token.

stat function

Specifies the name of a character (9) variable string containing the stat function to be performed.

ssa

Specifies the name of a character variable that contains the SSAs to be used for the call. Up to 15 SSAs can be specified, one of which is *rootssa*.

rootssa

Specifies the name of a character variable that contains a root segment search argument.

rsa

Specifies the name of a character variable that contains the record search argument.

psb name

Specifies the name of a character (8) containing the PSB name to be used for the call.

uibptr

Specifies the name of the user interface block (UIB).

sysserve

Specifies the name of a character (8) variable character string to be used for the call.

Example of a DL/I call format**Using the DL/I CEETDLI interface:**

```
CALL    CEETDLI (parmcount,function,db pcb,i/o area,ssa1);
```

Using the DL/I AIBTDLI interface:

```
CALL    AIBTDLI (parmcount,function,aib,i/o area,ssa1);
```

Using the DL/I language-specific interface:

```
%INCLUDE CEEIBMAW;
CALL    PLITDLI (parmcount,function,db pcb,i/o area,ssa1);
```

Related reference

[DL/I calls for database management \(Application Programming APIs\)](#)

[DL/I calls for IMS DB system services \(Application Programming APIs\)](#)

Specifying the I/O PCB mask

After your program issues a call with the I/O Program Communications Block (I/O PCB), IMS returns information about the results of the call to the I/O PCB. To determine the results of the call, your program must check the information that IMS returns.

Issuing a system service call requires an I/O PCB. Because the I/O PCB resides outside your program, you must define a mask of the PCB in your program to check the results of IMS calls. The mask must contain the same fields, in the same order, as the I/O PCB. Your program can then refer to the fields in the PCB through the PCB mask.

The following table shows the fields that the I/O PCB contains, their lengths, and the applicable environment for each field.

Table 34. I/O PCB mask

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
Logical terminal name ¹	8	X		X		
Reserved for IMS ²	2	X		X		
Status code ³	2	X	X	X	X	X

Table 34. I/O PCB mask (continued)

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
8-Byte Local date and time ⁴						
Date	4	X		X		
Time	4	X		X		
Input message sequence number ⁵	4	X		X		
Message output descriptor name ⁶	8	X		X		
Userid ⁷	8	X		X		
Group name ⁸	8	X		X		
12-Byte Time Stamp ⁹						
Date	4	X		X		
Time	6	X		X		
UTC Offset	2	X		X		
Userid Indicator¹⁰	1	X		X		
Reserved for IMS²	3					

Note:

1. Logical Terminal Name

This field contains the name of the terminal that sent the message. When your program retrieves an input message, IMS places the name of the logical terminal that sent the message in this field. When you want to send a message back to this terminal, you refer to the I/O PCB when you issue the ISRT call, and IMS takes the name of the logical terminal from the I/O PCB as the destination.

2. Reserved for IMS

These fields are reserved.

3. Status Code

IMS places the status code describing the result of the DL/I call in this field. IMS updates the status code after each DL/I call that the program issues. Your program should always test the status code after issuing a DL/I call.

The three status code categories are:

- Successful status codes or status codes with exceptional but valid conditions. This category does not contain errors. If the call was completely successful, this field contains blanks. Many of the codes in this category are for information only. For example, a QC status code means that no more messages exist in the message queue for the program. When your program receives this status code, it should terminate.
- Programming errors. The errors in this category are usually ones that you can correct. For example, an AD status code indicates an invalid function code.
- I/O or system errors.

For the second and third categories, your program should have an error routine that prints information about the last call that was issued program termination. Most installations have a standard error routine that all application programs at the installation use.

4. Local Date and Time

The current local date and time are in the prefix of all input messages except those originating from non-message-driven BMPs. The local date is a packed-decimal, right-aligned date, in the format yyddd. The local time is a packed-decimal time in the format hhmmss.t. The current local date and time indicate when IMS received the entire message and enqueued it as input for the program, rather than the time that the application program received the message. To obtain the application processing time, you must use the time facility of the programming language you are using.

For a conversation, for an input message originating from a program or for a message received using Multiple System Coupling (MSC), the time and date indicate when the original message was received from the terminal.

Note: Be careful when comparing the local date and time in the I/O PCB with the current time returned by the operating system. The I/O PCB date and time may not be consistent with the current time. It may even be greater than the current time for the following reasons:

- The time stamp in the I/O PCB is the local time that the message was received by IMS. If the local time was changed after the message arrived, it is possible for the current time to appear to be earlier than the I/O PCB time. This effect would be likely to occur in the hour immediately after the fall time change, when the clock is set back by one hour.
- The time stamp in the I/O PCB is derived from an internal IMS time stamp stored with the message. This internal time stamp is in Coordinated Universal Time (UTC), and contains the time zone offset that was in effect at the time the message was enqueued. This time zone offset is added to the UTC time to obtain the local time that is placed in the I/O PCB. However, the time zone offset that is stored is only fifteen minutes. If the real time zone offset was not an integer multiple of fifteen minutes, the local time passed back in the I/O PCB will differ from the actual time by plus or minus 7.5 minutes. This could cause the I/O PCB time to be later than the current time. See *IMS Version 15.3 Operations and Automation* for further explanation.

Concerns about the value in the local time stamp in the I/O PCB can be reduced by using the extended time stamp introduced in IMS V6. The system administrator can choose the format of the extended time stamp to be either local time or UTC. In some situations, it may be advantageous for the application to request the time in UTC from the operating system and compare it to the UTC form of the extended time stamp. This is an option available in installations where there is no ETR to keep the IMS UTC offset in sync with the z/OS UTC offset over changes in local time.

5. Input Message Sequence Number

The input message sequence number is in the prefix of all input messages except those originating from non-message-driven BMPs. This field contains the sequence number IMS assigned to the input message. The number is binary. IMS assigns sequence numbers by physical terminal, which are continuous since the time of the most recent IMS startup.

6. Message Output Descriptor Name

You only use this field when you use MFS. When you issue a GU call with a message output descriptor (MOD), IMS places its name in this area. If your program encounters an error, it can change the format of the screen and send an error message to the terminal by using this field. To do this, the program must change the MOD name by including the MOD name parameter on an ISRT or PURG call.

Although MFS does not support APPC, LU 6.2 programs 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 is to be formatted.

Related reading: For more information on the MOD name and the LTERM interface, see *IMS Version 15.3 Communications and Connections*.

7. Userid

The use of this field is connected with RACF signon security. If signon is not active in the system, this field contains blanks.

If signon is active in the system, the field contains one of the following:

- The user's identification from the source terminal.
- The LTERM name of the source terminal if signon is not active for that terminal.
- The authorization ID. For batch-oriented BMPs, the authorization ID is dependent on the value specified for the BMPUSID= keyword in the DFSDCxxx PROCLIB member:
 - If BMPUSID=USERID is specified, the value from the USER= keyword on the JOB statement is used.
 - If USER= is not specified on the JOB statement, the program's PSB name is used.
 - If BMPUSID=PSBNAME is specified, or if BMPUSID= is not specified at all, the program's PSB name is used. If PSBNAME is not defined to RACF, the Userid of the current Address Space will be used; this will be the Home Dependent Region one, or the Control Region one if LSO=Y or if PARDLI=1 has been specified for the BMP. Userid of the current Address Space will be used also if DFSBSEX0 has returned RC08.

Related Reading: For more information about authorizing resource use in a dependent region, see *IMS Version 15.3 System Administration*.

8. Group Name

The group name, which is used by DB2 to provide security for SQL calls, is created through IMS transactions.

Three instances that apply to the group name are:

- If you use RACF and signon on your IMS system, the RACROUTE SAF (extract) call returns an eight-character group name.
- If you use your own security package on your IMS system, the RACROUTE SAF call returns any eight-character name from the package and treats it as a group name. If the RACROUTE SAF call returns a return code of 4 or 8, a group name was not returned, and IMS blanks out the group name field.
- If you use LU 6.2, the transaction header can contain a group name.

Related reading: See *IMS Version 15.3 Communications and Connections* for more information on LU 6.2.

9. 12-Byte Time Stamp

This field contains the current date and time fields, but in the IMS internal packed-decimal format. The time stamp has the following parts:

Date

yyyymmddf

This packed-decimal date contains the year (yyyy), day of the year (ddd), and a valid packed-decimal + sign such as (f).

Time

hhmmss.thmij

This packed-decimal time consists of hours, minutes, and seconds (hhmmss) and fractions of the second to the microsecond (thmij). No packed-decimal sign is affixed to this part of the time stamp.

UTC Offset

aaq\$

The packed-decimal UTC offset is prefixed by 4 bits of attributes (a). If the 4th bit of (a) is 0, the time stamp is UTC; otherwise, the time stamp is local time. The control region parameter,

TSR=(U/L), specified in the DFSPBxxx PROCLIB member, controls the representation of the time stamp with respect to local time versus UTC time.

The offset value (qq\$) is the number of quarter hours of offset to be added to UTC or local time to convert to local or UTC time respectively.

The offset sign (\$) follows the convention for a packed-decimal plus or minus sign.

Field 4 on the I/O PCB Mask always contains the local date and time. For a description of field 4, see the notes for the previous table.

Related reading: For a more detailed description of the internal packed-decimal time-format, see *IMS Version 15.3 System Utilities*.

10. Userid Indicator

The Userid Indicator is provided in the I/O PCB and in the response to the INQY call. The Userid Indicator contains one of the following:

- U - The user's identification from the source terminal during signon
- L - The LTERM name of the source terminal if signon is not active
- P - The PSBNAME of the source BMP or transaction
- O - Other name

The value contained in the Userid Indicator field indicates the contents of the userid field.

Specifying the DB PCB mask

IMS describes the results of the calls your program issues in the DB PCB that is referenced in the call. To determine the success or failure of the DL/I call, the application program includes a mask of the DB PCB and then references the fields of the DB PCB through the mask.

A DB PCB mask must contain the fields shown in the following table. (Your program can look at, but not change, the fields in the DB PCB.) The fields in your DB PCB mask must be defined in the same order and with the same length as the fields shown here. When you code the DB PCB mask, you also give it a name, but the name is not part of the mask. You use the name (or the pointer, for PL/I) when you reference each of the PCBs your program processes. A GSAM DB PCB mask is slightly different from other DB PCB masks.

Of the nine fields, only five are important to you as you construct the program. These are the fields that give information about the results of the call. They are the segment level number, status code, segment name, length of the key feedback area, and key feedback area. The status code is the field your program uses most often to find out whether the call was successful. The key feedback area contains the data from the segments you have specified; the level number and segment name help you determine the segment type you retrieved after an unqualified GN or GNP call, or they help you determine your position in the database after an error or unsuccessful call.

Table 35. DB PCB mask

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
Database name ¹	8	X	X		X	
Segment level number ²	2	X	X		X	
Status code ³	2	X	X		X	
Processing options ⁴	4	X	X		X	
Reserved for IMS ⁵	4	X	X		X	
Segment name ⁶	8	X	X		X	
Length of key feedback area ⁷	4	X	X		X	

Table 35. DB PCB mask (continued)

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
Number of sensitive segments ⁸	4	X	X		X	
Key feedback area ⁹	var length	X	X		X	

Note:

1. This contains the name of the database. This field is 8 bytes long and contains character data.

2. Segment Level Number

This field contains numeric character data. It is 2 bytes long and right-justified. When IMS retrieves the segment you have requested, IMS places the level number of that segment in this field. If you are retrieving several segments in a hierarchic path with one call, IMS places the number of the lowest-level segment retrieved. If IMS is unable to find the segment that you request, it gives you the level number of the last segment it encounters that satisfied your call.

3. Status Code

After each DL/I call, this field contains the two-character status code that describes the results of the DL/I call. IMS updates this field after each call and does not clear it between calls. The application program should test this field after each call to find out whether the call was successful.

When the program is initially scheduled, this field contains a data-availability status code, which indicates any possible access constraint based on segment sensitivity and processing options.

Related Reading: For more information on these status codes, see the topic "INIT Call" in *IMS Version 15.3 Application Programming APIs*.

During normal processing, four categories of status codes exist:

- Successful or exceptional but valid conditions. If the call was completely successful, this field contains blanks. Many of the codes in this category are for information only. For example, GB means that IMS has reached the end of the database without satisfying the call. This situation is expected in sequential processing and is not usually the result of an error.
- Errors in the program. For example, AK means that you have included an invalid field name in a segment search argument (SSA). Your program should have error routines available for these status codes. If IMS returns an error status code to your program, your program should terminate. You can then find the problem, correct it, and restart your program.
- I/O or system error. For example, an AO status code means that there has been an I/O error concerning OSAM, BSAM, or VSAM. If your program encounters a status code in this category, it should terminate immediately. This type of error cannot normally be fixed without a system programmer, database administrator, or system administrator.
- Data-availability status codes. These are returned only if your program has issued the INIT call indicating that it is prepared to handle such status codes. "Status Code Explanations" in *IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes* describes possible causes and corrections in more detail.

4. Processing Options

This is a 4-byte field containing a code that tells IMS what type of calls this program can issue. It is a security mechanism in that it can prevent a particular program from updating the database, even though the program can read the database. This value is coded in the PROCOPT parameter of the PCB statement when the PSB for the application program is generated. The value does not change.

5. Reserved for IMS

This 4-byte field is used by IMS for internal linkage. It is not used by the application program.

6. Segment Name

After each successful call, IMS places in this field the name of the last segment that satisfied the call. When a retrieval is successful, this field contains the name of the retrieved segment. When a retrieval is unsuccessful, this field contains the last segment along the path to the requested segment that would satisfy the call. The segment name field is 8 bytes long.

When a program is initially scheduled, the name of the database type is put in the SEGNAME field. For example, the field contains DEDB when the database type is DEDB; GSAM when the database type is GSAM; HDAM, or PHDAM when the database type is HDAM or PHDAM.

7. Length of Key Feedback Area

This is a 4-byte binary field that gives the current length of the key feedback area. Because the key feedback area is not usually cleared between calls, the program needs to use this length to determine the length of the relevant current concatenated key in the key feedback area.

8. Number of Sensitive Segments

This is a 4-byte binary field that contains the number of segment types in the database to which the application program is sensitive.

9. Key Feedback Area

At the completion of a retrieval or ISRT call, IMS places the concatenated key of the retrieved segment in this field. The length of the key for this request is given in the 4-byte field. If IMS is unable to satisfy the call, the key feedback area contains the key of the segment at the last level that was satisfied. A segment's concatenated key is made up of the keys of each of its parents and its own key. Keys are positioned left to right, starting with the key of the root segment and following the hierarchic path. IMS does not normally clear the key feedback area. IMS sets this length of the key feedback area to indicate the portion of the area that is valid at the completion of each call. Your program should not use the content of the key feedback area that is not included in the key feedback area length.

Related concepts

“Data areas in GSAM databases” on page 226

Generalized Sequential Access Method (GSAM) databases are available only to application programs that can run as batch programs, batch-oriented BMPs, transaction-oriented BMPS, or JBPs. The program communication block (PCB) mask and the record search argument (RSA) that you use in a GSAM database call have special formats.

Specifying the AIB mask

The application interface block (AIB) is used by your program to communicate with IMS, when your application does not have a PCB address or the call function does not use a PCB.

The application program can use the returned PCB address, when available, to inspect the status code in the PCB and to obtain any other information needed by the application program. The AIB mask enables your program to interpret the control block defined. The AIB structure must be defined in working storage, on a fullword boundary, and initialized according to the order and byte length of the fields as shown in the following table. The table's notes describe the contents of each field.

Table 36. AIB fields

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
AIB identifier ¹	8	X	X	X	X	X
DFS AIB allocated length ²	4	X	X	X	X	X
Subfunction code ³	8	X	X	X	X	X
Resource name 1 ⁴	8	X	X	X	X	X
Resource name 2 ⁵	8	X	X	X	X	X
Reserved ⁶	8					

Table 36. AIB fields (continued)

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
Maximum output area length ⁷	4	X	X	X	X	X
Output area length used ⁸	4	X	X	X	X	X
Resource field ⁹	4	X	X	X	X	X
Optional area length ¹⁰	4	X	X	X	X	X
Reserved ¹¹	4	X	X	X	X	X
Return code ¹²	4	X	X	X	X	X
Reason code ¹³	4	X	X	X	X	X
Error code extension ¹⁴	4	X	X	X		
Resource address 1 ¹⁵	4	X	X	X	X	X
Resource address 2 ¹⁶	4	X	X	X	X	X
Resource address 3 ¹⁷	4	X	X	X	X	X
User defined token ¹⁸	16	X	X	X	X	X
Return token ¹⁹	8	X	X		X	
Reserved ²⁰	16					

AIB Identifier (AIBID)

This 8-byte field contains the AIB identifier. You must initialize AIBID in your application program to the value DFSAIBbb before you issue DL/I calls. This field is required. When the call is completed, the information returned in this field is unchanged.

DFSAIB Allocated Length (AIBLEN)

This field contains the actual 4-byte length of the AIB as defined by your program. You must initialize AIBLEN in your application program before you issue DL/I calls. The minimum length required is 128 bytes. When the call is completed, the information returned in this field is unchanged. This field is required.

Subfunction Code (AIBSFUNC)

This 8-byte field contains the subfunction code for those calls that use a subfunction. You must initialize AIBSFUNC in your application program before you issue DL/I calls. When the call is completed, the information returned in this field is unchanged.

Resource Name (AIBRSNM1)

This 8-byte field contains the name of a resource. The resource varies depending on the call. You must initialize AIBRSNM1 in your application program before you issue DL/I calls. When the call is complete, the information returned in this field is unchanged. This field is required.

For PCB related calls where the AIB is used to pass the PCB name instead of passing the PCB address in the call list, this field contains the PCB name. The PCB name for the I/O PCB is IOPCBbb. The PCB name for other types of PCBs is defined in the PCBNAME= parameter in PSBGEN.

Resource Name 2 (AIBRSNM2)

This 8-byte field contains the name of a resource. The resource varies depending on the call. You must initialize AIBRSNM2 in your application program before you issue DL/I calls.

Reserved

This 8-byte field is reserved.

Maximum Output Area Length (AIBOALEN)

This 4-byte field contains the length of the output area in bytes that was specified in the call list. You must initialize AIBOALEN in your application program for all calls that return data to the output area. When the call is completed, the information returned in this area is unchanged.

Used Output Area Length (AIBOAUSE)

This 4-byte field contains the length of the data returned by IMS for all calls that return data to the output area. When the call is completed this field contains the length of the I/O area used for this call.

Resource Field (AIBRSFLD)

This 4-byte field contains a resource. The resource varies depending on the call. You must initialize AIBRSFLD in your application program before you issue DL/I calls.

Optional Area Length (AIBOPLEN)

This 4-byte field contains a resource. The resource varies depending on the call. You must initialize AIBOPLEN in your application program before you issue DL/I calls.

Reserved

This 4-byte field is reserved.

Return code (AIBRETRN)

When the call is completed, this 4-byte field contains the return code.

Reason Code (AIBREASN)

When the call is completed, this 4-byte field contains the reason code.

Error Code Extension (AIBERRXT)

This 4-byte field contains additional error information depending on the return code in AIBRETRN and the reason code in AIBREASN.

Resource Address 1 (AIBRSA1)

When the call is completed, this 4-byte field contains call-specific information. For PCB related calls where the AIB is used to pass the PCB name instead of passing the PCB address in the call list, this field returns the PCB address.

Resource Address 2 (AIBRSA2)

When the call is completed, this 4-byte field contains call-specific information.

Resource Address 3 (AIBRSA3)

When the call is completed, this 4-byte field contains call-specific information.

User Defined Token (AIBUTKN)

This 16-byte field contains a user defined token. The token varies depending on the call.

AIB return token (AIBRTKN)

AIB return token. This 8-byte field contains a token returned by a DL/I call. The usage is specific to the DL/I call for which the token was returned.

Reserved

This 16-byte field is reserved.

Specifying the AIB mask for ODBA applications

The following table describes the fields for specifying the application interface block (AIB) mask for ODBA applications.

The notes that follow describe the contents of each field.

Table 37. AIB fields for use of ODBA applications

AIB Fields	Byte Length	DB/DC	IMS DB	DCCTL	DB Batch	TM Batch
AIB identifier	8	X	X	X	X	X
DFSAIB allocated length	4	X	X	X	X	X
Subfunction code	8	X	X	X	X	X

Table 37. AIB fields for use of ODBA applications (continued)

AIB Fields	Byte Length	DB/DC	IMS DB	DCCTL	DB Batch	TM Batch
Resource name #1	8	X	X	X	X	X
Resource name #2	8					
Reserved 1	8	X				
Maximum output area length	4	X	X	X	X	X
Output area length used	4	X	X	X	X	X
Reserved 2	12					
Return code	4	X	X	X	X	X
Reason code	4	X	X	X	X	X
Error code extension	4	X				
Resource address #1	4	X	X	X	X	X
Resource address #2	4					
Resource address #3	4					
AIB return token	8	X	X		X	
Reserved 3	32					
Reserved for ODBA	136					

AIB Identifier (AIBID)

This 8-byte field contains the AIB identifier. You must initialize AIBID in your application program to the value DFSAIBbb before you issue DL/I calls. This field is required. When the call is completed, the information returned in this field is unchanged.

DFS AIB Allocated Length (AIBLEN)

This field contains the actual 4-byte length of the AIB as defined by your program. You must initialize AIBLEN in your application program before you issue DL/I calls. The minimum length required is 264 bytes for ODBA. When the call is completed, the information returned in this field is unchanged. This field is required.

Subfunction Code (AIBSFUNC)

This 8-byte field contains the subfunction code for those calls that use a subfunction. You must initialize AIBSFUNC in your application program before you issue DL/I calls. When the call is completed, the information returned in this field is unchanged.

Resource Name (AIBRSNM1) #1

This 8-byte field contains the name of a resource. The resource varies depending on the call. You must initialize AIBRSNM1 in your application program before you issue DL/I calls. When the call is complete, the information returned in this field is unchanged. This field is required.

For PCB related calls where the AIB is used to pass the PCB name instead of passing the PCB address in the call list, this field contains the PCB name. The PCB name for the I/O PCB is IOPCBbb. The PCB name for other types of PCBs is defined in the PCBNAME= parameter in PSBGEN.

Resource Name (AIBRSNM2) #2

Specify a 4-character ID of ODBA startup table DFSxxxx0, where xxxx is a 4-character ID.

Reserved 1

This 8-byte field is reserved.

Maximum Output Area Length (AIBOALEN)

This 4-byte field contains the length of the output area in bytes that was specified in the call list. You must initialize AIBOALEN in your application program for all calls that return data to the output area. When the call is completed, the information returned in this area is unchanged.

Used Output Area Length (AIBOAUSE)

This 4-byte field contains the length of the data returned by IMS for all calls that return data to the output area. When the call is completed this field contains the length of the I/O area used for this call.

Reserved 2

This 12-byte field is reserved.

Return code (AIBRETRN)

When the call is completed, this 4-byte field contains the return code.

Reason Code (AIBREASN)

When the call is completed, this 4-byte field contains the reason code.

Error Code Extension (AIBERRXT)

This 4-byte field contains additional error information depending on the return code in AIBRETRN and the reason code in AIBREASN.

Resource Address (AIBRSA1) #1

When the call is completed, this 4-byte field contains call-specific information. For PCB related calls where the AIB is used to pass the PCB name instead of passing the PCB address in the call list, this field returns the PCB address.

Resource Address (AIBRSA2) #2

This 4-byte field is reserved for ODBA.

Resource Address (AIBRSA3) #3

This 4-byte token, returned on the APSB call, is required for subsequent DLI calls and the DPSB call related to this thread.

AIB return token (AIBRTKN)

AIB return token. This 8-byte field contains a token returned by a DL/I call. The usage is specific to the DL/I call for which the token was returned.

Reserved 3

This 32-byte field is reserved.

Reserved for ODBA

This 136-byte field is reserved for ODBA.

The application program can use the returned PCB address, when available, to inspect the status code in the PCB and to obtain any other information needed by the application program.

COBOL AIB Mask Example

```

01  AIB.
   02  AIBRID          PIC x(8).
   02  AIBRLN          PIC 9(9) USAGE BINARY.
   02  AIBRSFUNC       PIC x(8).
   02  AIBRSNM1        PIC x(8).
   02  AIBRSNM2        PIC x(8).
   02  AIBRSNM3        PIC x(8).
   02  AIBOALEN        PIC 9(9) USAGE BINARY.
   02  AIBOAUSE        PIC 9(9) USAGE BINARY.
   02  AIBRESV2        PIC x(12).
   02  AIBRETRN        PIC 9(9) USAGE BINARY.
   02  AIBREASN        PIC 9(9) USAGE BINARY.
   02  AIBERRXT        PIC 9(9) USAGE BINARY.
   02  AIBRESA1        USAGE POINTER.
   02  AIBRESA2        USAGE POINTER.
   02  AIBRESA3        USAGE POINTER.
   02  AIBRESV4        PIC x(40).
   02  AIBRSAVE        OCCURS 18 TIMES USAGE POINTER.
   02  AIBRTOKN        OCCURS 6 TIMES USAGE POINTER.
   02  AIBRTOKC        PIC x(16).

```

```

02 AIBRTOKV          PIC x(16).
02 AIBRTOKA          OCCURS 2 TIMES PIC 9(9) USAGE BINARY.

```

Assembler AIB Mask Example

```

DFS AIB      DSECT
AIBID      DS      CL8'DFSAIB'
AIBLEN     DS      F
AIBSFUNC   DS      CL8
AIBRSNM1   DS      CL8
AIBRSVM2   DS      CL8
AIBRSNM3   DS      CL8
           DS      2F
AIBOALEN   DS      F
AIBOAUSE   DS      F
           DS      2F
           DS      H
           DS      H
AIBRETRN   DS      F
AIBREASN   DS      F
AIBRRXT    DS      F
AIBRSA1    DS      A
AIBRSA2    DS      A
AIBRSA3    DS      A
           DS      10F
AIBLL      EQU     *-DFSAIB
AIBSAVE    DS      18F
AIBTOKN    DS      6F
AIBTOKC    DS      CL16
AIBTOKV    DS      XL16
AIBTOKA    DS      2F
AIBAERL    EQU     *-DFSAIB

```

Specifying the UIB (CICS online programs only)

The interface between your CICS online program and DL/I passes additional information to your program in a user interface block (UIB). The UIB contains the address of the PCB list and any return codes your program must examine before checking the status code in the DB PCB.

When you issue the PCB call to obtain a PSB for your program, a UIB is created for your program. As with any area outside your program, you must include a definition of the UIB and establish addressability to it. CICS provides a definition of the UIB for all programming languages:

- In COBOL programs, use the COPY DLIUIB statement.
- In PL/I programs, use a %INCLUDE DLIUIB statement.
- In assembler language programs, use the DLIUIB macro.

Three fields in the UIB are important to your program: UIBPCBAL, UIBFCTR, and UIBDLTR. UIBPCBAL contains the address of the PCB address list. Through it you can obtain the address of the PCB you want to use. Your program must check the return code in UIBFCTR (and possibly UIBDLTR) before checking the status code in the DB PCB. If the contents of UIBFCTR and UIBDLTR are not null, the content of the status code field in the DB PCB is not meaningful. The return codes are described in the topic "CICS-DL/I user interface block return codes" in *IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes*.

Immediately after the statement that defines the UIB in your program, you must define the PCB address list and the PCB mask.

The following code example shows how to use the COPY DLIUIB statement in a VS COBOL II program:

Defining the UIB, PCB address list, and the PCB mask for VS COBOL II

```

LINKAGE SECTION.

    COPY DLIUIB.
01  OVERLAY-DLIUIB REDEFINES DLIUIB.
    02  PCBADDR USAGE IS POINTER.
    02  FILLER PIC XX.

01  PCB-ADDRESSES.

```

```

02 PCB-ADDRESS-LIST
    USAGE IS POINTER OCCURS 10 TIMES.
01 PCB1.
02 PCB1-DBD-NAME PIC X(8).
02 PCB1-SEG-LEVEL PIC XX.
.
.
.

```

The COBOL COPY DLIUIB copybook

```

01 DLIUIB.
*                               Address of the PCB addr list
02 UIBPCBAL PIC S9(8) COMP.
*                               DL/I return codes
02 UIBRCODE.
*                               Return codes
03 UIBFCTR PIC X.
88 FCNORESP VALUE ' '.
88 FCNOTOPEN VALUE ' '.
88 FCINVREQ VALUE ' '.
88 FCINVPCB VALUE ' '.
*                               Additional information
03 UIBDLTR PIC X.
88 DLPSBNF VALUE ' '.
88 DLTASKNA VALUE ' '.
88 DLPSBSCH VALUE ' '.
88 DLLANGCON VALUE ' '.
88 DLPSBFAIL VALUE ' '.
88 DLPSBNA VALUE ' '.
88 DLTERMNS VALUE ' '.
88 DLFUNCNS VALUE ' '.
88 DLINA VALUE ' '.

```

The values placed in level 88 entries are not printable. They are described in the topic "CICS-DL/I User Interface Block Return Codes" in *IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes*. The meanings of the field names and their hexadecimal values are shown below:

FCNORESP

Normal response Value X'00'

FCNOTOPEN

Not open Value X'0C'

FCINVREQ

Invalid request Value X'08'

FCINVPCB

Invalid PCB Value X'10'

DLPSBNF

PSB not found Value X'01'

DLTASKNA

Task not authorized Value X'02'

DLPSBSCH

PSB already scheduled Value X'03'

DLLANGCON

Language conflict Value X'04'

DLPSBFAIL

PSB initialization failed Value X'05'

DLPSBNA

PSB not authorized Value X'06'

DLTERMNS

Termination not successful Value X'07'

DLFUNCNS

Function unscheduled Value X'08'

DLINA

DL/I not active Value X'FF'

The following code example shows how to define the UIB, PCB address list, and PCB mask for PL/I.

Defining the UIB, PCB address list, and the PCB mask for PL/I

```
DCL UIBPTR PTR; /* POINTER TO UIB */
DCL 1 DLIUIB UNALIGNED BASED(UIBPTR), /* EXTENDED CALL USER INTFC BLK */
    2 UIBPCBAL PTR, /* PCB ADDRESS LIST */
    2 UIBRCODE, /* DL/I RETURN CODES */
    3 UIBFCTR BIT(8) ALIGNED, /* RETURN CODES */
    3 UIBDLTR BIT(8) ALIGNED; /* ADDITIONAL INFORMATION */
```

The following code example shows how to define the UIB, PCB address list, and PCB mask for assembler language.

Defining the UIB, PCB address list, and the PCB mask for assembler language

DLIUIB	DSECT		
UIB	DS	0F	EXTENDED CALL USER INTFC BLK
UIBPCBAL	DS	A	PCB ADDRESS LIST
UIBRCODE	DS	0XL2	DL/I RETURN CODES
UIBFCTR	DS	X	RETURN CODE
UIBDLTR	DS	X	ADDITIONAL INFORMATION
	DS	2X	RESERVED
	DS	0F	LENGTH IS FULLWORD MULTIPLE
UIBLEN	EQU	*-UIB	LENGTH OF UIB

Related reference

[“Coding a CICS online program in COBOL” on page 184](#)

The following code examples are skeleton online programs in Enterprise COBOL. They show examples of how to define and set up addressability to the UIB.

[“Coding a CICS online program in PL/I” on page 192](#)

The following code example is a skeleton CICS online program in PL/I. It shows you how to define and establish addressability to the UIB.

[“Coding a CICS online program in assembler language” on page 178](#)

The following code example in assembler language shows how you define and establish addressability to the UIB.

[“Language specific entry points” on page 227](#)

In your application program written in assembler language, C, COBOL, Pascal, or PL/I, control is passed from IMS through an *entry point*.

Specifying the I/O areas

Use an I/O area to pass segments between the application program and IMS.

What the I/O area contains depends on the type of call you are issuing:

- When you retrieve a segment, IMS places the segment you requested in the I/O area.
- When you add a new segment, you first build the new segment in the I/O area.
- Before modifying a segment, your program must first retrieve it. When you retrieve the segment, IMS places the segment in an I/O area.

The format of the record segments you pass between your program and IMS can be fixed length or variable length. Only one difference is important to the application program: a message segment containing a 2-byte length field (or 4 bytes for the PLITDLI interface) at the beginning of the data area of the segment.

The I/O area for IMS calls must be large enough to hold the largest segment your program retrieves from or sends to IMS.

If your program issues any Get or ISRT calls that use the D command code, the I/O area must be large enough to hold the largest path of segments that the program retrieves or inserts.

Formatting segment search arguments (SSAs)

Segment search arguments in your assembler language, C language, COBOL, Java, Pascal, and PL/I application programs must be coded according to the following rules and formats.

SSA coding rules

Use the following rules for coding a segment search argument.

- Define the SSA in the data area of your program.
- The segment name field must:
 - Be 8 bytes long. If the name of the segment you are specifying is less than 8 bytes long, it should be left justified and padded on the right with blanks.
 - Contain a segment name that has been defined in the DBD that your application program uses. In other words, make sure you use the exact segment name, or your SSA will be invalid.
 - Or, if the DL/I call uses command code O, the segment field name is the starting offset and length of the data that you want to retrieve. The starting offset is relative to the physical segment definition and starts with 1. The maximum length that can be retrieved is the maximum segment size for the database type, and the minimum length is 1. The two fields are specified instead of a standard field name in the following format: 'ooooLLLL'. oooo is the offset position and LLLL is the length of the data that you want to retrieve.
- If the SSA contains only the segment name, byte 9 must contain a blank.
- If the SSA contains one or more command codes:
 - Byte 9 must contain an asterisk (*).
 - The last command code must be followed by a blank unless the SSA contains a qualification statement. If the SSA contains a qualification statement, the command code must be followed by the left parenthesis of the qualification statement.
- If the SSA contains a qualification statement:
 - The qualification statement must begin with a left parenthesis and end with a right parenthesis.
 - There must not be any blanks between the segment name or command codes, if used, and the left parenthesis.
 - The field name must be 8 bytes long. If the field name is less than 8 bytes, it must be left justified and padded on the right with blanks. The field name must have been defined for the specified segment type in the DBD the application program is using.
 - The relational operator follows the field name. It must be 2 bytes long and can be represented alphabetically or symbolically. The following table lists the relational operators.

Table 38. Relational operators

Symbolic	Alphabetic	Meaning
=b=	EQ	Equal to
>= or =>	GE	Greater than or equal to
<= or =<	LE	Less than or equal to
>b>	GT	Greater than
<b<	LT	Less than
≠ or =≠	NE	Not equal to

- The comparative value follows the relational operator. The length of this value must be equal to the length of the field that you specified in the field name. This length is defined in the DBD. The comparative value must include leading zeros for numeric values or trailing blanks for alphabetic values as necessary. The comparative value cannot include any parenthesis.
- If you are using multiple qualification statements within one SSA (Boolean qualification statements), the qualification statements must be separated by one of these symbols:

*** or &**

Dependent AND

+ or |

Logical OR

#

Independent AND

One of these symbols must appear between the qualification statements that the symbol connects.

- The last qualification statement must be followed by a right parenthesis.

An SSA created by the application program must not exceed the space allocated for the SSA in the PSB.

Related reading: For additional information about defining the PSB SSA size, see the explanation of the PSBGEN statement in *IMS Version 15.3 Database Utilities*.

SSA coding formats

Use the following formats to code segment search arguments in assembler language, C language, COBOL, Pascal, and PL/I.

Assembler language SSA definition examples

The following code example shows how you would define a qualified SSA without command codes. If you want to use command codes with this SSA, code the asterisk (*) and command codes between the 8-byte segment name field and the left parenthesis that begins the qualification statement.

```
*      CONSTANT AREA
:
SSANAME DS    0CL26
ROOT    DC    CL8'ROOT    '
        DC    CL1'('
        DC    CL8'KEY      '
        DC    CL2'='
NAME    DC    CLn'vv...v'
        DC    CL1')'
```

This SSA looks like this:

```
ROOTbbbb(KEYbbbbbb=vv...v)
```

C language SSA definition examples

An unqualified SSA that does not use command codes looks like this in C:

```
const struct {
    char seg_name_u[8];
    char blank[1];
} unqual_ssa = {"NAME", " "};
```

You can use an SSA that is coded like this for each DL/I call that needs an unqualified SSA by supplying the name of the segment type you want during program execution. Note that the string size declarations are such that the C null terminators do not appear within the structure.

You can, of course, declare this as a single string:

```
const char unqual_ssa[] = "NAME      "; /* 8 chars + 1 blank */
```

DL/I ignores the trailing null characters.

You can define SSAs in any of the ways explained for the I/O area.

The easiest way to create a qualified SSA is using the **sprintf** function. However, you can also define it using a method similar to that used by COBOL or PL/I.

The following is an example of a qualified SSA without command codes. To use command codes with this SSA, code the asterisk (*) and command codes between the 8-byte segment name field and the left parenthesis that begins the qualification statement.

```
struct {
    seg_name      char[8];
    seg_qual      char[1];
    seg_key_name  char[8];
    seg_opr       char[2];
    seg_key_value char[n];
    seg_end_char  char[1];
} qual_ssa = {"ROOT", "(", "KEY", "=", "vv...vv", ")"};
```

Another way is to define the SSA as a string, using **sprintf**. Remember to use the preprocessor directive **#include <stdio.h>**.

```
char qual_ssa[8+1+8+2+6+1+1]; /* the final 1 is for the */
                                /* trailing '\0' of string */
sprintf(qual_ssa,
        "ROOT", "KEY", "=", "vvvvv");
```

Alternatively, if only the value were changing, the **sprintf** call can be:

```
sprintf(qual_ssa,
        "ROOT (KEY =, "vvvvv");
/* 12345678 12345678 */
```

In both cases, the SSA looks like this:

```
ROOTbbbb(KEYbbbbbb=vv...v)
```

COBOL SSA definition examples

An unqualified SSA that does not use command codes looks like this in COBOL:

```
DATA DIVISION.
WORKING-STORAGE SECTION.
:
01 UNQUAL-SSA.
   02 SEG-NAME    PICTURE X(08)  VALUE '.....'.
   02 FILLER      PICTURE X       VALUE ' '.
```

By supplying the name of the segment type you want during program execution, you can use an SSA coded like the one in this example for each DL/I call that needs an unqualified SSA.

Use a 01 level working storage entry to define each SSA that the program is to use. Then use the name you have given the SSA as the parameter in the DL/I call, in this case:

```
UNQUAL-SSA,
```

The following SSA is an example of a qualified SSA that does not use command codes. If you use command codes in this SSA, code the asterisk (*) and the command code between the 8-byte segment name field and the left parenthesis that begins the qualification statement.

```
DATA DIVISION.
WORKING-STORAGE SECTION.
:
01 QUAL-SSA-MAST.
   02 SEG-NAME-M    PICTURE X(08)  VALUE 'ROOT' .
   02 BEGIN-PAREN-M PICTURE X       VALUE '(' .
   02 KEY-NAME-M    PICTURE X(08)  VALUE 'KEY' .
```

02 REL-OPER-M	PICTURE X(02)	VALUE '='.
02 KEY-VALUE-M	PICTURE X(n)	VALUE 'vv...v'.
02 END-PAREN-M	PICTURE X	VALUE ')'

The SSA looks like this:

```
ROOTbbbb(KEYbbbbbb=vv...v)
```

Pascal SSA definition examples

An unqualified SSA that does not use command codes looks like this in Pascal:

```
type
  STRUCT = record
    SEG_NAME : ALFA;
    BLANK    : CHAR;
  end;
const
  UNQUAL_SSA = STRUCT('NAME', '');
```

You can also declare this SSA as a single string:

```
const
  UNQUAL_SSA = 'NAME  ';
```

The SSA shown in the following example is a qualified SSA that does not use command codes. If you use command codes in this SSA, code the asterisk (*) and the command code between the 8-byte segment name field and the left parenthesis that begins the qualification statement.

```
type
  STRUCT = record
    SEG_NAME      : ALFA;
    SEG_QUAL      : CHAR;
    SEG_KEY_NAME  : ALFA;
    SEG_OPR       : CHAR;
    SEG_KEY_VALUE : packed array[1..n] of CHAR;
    SEG_END_CHAR  : CHAR;
  end;
const
  QUAL_SSA = STRUCT('ROOT', '(', 'KEY', '=', 'vv...v', ')');
```

This SSA looks like this:

```
ROOTbbbb(KEYbbbbbb=vv...v)
```

PL/I SSA definition examples

An unqualified SSA that does not use command codes looks like this in PL/I:

```
DCL 1 UNQUAL_SSA    STATIC UNALIGNED,
    2   SEG_NAME_U CHAR(8) INIT('NAME  '),
    2   BLANK      CHAR(1) INIT(' ');
```

You can use a SSA that is coded like this for each DL/I call that needs an unqualified SSA by supplying the name of the segment type you want during program execution.

In PL/I you define SSAs in structure declarations. The unaligned attribute is required for SSA data interchange with IMS. The SSA character string must reside contiguously in storage. For example, assignment of variable key values might cause IMS to construct an invalid SSA if the key value has changed the aligned attribute.

A separate SSA structure is required for each segment type that the program accesses because the value of the key fields differs among segment types. After you have initialized the fields (other than the key values), the SSA should not need to be changed again. You can define SSAs in any of the ways explained for the I/O area.

The following is an example of a qualified SSA without command codes. If you use command codes in this SSA, code the asterisk (*) and command codes between the 8-byte segment name field and the left parenthesis that begins the qualification statement.

```
DCL 1   QUAL_SSA          STATIC UNALIGNED,
      2   SEG_NAME        CHAR(8) INIT('ROOT  '),
      2   SEG_QUAL        CHAR(1) INIT('('),
      2   SEG_KEY_NAME    CHAR(8) INIT('KEY   '),
      2   SEG_OPR         CHAR(2) INIT('= '),
      2   SEG_KEY_VALUE   CHAR(n) INIT('vv...v'),
      2   SEG_END_CHAR    CHAR(1) INIT(')');
```

This SSA looks like this:

```
ROOTbbbb(KEYbbbbbb=vv...v)
```

Related concepts

[“Specifying segment search arguments using the SSAList interface” on page 646](#)

The SSAList interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Data areas in GSAM databases

Generalized Sequential Access Method (GSAM) databases are available only to application programs that can run as batch programs, batch-oriented BMPs, transaction-oriented BMPS, or JBP. The program communication block (PCB) mask and the record search argument (RSA) that you use in a GSAM database call have special formats.

GSAM DB PCB masks are slightly different from other DB PCB masks. The fields that are different are the length of the key feedback area and the key feedback area. Also, an additional field exists that gives the length of the record being retrieved or inserted when using undefined-length records.

The RSA is an 8-byte token for basic format data sets or 12-byte token for large format data sets that can be returned on GN and ISRT calls. The application program can save the RSA for use in a subsequent GU call.

Related concepts

[“Processing GSAM databases” on page 283](#)

GSAM databases are available to application programs that can run online in IMS batch message processing (BMP) regions (message-driven or non-message-driven) or Java batch processing (JBP) regions or standalone in DLIBATCH regions.

Related reference

[“Specifying the DB PCB mask” on page 212](#)

IMS describes the results of the calls your program issues in the DB PCB that is referenced in the call. To determine the success or failure of the DL/I call, the application program includes a mask of the DB PCB and then references the fields of the DB PCB through the mask.

AIBTDLI interface

Use AIBTDLI as the interface between your application program and IMS.

Restriction: No fields in the application interface block (AIB) can be used by the application program except as defined by IMS.

When you use the AIBTDLI interface, you specify the program communication block (PCB) requested for the call by placing the PCB name (as defined by PSBGEN) in the resource name field of the AIB. You do not specify the PCB address. Because the AIB contains the PCB name, your application program can refer to the PCB name rather than the PCB address. Your application program does not need to know the relative PCB position in the PCB list. At completion of the call, the AIB returns the PCB address that corresponds to the PCB name passed by the application program.

The names of DB PCBs and alternate PCBs are defined by the user during PSBGEN. All I/O PCBs are generated with the PCB name *bbb*. For a generated program specification block (GPSB), the I/O PCB is generated with the PCB name *IOPCBbbb*, and the modifiable alternate PCB is generated with the PCB name *TPPCB1bb*.

The ability to pass the PCB name means that you do not need to know the relative PCB number in the PCB list. In addition, the AIBTDLI interface enables your application program to make calls on PCBs that do not reside in the PCB list. The **LIST=** keyword, which is defined in the PCB macro during PSBGEN, controls whether the PCB is included in the PCB list.

The AIB resides in user-defined storage that is passed to IMS for DL/I calls that use the AIBTDLI interface. Upon call completion, IMS updates the AIB. Allocate at least 128 bytes of storage for the AIB.

Related concepts

[“PCB masks for GSAM databases” on page 283](#)

For the most part, you process GSAM databases in the same way that you process other IMS databases. You use calls that are very similar to DL/I calls to communicate your requests. GSAM describes the results of those calls in a GSAM DB PCB.

Related reference

[“Application programming for PL/I” on page 360](#)

Application programs in PL/I use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

[“Application programming for Pascal” on page 357](#)

Application programs in Pascal use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

[“Application programming for C language” on page 352](#)

Application programs in C use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

[“Application programming for assembler language” on page 349](#)

Application programs in assembly language use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

[“Assembler language application programming” on page 195](#)

Application programs in assembly language use the following format, parameters, and DL/I calls to communicate with IMS databases.

Language specific entry points

In your application program written in assembler language, C, COBOL, Pascal, or PL/I, control is passed from IMS through an *entry point*.

Your entry point must refer to the PCBs in the order in which they have been defined in the PSB. When you code each DL/I call, you must provide the PCB you want to use for that call. In all cases except CICS online, the list of PCBs that the program can access is passed to the program at its entry point. For CICS online, you must first schedule a PSB as described in the topic “System Service Call: PCB” in *IMS Version 15.3 Application Programming APIs*.

Application interfaces that use the AIB structure (AIBTDLI or CEETDLI), such as Java application interfaces, use the PCB name rather than the PCB structure and do not require the PCB list to be passed at entry to the application.

In a CICS online program, you do not obtain the address of the PCBs through an entry statement, but through the user interface block (UIB).

Leave the value blank if the application has been enabled for the IBM Language Environment® for z/OS & VM.

Assembler language entry point

You can use any name for the entry statement to an assembler language DL/I program. When IMS passes control to the application program, register 1 contains the address of a variable-length fullword parameter list. Each word in the list contains the address of a PCB. Save the content of register 1 before you overwrite it. IMS sets the high-order byte of the last fullword in the list to X'80' to indicate the end of the list. Use standard z/OS linkage conventions with forward and backward chaining.

C language entry point

When IMS passes control to your program, it passes the addresses, in the form of pointers, for each of the PCBs that your program uses. The usual **argc** and **argv** arguments are not available to a program that is invoked by IMS. The IMS parameter list is made accessible by using the **__pcblist** macro. You can directly reference the PCBs by **__pcblist[0]**, **__pcblist[1]**, or you can define macros to give these more meaningful names. Note that I/O PCBs must be cast to get the proper type:

```
(IO_PCB_TYPE *) (__pcblist[0])
```

The entry statement for a C language program is the **main** statement.

```
#pragma runopts(env(IMS),plist(IMS))
#include <ims.h>

main()
{
...
}
```

The **env** option specifies the operating environment in which your C language program is to run. For example, if your C language program is invoked under IMS and uses IMS facilities, specify **env(IMS)**. The **plist** option specifies the format of the invocation parameters that is received by your C language program when it is invoked. When your program is invoked by a system support services program, the format of the parameters passed to your main program must be converted into the C language format: **argv**, **argc**, and **envp**. To do this conversion, you must specify the format of the parameter list that is received by your C language program. The **ims.h** include file contains declarations for PCB masks.

You can finish in three ways:

- End the main procedure without an explicit **return** statement.
- Execute a **return** statement from **main**.
- Execute an **exit** or an **abort** call from anywhere, or alternatively issue a **longjmp** back to **main**, and then do a normal return.

One C language program can pass control to another by using the **system** function. The normal rules for passing parameters apply; in this case, the **argc** and **argv** arguments can be used to pass information. The initial **__pcblist** is made available to the invoked program.

COBOL entry point

The procedure statement must refer to the I/O PCB first, then to any alternate PCB it uses, and finally to the DB PCBs it uses. The alternate PCBs and DB PCBs must be listed in the order in which they are defined in the PSB.

```
PROCEDURE DIVISION USING PCB-NAME-1 [, ..., PCB-NAME-N]
```

In previous versions of IMS, USING might be coded on the entry statement to reference PCBs. However, IMS continues to accept such coding on the entry statement.

Recommendation: Use the procedure statement rather than the entry statement to reference the PCBs.

Pascal entry point

The entry point must be declared as a REENTRANT procedure. When IMS passes control to a Pascal procedure, the first address in the parameter list is reserved for Pascal's use, and the other addresses are the PCBs the program uses. The PCB types must be defined before this entry statement. The IMS interface routine PASTDLI must be declared with the GENERIC directive.

```
procedure ANYNAME(var SAVE: INTEGER;
                  var pcb1-name: pcb1-name-type[;
                  ...
                  var pcbn-name: pcbn-name-type]); REENTRANT;
procedure ANYNAME;
(* Any local declarations *)
  procedure PASTDLI; GENERIC;
begin
  (* Code for ANYNAME *)
end;
```

PL/I entry point

The entry statement must appear as the first executable statement in the program. When IMS passes control to your program, it passes the addresses of each of the PCBs your program uses in the form of pointers. When you code the entry statement, make sure you code the parameters of this statement as pointers to the PCBs, and not the PCB names.

```
anyname: PROCEDURE (pcb1_ptr [..., pcbn_ptr]) OPTIONS (MAIN);
:
RETURN;
```

The entry statement can be any valid PL/I name.

CEETDLI, AIBTDLI, and AERTDLI interface considerations

The following considerations apply for CEETDLI, AIBTDLI, and AERTDLI.

The considerations for CEETDLI are:

- For PL/I programs, the CEETDLI entry point is defined in the CEEIBMAW include file. Alternatively, you can declare it yourself, but it must be declared as an assembler language entry (DCL CEETDLI OPTIONS(ASM);).
- For C language application programs, you must specify **env(IMS)** and **plist(IMS)**; these specifications enable the application program to accept the PCB list of arguments. The CEETDLI function is defined in <leawi.h>; the CTDLI function is defined in <ims.h>.

The considerations for AIBTDLI are:

- When using the AIBTDLI interface for C/MVS, Enterprise COBOL, or PL/I language application programs, the language run-time options for suppressing abend interception (that is, NOSPIE and NOSTAE) must be specified. However, for Language Environment-conforming application programs, the NOSPIE and NOSTAE restriction is removed.
- The AIBTDLI entry point for PL/I programs must be declared as an assembler language entry (DCL AIBTDLI OPTIONS(ASM);).
- For C language applications, you must specify **env(IMS)** and **plist(IMS)**; these specifications enable the application program to accept the PCB list of arguments.

The considerations for AERTDLI are:

- When using the AERTDLI interface for C/MVS, COBOL, or PL/I language application programs, the language run-time options for suppressing abend interception (that is, NOSPIE and NOSTAE) must be specified. However, for Language Environment-conforming application programs, the NOSPIE and NOSTAE restriction is removed.
- The AERTDLI entry point for PL/I programs must be declared as an assembler language entry (DCL AERTDLI OPTIONS(ASM);).

- For C language applications, you must specify **env(IMS)** and **plis(IMS)**. These specifications enable the application program to accept the PCB list of arguments.
- AERTDLI must receive control with 31 bit addressability.

Related reference

[“Specifying the UIB \(CICS online programs only\)” on page 219](#)

The interface between your CICS online program and DL/I passes additional information to your program in a user interface block (UIB). The UIB contains the address of the PCB list and any return codes your program must examine before checking the status code in the DB PCB.

Program communication block (PCB) lists

In your application program, code your PCB or GPSB list in the following format.

PCB list format

The following example shows the general format of a PCB list.

```
[IOPCB]
[Alternate PCB ... Alternate PCB]
[DB PCB ... DB PCB]
[GSAM PCB ... GSAM PCB]
```

Each PSB must contain at least one PCB. An I/O PCB is required for most system service calls. An I/O PCB or alternate PCB is required for transaction management calls. (Alternate PCBs can exist in IMS TM.) DB PCBs for DL/I databases are used only with the IMS Database Manager under DBCTL. GSAM PCBs can be used with DCCTL.

Format of a GPSB PCB list

A generated program specification block (GPSB) takes this format:

```
[IOPCB]
[Alternate PCB]
```

A GPSB contains only an I/O PCB and one modifiable alternate PCB. (A modifiable alternate PCB enables you to change the destination of the alternate PCB while the program is running.) A GPSB can be used by all transaction management application programs, and permits access to the specified PCBs without the need for a specific PSB for the application program.

The PCBs in a GPSB have predefined PCB names. The name of the I/O PCB is IOPCB. The name of the alternate PCB is TPPCB1bb. The minimum size of the I/O work area that IMS generates for GPSBs in a DBCTL environment is 600 bytes.

PCB summary

If you intend to issue system service requests, be aware of the differences between I/O PCBs and alternate PCBs in various types of application programs.

DB Batch Programs

If CMPAT=Y is specified in PSBGEN, the I/O PCB is present in the PCB list; otherwise, the I/O PCB is not present, and the program cannot issue system service calls. Alternate PCBs are always included in the list of PCBs that IMS supplies to the program.

BMPs, MPPs, and IFPs

The I/O PCB and alternate PCBs are always passed to BMPs, MPPs, and IFPs.

The PCB list always contains the address of the I/O PCB, followed by the addresses of any alternate PCBs, followed by the addresses of the DB PCBs.

CICS Online Programs with DBCTL

If you specify the IOPCB option on the PCB call, the first PCB address in your PCB list is the I/O PCB, followed by any alternate PCBs, followed by the addresses of the DB PCBs.

If you do not specify the I/O PCB option, the first PCB address in your PCB list points to the first DB PCB.

The following table summarizes the I/O PCB and alternate PCB information.

Table 39. I/O PCB and alternate PCB information summary

Environment	CALL DL/I	
	I/O PCB address in PCB list	Alternate PCB address in PCB list
MPP	Yes	Yes
IFP	Yes	Yes
BMP	Yes	Yes
DB Batch¹	No	Yes
DB Batch²	Yes	Yes
TM Batch³	Yes	Yes
CICS DBCTL⁴	No	No
CICS DBCTL⁵	Yes	Yes

Notes:

1. CMPAT = N specified.
2. CMPAT = Y specified.
3. CMPAT = Option. Default is always to Y, even when CMPAT = N is specified.
4. SCHD request issued without the IOPCB or SYSSERVE option.
5. SCHD request issued with the IOPCB or SYSSERVE for a CICS DBCTL request or for a function-shipped request which is satisfied by a CICS system using DBCTL.

The AERTDLI interface

You can make database calls with AIBs in your ODBA applications using the AERTDLI interface.

Requirement: Allocate 264 bytes of storage for the AIB.

When you use the AERTDLI interface, the AIB used for database calls must be the same AIB as used for the APSB call. Specify the PCB that is requested for the call by placing the PCB name (as defined by PSBGEN) in the resource name field of the AIB. You do not specify the PCB address. Because the AIB contains the PCB name, your application can refer to the PCB name rather than to the PCB address. The AERTDLI call allows you to select PCBs directly by name rather than by a pointer to the PCB. At completion of the call, the AIB returns the PCB address that corresponds to the PCB name that is passed by the application program.

For PCBs to be used in a AERTDLI call, you must assign a name in PSBGEN, either with PCBNAME= or with the name as a label on the PCB statement. PCBs that have assigned names are also included in the positional pointer list, unless you specify LIST=NO. During PSBGEN, you define the names of the DB PCBs and alternate PCBs. All I/O PCBs are generated with the PCB name IOPCBbbb.

Because you pass the PCB name, you do not need to know the relative PCB number in the PCB list. In addition, the AERTDLI interface enables your application program to make calls on PCBs that do not reside in the PCB list. The LIST= keyword, which is defined in the PCB macro during PSBGEN, controls whether the PCB is included in the PCB list.

The AIB resides in user-defined storage that is passed to IMS for DL/I calls that use the AERTDLI interface. When the call is completed, the AIB is updated by IMS. Because some of the fields in the AIB are used internally by IMS, the same APSB AIB must be used for all subsequent calls for that PSB.

Language environments

IBM Language Environment provides the strategic execution environment for running your application programs written in one or more high level languages.

It provides not only language-specific run-time support, but also cross-language run-time services for your applications, such as support for initialization, termination, message handling, condition handling, storage management, and National Language Support. Many of Language Environment's services are accessible explicitly through a set of Language Environment interfaces that are common across programming languages; these services are accessible from any Language Environment-conforming program.

Language Environment-conforming programs can be compiled with the following compilers:

- IBM C++/MVS
- IBM COBOL
- IBM PL/I

By default, the Language Environment infrastructure uses the 31-bit addressing mode. By specifying JVM=64, the Language Environment is changed to use the 64-bit addressing mode. By specifying JVM=3164, the dependent region will initialize a 31-bit Language Environment and a secondary 64-bit Language Environment to support 31-bit COBOL and 64-bit Java interoperability.

Language Environment supports C, C++, and assembly language interoperability in a 64-bit addressing mode, but does not support COBOL and PL/I interoperability in a 64-bit addressing mode. Do not switch to JVM=64 if your Java application invokes either COBOL or PL/I. If the regions are switched to use JVM=64 inadvertently, and incompatible interoperable applications are running, the application might receive system or user abends.

The CEETDLI interface to IMS

The language-independent CEETDLI interface to IMS is provided by Language Environment. It is the only IMS interface that supports the advanced error handling capabilities provided by Language Environment. The CEETDLI interface supports the same functionality as the other IMS application interfaces, and it has the following characteristics:

- The parmcount variable is optional.
- Length fields are 2 bytes long.
- Direct pointers are used.

Related reading: For more information about Language Environment, see *z/OS Language Environment Programming Guide*.

LANG= option on PSBGEN for PL/I compatibility

For IMS PL/I applications running in a compatibility mode that uses the PLICALLA entry point, you must specify LANG=PLI on the PSBGEN. Your other option is to change the entry point and add SYSTEM(IMS) to the EXEC PARM of the compile step so that you can specify LANG=blank or LANG=PLI on the PSBGEN. The following table summarizes when you can use LANG=blank and LANG=PLI.

Table 40. Using LANG= option in a Language Environment for PL/I compatibility

Compile exec statement is PARM=(...,SYSTEM(IMS)...	Entry point name is PLICALLA	Valid LANG= value
Yes	Yes	LANG=PLI

Table 40. Using LANG= option in a Language Environment for PL/I compatibility (continued)

Compile exec statement is PARM=(...,SYSTEM(IMS)...	Entry point name is PLICALLA	Valid LANG= value
Yes	No	LANG=blank or LANG=PLI
No	No	Note: Not valid for IMS PL/I applications
No	Yes	LANG=PLI

PLICALLA is only valid for PL/I compatibility with Language Environment. If a PL/I application using PLICALLA entry at bind time is bound using Language Environment with the PLICALLA entry, the bind will work; however, you must specify LANG=PLI in the PSB. If the application is re-compiled using PL/I for z/OS & VM Version 1 Release 1 or later, and then bound using Language Environment Version 1 Release 2 or later, the bind will fail. You must remove the PLICALLA entry statement from the bind.

Special DL/I situations for IMS DB programming

Special cases during application programming for IMS DB include usage of the GUR call, program scheduling against HALDBs, mixed language programming, using the extended addressing capabilities of z/OS, and setting COBOL compiler options for preloaded programs.

GUR call

The get unique record (GUR) DL/I call is a special case because it always accesses the IMS catalog database. When the catalog is enabled, IMS dynamically attaches the catalog PCB on behalf of your application program. Your application program can use the GUR call to get catalog data in the form of a single XML instance document for a particular catalog record. You can also issue other DL/I read calls to process the catalog database in the same way as any other database. The GUR call is provided to reduce the number of processing steps required to retrieve a complete catalog record for a DBD or PSB.

Restriction: The use of SSA command codes is not allowed.

Application program scheduling against HALDBs

Application programs are scheduled against HALDBs the same way they are against non-HALDBs. Scheduling is based on the availability status of the HALDB master and is not affected by individual partition access and status.

The application programmer needs to be aware of changes to the handling of unavailable data for HALDBs. The feedback on data availability at PSB schedule time shows the availability of the HALDB master, not of the partitions. However, the error settings for data unavailability of a partition at the first reference to the partition during the processing of a DL/I call are the same as those of a non-HALDB, namely status code BA or pseudo ABENDU3303.

For example, if you issue the IMS /DBR command to half of the partitions to take them offline, the remaining partitions are available to the programs.

When an application program accesses a partition, that partition is considered to be in use by the application for the duration of that instance of the application. DBDUMP, DBRECOVERY, and START commands can operate against a partition currently not in use. The command is not processed for any partition that is being accessed by a BMP. A DFS0565I message is issued for partitions that are in use by a BMP. An exception to this rule is a partition where the accessing BMP issued a CHKP call and has not issued any subsequent DL/I calls. If an application attempts to access data from a stopped partition, a pseudo abend ABENDU3303 results or the application receives a BA status code. If the partition is started with the STA DB command before the application attempts to access data in that partition again, the DL/I call is processed successfully.

Mixed-language programming

When an application program uses the Language Environment language-independent interface, CEETDLI, IMS does not need to know the language of the calling program.

When the application program calls IMS in a language-dependent interface, IMS determines the language of the calling program according to the entry name that is specified in the CALL statement. That is, IMS assumes that the program is:

- Assembler language when the application program uses CALL ASMTDLI
- C language when the application program uses rc=CTDLI
- COBOL when the application program uses CALL CBLTDLI
- Pascal when the application program uses CALL PASTDLI
- PL/I when the application program uses CALL PLITDLI

For example, if a PL/I program calls an assembler language subroutine and the assembler language subroutine makes DL/I calls by using CALL ASMTDLI, the assembler language subroutine should use the assembler language calling convention, not the PL/I convention.

In this situation, where the I/O area uses the LLZZ format, LL is a halfword, not the fullword that is used for PL/I.

Extended addressing capabilities of z/OS

The two modes in z/OS with extended addressing capabilities are: the addressing mode (AMODE) and the residency mode (RMODE). IMS places no constraints on the RMODE and AMODE of an application program. The program can reside in the extended virtual storage area. The parameters that are referenced in the call can also be in the extended virtual storage area.

COBOL compiler options for preloaded programs

If you compile your COBOL program with the VS COBOL II compiler and preload it, you must use the COBOL compiler options RES and RENT.

Application programming with the IMS catalog

The IMS catalog database is accessible to standard IMS DB application programs when it is enabled for your IMS system.

Information in the IMS catalog

The IMS catalog database stores application and database metadata in a format that is accessible to standard IMS DB application programs. This information includes database definitions, program specifications, and user comments. Any application program can read this information, but the catalog database is write-protected and can be updated only by authorized system utilities such as the IMS catalog populate utility (DFS3PU00).

By default, the IMS catalog is named DFSCD000. The *DFSC* prefix is replaced with an alias prefix if one is defined to IMS.

Information in the IMS catalog secondary index

The IMS catalog secondary index contains a single segment type, DBDPSB. It is logically linked to the DBDXREF segment type in the IMS catalog database, which is included in all catalog records for IMS PSBs. You can use the catalog secondary index to determine which IMS programs reference a specific user database without processing the entire IMS catalog.

By default, the IMS catalog is named DFSCX000. The *DFSC* prefix is replaced with an alias prefix if one is defined to IMS.

IMS catalog PSBs and PCBs for application programs

IMS does not require user PSBs to contain a PCB for the IMS catalog database or secondary index. The catalog PSBs DFSCP000, DFSCP002, and DFSCP003 are dynamically attached to any user PSB that makes a DL/I call to the catalog database or issues an INIT DB QUERY call. Each PSB is intended for use by a different type of application program:

DFSCP000

High-level assembler and COBOL applications

DFSCP002

PL/I applications

DFSCP003

PASCAL applications

Restriction: The IMS catalog PSBs are not dynamically attached to generated PSBs or GSAM-only PSBs.

The following PCBs are included to support different catalog processing models:

DFSCAT00

The primary PCB to access all data in the DFSCD000 (IMS catalog) database. Use this PCB to perform standard catalog processing.

DFSCATSX

This PCB provides a SENSEG for the DBDXREF segment type in catalog PSB records and uses PROCSEQ=DFSCX000. Use this PCB to perform faster processing of the catalog database via the catalog secondary index.

DFSCATX0

This PCB provides a SENSEG for the DBDPSB segment type in catalog secondary index records. Use this PCB to process the catalog secondary index directly.

All catalog PCBs are resident. All catalog processing is performed with PROCOPT=GP.

IMS automatically increases the space allocated for the user PSB to attach the catalog PSBs. 96 bytes of additional space are allocated for each user PSB in the PSB CSA storage pool. The catalog PSB itself occupies 12kb in the DLIPSB pool and 500 bytes CSAPSB pool for each user PSB that is using the catalog PSBs. You might need to increase the size of your storage pools, up to the maximum size of the catalog PSB in each pool multiplied by the number of user PSBs that concurrently access the catalog.

You can use the formula in the following table to find the increase in quantities of DLMP, DPSB, or PSBW after you enable the IMS catalog.

	Increase in quantity of		
	DLMP	DPSB	PSBW
The number of active PST in parallel (PPP)	$PPP \times 288_{(b)}$	0	0
The number of dynamic attachment happening in parallel (QQQ) _(a)	$(QQQ \times 144)_{(c)} + (QQQ - 1) \times 1,248_{(d)}$	$(QQQ - 1) \times 17,888_{(e)}$	$QQQ \times 15,792_{(f)}$

	Increase in quantity of		
	DLMP	DPSB	PSBW
(a) The dynamic attachment of IMS catalog happens when: • An application for the PSB issues a GUR call against IMS catalog HALDB for the first time in an IMS system. • An application for the PSB issues INQY DBQUERY call. (b) 288(x120) = Attached catalog DBPCBs 288(x120) (c) 144(x90) = DFSCP000 prefix 32(x20) + copied PDIR of DFSCP000 112(x70) (d) 1,248(x4E0) = Copied DFSCP000 960(x3C0) + PSB3INFO 288(x120) (e) 17,888(x45E0) = Copied DFSCP000 17,888(x45E0) (f) 15,792(x3BD0) = Work pool for DFSCP000 15,792(x3BD0)			

GUR call

Your application program can use the Get Unique Record (GUR) DL/I call to get catalog data in the form of a single XML instance document for a particular catalog record. You can also issue other DL/I read calls to process the catalog database in the same way as any other IMS database. The GUR call is provided to reduce the number of processing steps required to retrieve a complete catalog record for a DBD or PSB.

Restriction: The use of SSA command codes is not allowed.

Related concepts

[Format of records in the IMS catalog database \(Database Administration\)](#)

[IMS catalog secondary index \(Database Administration\)](#)

Related reference

[“Special DL/I situations for IMS DB programming” on page 233](#)

Special cases during application programming for IMS DB include usage of the GUR call, program scheduling against HALDBs, mixed language programming, using the extended addressing capabilities of z/OS, and setting COBOL compiler options for preloaded programs.

[GUR call \(Application Programming APIs\)](#)

Chapter 13. Database versioning and application programming

When database versioning is enabled in an IMS system, IMS can maintain multiple versions of the structural definition of a database so that existing application programs can continue to access a database after the database is modified to support new application programs.

When a new version of a database is defined, the database administrator specifies a version number for the new database definition. The version number is then used to request access to that version of the database.

When multiple version of a database are available, if a specific database version is not specified for an application program, IMS provides access to the current version of the database by default. The current version of a database has the highest version number and contains the latest changes to the database. This IMS system default can be changed so that IMS provides access to version 0 of the database instead.

The IMS system default can be overridden at the program specification block (PSB) level by specifying the DBLEVEL parameter in the PSBGEN statement during PSB generation.

If an application program requires a specific database version, that version number can be specified explicitly either on the DBVER parameter of a PCB statement when the PCB is defined or at runtime by issuing the DL/I INIT VERSION call.

If the requested version of a database definition cannot be found or if database versioning is not enabled when a version is requested, IMS terminates the program with abend 3303 and issues message DFS3303I, which contains details regarding the cause of the abend. Optionally, application programs can issue the INIT STATUS GROUPE call to receive a BA status code instead of abend 3303.



Attention: When a new version of a database is created, before application programs update the new version of the database, confirm that the prior versions of the database can still be accessed.

Database versioning supports only certain changes to a database definition. If unsupported changes are made to a database, application programs will not be able to access the prior versions of the database. Only the current version of the database is accessible.

For most database types, the unsupported changes are not detected until an application program that uses a prior version of the database is scheduled. However, if the HALDB alter function is used to apply the structural changes to a HALDB database, IMS detects unsupported database changes during alter processing.

If a new version of a database contains unsupported changes, either all application programs need to be updated to use the database structure of the new version or the database definition needs to be changed to remove the unsupported structure change.

Batch application programs and database versioning

You can enable database versioning for offline DL/I batch application programs that run in DLIBATCH or DBBBATCH regions by specifying DBVERSION=Y in a DFSDFxxx member in the IMS.PROCLIB data set.

The DLIBATCH or DBBBATCH application programs reference the DFSDFxxx member by specifying the DFSDF=xxx parameter in the EXEC statement of their JCL. For example:

```
//STEP1 EXEC PGM=DFSRR00,REGION=0M,  
// PARM=(DLI,DFSDDL0,PSBCJK03,,01,,,,,,,,,BCH1,,Y,Y,,,,,,,,,,,,,  
// ,,,,,,DFSDF=C35')
```

Important: DLIBATCH application programs use PSB and DBD libraries instead of an ACB library. When using database versioning, DLIBATCH application programs must use the DBD library that contains the DBD member that matches the current physical database structure.

Related concepts

[Database versioning \(Database Administration\)](#)

Related tasks

[Altering the definition of an online HALDB database \(Database Administration\)](#)

Related reference

[INIT call \(Application Programming APIs\)](#)

[PSBGEN statement \(System Utilities\)](#)

Chapter 14. Establishing a DL/I interface from COBOL or PL/I

To establish a DL/I interface from COBOL or PL/I, use either the CBLTDLI procedure or the PLITDLI procedure.

CBLTDLI

The following control statements are necessary to establish a COBOL to DL/I interface. The block size of the following members must be less than or equal to 3200.

```
LIBRARY SDFSRESL(CBLTDLI)    DL/I LANGUAGE INTERFACE
LIBRARY SDFSRESL(DFHEI01)   HLPI LANGUAGE INTERFACE
LIBRARY SDFSRESL(DFHEI1)    HLPI LANGUAGE INTERFACE
```

PLITDLI

The following control statements are necessary to establish a PL/I to DL/I interface. The blocksize of the following members must be less than or equal to 3200.

```
LIBRARY SDFSRESL(PLITDLI)    DL/I LANGUAGE INTERFACE
LIBRARY SDFSRESL(DFHEI01)   HLPI LANGUAGE INTERFACE
LIBRARY SDFSRESL(DFHEI1)    HLPI LANGUAGE INTERFACE
ENTRY  PLICALLA
```

PLITDLI is valid when using the PL/I Optimizing Compiler.

Chapter 15. Current position in the database after each call

Positioning means that DL/I tracks your place in the database after each call that you issue. By tracking your position in the database, DL/I enables you to process the database sequentially.

Current position after successful calls

Position is important when you process the database sequentially by issuing GN, GNP, GHN, and GHNP calls.

Current position is where IMS starts its search for the segments that you specify in the calls.

This section explains current position for successful calls. Current position is also affected by an unsuccessful retrieval or ISRT call.

Before you issue the first call to the database, the current position is the place immediately before the first root segment occurrence in the database. This means that if you issue an unqualified GN call, IMS retrieves the first root segment occurrence. It is the **next** segment occurrence in the hierarchy that is defined by the DB PCB that you referenced.

Certain calls cancel your position in the database. You can reestablish this position with the GU call. Because the CHKP and SYNC (commit point) calls cancel position, follow either of these calls with a GU call. The ROLS and ROLB calls also cancel your position in the database.

When you issue a GU call, your current position in the database does not affect the way that you code the GU call or the SSA you use. If you issue the same GU call at different points during program execution (when you have different positions established), you will receive the same results each time you issue the call. If you have coded the call correctly, IMS returns the segment occurrence you requested regardless of whether the segment is before or after the current position.

Exception: If a GU call does not have SSAs for each level in the call, it is possible for IMS to return a different segment at different points in your program. This is based on the position at each level.

For example, suppose you issue the following call against the data structure shown in the following figure.

```
GU  Abbbbbbb(AKEYbbbbA1)
    Bbbbbbbb(BKEYbbbb=bB11)
    Dbbbbbbb(DKEYbbbbD111)
```

The structure in the figure contains six segment types: A, B, C, D, E, and F. [Figure 49 on page 242](#) shows one database record, the root of which is A1.




Figure 49. Current position hierarchy

When you issue this call, IMS returns the D segment with the key D111, regardless of where your position is when you issue the call. If this is the first call your program issues (and if this is the first database record in the database), current position before you issue the call is immediately before the first segment occurrence in the database—just before the A segment with the key of A1. Even if current position is past segment D111 when you issue the call (for example, just before segment F111), IMS still returns the segment D111 to your program. This is also true if the current position is in a different database record.

When you issue GN and GNP calls, current position in the database affects the way that you code the call and the SSA. That is because when IMS searches for a segment described in a GN or GNP call, it starts the search from current position and can only search forward in the database. IMS cannot look behind that segment occurrence to satisfy a GN or GNP. These calls can only move forward in the database when trying to satisfy your call, unless you use the F command code, the use of which is described in the topic "F Command Code" in *IMS Version 15.3 Application Programming APIs*.

If you issue a GN call for a segment occurrence that you have already passed, IMS starts searching at the current position and stops searching when it reaches the end of the database (resulting in a GB status code), or when it determines from your SSA that it cannot find the segment you have requested (GE status code).

Current position affects ISRT calls when you do not supply qualified SSAs for the parents of the segment occurrence that you are inserting. If you supply only the unqualified SSA for the segment occurrence, you must be sure that your position in the database is where you want the segment occurrence to be inserted.

Related concepts

[A command code \(Application Programming APIs\)](#)

[G command code \(Application Programming APIs\)](#)

["Current position after unsuccessful calls" on page 246](#)

IMS establishes another kind of position when you issue retrieval and ISRT calls. This is position on one segment occurrence at each hierarchic level in the path to the segment that you are retrieving or inserting. Not every DL/I call that your program issues will be completely successful. When a call is unsuccessful, you should understand how to determine your position in the database after that call.

Position after retrieval calls

After you issue any kind of successful retrieval call, position immediately follows the segment occurrence you just retrieved or the lowest segment occurrence in the path if you retrieved several segment occurrences using the D command code. When you use the D command code in a retrieval call, a successful call is one that IMS completely satisfies.

For example, if you issue the following call against the database shown in the previous figure, IMS returns the C segment occurrence with the key of C111. Current position is immediately **after** C111. If you then issue an unqualified GN call, IMS returns the C112 segment to your program.

```
GU  Abbbbbbb(AKEYbbbbEQA1)
    Bbbbbbbb(BKEYbbbbEQB11)
    Cbbbbbbb(CKEYbbbbEQC111)
```

Your current position is the same after retrieving segment C111, whether you retrieve it with GU, GN, GNP, or any of the Get Hold calls.

If you retrieve several segment occurrences by issuing a Get call with the D command code, current position is immediately after the lowest segment occurrence that you retrieved. If you issue the GU call as shown in the example above, but include the D command code in the SSA for segments A and B, the current position is still immediately after segment C111. C111 is the last segment that IMS retrieves for this call. With the D command code, the call looks like this:

```
GU  Abbbbbbb(AKEYbbbbEQA1)
    Bbbbbbbb(BKEYbbbbEQB11)
    Cbbbbbbb*D(CKEYbbbbEQC111)
```

You do not need the D command code on the SSA for the C segment because IMS always returns to your I/O area the segment occurrence that is described in the last SSA.

Position after DLET

After a successful DLET call, position immediately follows the segment occurrence you deleted. This is true when you delete a segment occurrence with or without dependents.

For example, if you issue the call shown in the following code example to delete segment C111, current position is immediately after segment C111. Then, if you issue an unqualified GN call, IMS returns segment C112.

```
GHU  Abbbbbbb(AKEYbbbb=bA1)
      Bbbbbbbb(BKEYbbbb=bB11)
      Cbbbbbbb(CKEYbbbb=bC111)
DLET
```

The following figure shows what the hierarchy looks like after this call. The successful DLET call has deleted segment C111.




Figure 50. Hierarchy after deleting a segment

When you issue a successful DLET call for a segment occurrence that has dependents, IMS deletes the dependents, and the segment occurrence. Current position still immediately follows the segment occurrence you deleted. An unqualified GN call returns the segment occurrence that followed the segment you deleted.

For example, if you delete segment B11 in the hierarchy shown in the previous figure, IMS deletes its dependent segments, C112 and D111, as well. Current position immediately follows segment B11, just before segment B12. If you then issue an unqualified GN call, IMS returns segment B12. The following figure shows what the hierarchy looks like after you issued this call.




Figure 51. Hierarchy after deleting a segment and dependents

Because IMS deletes the segment's dependents, you can think of current position immediately following the last (lowest, right-most) dependent. In the example in the first figure, this immediately follows segment D111. But if you then issue an unqualified GN call, IMS still returns segment B12. You can think of position in either place—the results are the same either way. An exception to this can occur for a DLET that follows a GU path call, which returned a GE status code.

Related concepts

[“Current position after unsuccessful calls” on page 246](#)

IMS establishes another kind of position when you issue retrieval and ISRT calls. This is position on one segment occurrence at each hierarchic level in the path to the segment that you are retrieving or inserting. Not every DL/I call that your program issues will be completely successful. When a call is unsuccessful, you should understand how to determine your position in the database after that call.

Position after REPL

A successful REPL call does not change your position in the database. Current position is just where it was before you issued the REPL call.

It immediately follows the lowest segment that is retrieved by the Get Hold call that you issued before the REPL call.

For example, if you retrieve segment B13 in the previous figure using a GHU instead of a GU call, change the segment in the I/O area, and then issue a REPL call, current position immediately follows segment B13.

Position after ISRT

After you add a new segment occurrence to the database, current position immediately follows the new segment occurrence.

For example, in the following figure, if you issue the following call to add segment C113 to the database, current position immediately follows segment C113. An unqualified call would retrieve segment D111.

```
ISRT  Abbbbbbb(AKEYbbbb=bA1)
      Bbbbbbbb(BKEYbbbb=bB11)
      Cbbbbbbb
```

If you are inserting a segment that has a unique key, IMS places the new segment in key sequence. If you are inserting a segment that has either a non-unique key or no key at all, IMS places the segment according to the rules parameter of the SEGM statement of the DBD for the database. the topic "ISRT Call" in *IMS Version 15.3 Application Programming APIs* explains these rules.

If you insert several segment occurrences using the D command code, current position immediately follows the lowest segment occurrence that is inserted.

For example, suppose you insert a new segment B (this would be B14), and a new C segment occurrence (C141), which is a dependent of B14. The following figure shows what the hierarchy looks like after these segment occurrences are inserted. The call to do this looks like this:

```
ISRT  Abbbbbbb(AKEYbbbb=bA1)
      Bbbbbbbb
*D    Cbbbbbbb
```

You do not need the D command code in the SSA for the C segment. On ISRT calls, you must include the D command code in the SSA for the only first segment you are inserting. After you issue this call, position immediately follows the C segment occurrence with the key of C141. Then, if you issue an unqualified GN call, IMS returns segment E11.

If your program receives an II status code as a result of an ISRT call (which means that the segment you tried to insert already exists in the database), current position is just **before** the duplicate of the segment that you tried to insert.




Figure 52. Hierarchy after adding new segments and dependents

Current position after unsuccessful calls

IMS establishes another kind of position when you issue retrieval and ISRT calls. This is position on one segment occurrence at each hierarchic level in the path to the segment that you are retrieving or inserting. Not every DL/I call that your program issues will be completely successful. When a call is unsuccessful, you should understand how to determine your position in the database after that call.

You need to know how IMS establishes this position to understand the U and V command codes described in the topic "General Command Codes for DL/I Calls" in *IMS Version 15.3 Application Programming APIs*. Also, you need to understand where your position in the database is when IMS returns a not-found status code to a retrieval or ISRT call.

Position after an unsuccessful DLET or REPL call

DLT and REPL calls do not affect current position. Your position in the database is the same as it was before you issued the call. However, an unsuccessful Get call or ISRT call does affect your current position.

To understand where your position is in the database when IMS cannot find the segment you have requested, you need to understand how DL/I determines that it cannot find your segment.

In addition to establishing current position after the lowest segment that is retrieved or inserted, IMS maintains a second type of position on one segment occurrence at each hierarchic level in the path to the segment you are retrieving or inserting.

For example, in the following figure, if you had just successfully issued the GU call with the SSA shown below, IMS has a position established at each hierarchic level.

```

GU  Abbbbbbb(AKEYbbbb=bA1)
    Bbbbbbbb(BKEYbbbbbbB11)
    Cbbbbbbb(CKEYbbbb=bC111)
  
```

Now DL/I has three positions, one on each hierarchic level in the call:

- One on the A segment with the key A1

- One on the B segment with the key B11
- One on the C segment with the key C111




Figure 53. DL/I positions

When IMS searches for a segment occurrence, it accepts the first segment occurrence it encounters that satisfies the call. As it does so, IMS stores the key of that segment occurrence in the key feedback area.

Position after an unsuccessful retrieval or ISRT call

Current position after a retrieval or ISRT call that receives a GE status code depends on how far IMS got in trying to satisfy the SSA in the call. When IMS processes an ISRT call, it checks for each of the parents of the segment occurrence you are inserting. An ISRT call is similar to a retrieval call, because IMS processes the call level by level, trying to find segment occurrences to satisfy each level of the call. When IMS returns a GE status code on a retrieval call, it means that IMS was unable to find a segment occurrence to satisfy one of the levels in the call. When IMS returns a GE status code on an ISRT call, it means that IMS was unable to find one of the parents of the segment occurrence you are inserting. These are called not-found calls.

When IMS processes retrieval and ISRT calls, it tries to satisfy your call until it determines that it cannot. When IMS first tries to find a segment matching the description you have given in the SSA and none exists under the first parent, IMS tries to search for your segment under another parent. How you code the SSA in the call determines whether IMS can move forward and try again under another parent.

For example, suppose you issue the following GN call to retrieve the C segment with the key of C113 in the hierarchy shown in the previous figure.

```

Abbbbbbb(AKEYbbbb=bA1)
Bbbbbbbb(BKEYbbbb=bB11)
Cbbbbbbb(CKEYbbbb=bC113)
  
```

When IMS processes this call, it searches for a C segment with the key equal to C113. IMS can only look at C segments whose parents meet the qualifications for the A and B segments. The B segment that is part of the path must have a key equal to B11, and the A segment that is part of the path must have a key equal to A1. IMS then looks at the first C segment. Its key is C111. The next C segment has a key of C112.

IMS looks for a third C segment occurrence under the B11 segment occurrence. No more C segment occurrences exist under B11.

Because you have specified in the SSA that the A and B segment occurrences in C's path must be equal to certain values, IMS cannot look for a C segment occurrence with a key of C113 under any other A or B segment occurrence. No more C segment occurrences exist under the parent B11; the parent of C must be B11, and the parent of B11 must be A1. IMS determines that the segment you have specified does not exist and returns a not-found (GE) status code.

When you receive the GE status code on this call, you can determine where your position is from the key feedback area, which reflects the positions that IMS has at the levels it was able to satisfy, in this case, A1 and B11.

After this call, current position immediately follows the last segment occurrence that IMS examined in trying to satisfy your call, in this case, C112. Then, if you issue an unqualified GN call, IMS returns D111.

The current position after this call is different if A and B have non-unique keys. Suppose A's key is unique and B's is non-unique. After IMS searches for a C113 segment under B11 and is unable to find one, IMS moves forward from B11 to look for another B segment with a key of B11. When IMS does not find one, DL/I returns a GE status code. Current position is further in the database than it was when both keys were unique. Current position immediately follows segment B11. An unqualified GN call would return B12.

If B is an unkeyed segment, there can be no position at B level. Position is set after the highest segment matching SSA qualification, in this case, A1.

If A and B both have non-unique keys, current position after the previous call immediately follows segment A1. Assuming no more segment A1s exist, an unqualified GN call would return segment A2. If other A1s exist, IMS tries to find a segment C113 under the other A1s.

But suppose you issue the same call with a greater-than-or-equal-to relational operator in the SSA for segment B:

```
GU  Abbbbbbb(AKEYbbbb=>bA1)
     Bbbbbbbb(BKEYbbbb=>B11)
     Cbbbbbbb(CKEYbbbb=>bC113)
```

IMS establishes position on segment A1 and segment B11. Because A1 and B11 satisfy the first two SSAs in the call, IMS stores their keys in the key feedback area. IMS searches for a segment C113 under segment B11. None is found. But this time, IMS can continue searching, because the key of the B parent can be greater than or equal to B11. The next segment is B12. Because B12 satisfies the qualification for segment B, IMS places B12's key in the key feedback area. IMS then looks for a C113 under B12 and does not find one. The same thing happens for B13: IMS places the key of B13 in the key feedback area and looks for a C113 under B13.

When IMS finds no more B segments under A1, it again tries to move forward to look for B and C segments that satisfy the call under another A parent. But this time it cannot; the SSA for the A segment specifies that the A segment must be equal to A1. (If the keys were non-unique, IMS could look for another A1 segment.) IMS then knows that it cannot find a C113 under the parents you have specified and returns a GE status code to your program.

In this example, you have not limited the search for segment C113 to only one B segment, because you have used the greater-than-or-equal-to operator. The position is further than you might have expected, but you can tell what the position is from the key feedback area. The last key in the key feedback area is the key of segment B13. The current position of IMS immediately follows segment B13. If you then issue an unqualified GN call, IMS returns segment E11.

Each of the B segments that IMS examines for this call satisfies the SSA for the B segment, so IMS places the key of each in the key feedback area. But if one or more of the segments IMS examines does not satisfy the call, IMS does not place the key of that segment in the key feedback area. This means that the position in the database might be further than the position reflected by the key feedback area. For example, suppose you issue the same call, but you qualify segment B on a data field in addition to the key field. To do this, you use multiple qualification statements for segment B.

Assume the data field you are qualifying the call on is called BDATA. Assume the value you want is 14, but that only one of the segments, B11, contains a value in BDATA of 14:

```
GN  Abbbbbbb(AKEYbbbb=bA1)
    Bbbbbbbb(BKEYbbbb>=B11*BDATABbb=b14)
    Cbbbbbbb(CKEYbbbb=bC113)
```

After you issue this call, the key feedback area contains the key for segment B11. If you continue issuing this call until you receive a GE status code, the current position immediately follows segment B13, but the key feedback area still contains only the key for segment B11. Of the B segments IMS examines, only one of them (B11) satisfies the SSA in the call.

When you use a greater-than or greater-than-or-equal-to relational operator, you do not limit the search. If you get a GE status code on this kind of call, and if one or more of the segments IMS examines does not satisfy an SSA, the position in the database may be further than the position reflected in the key feedback area. If, when you issue the next GN or GNP call, you want IMS to start searching from the position reflected in the key feedback area instead of from its real position, you can either:

- Issue a fully qualified GU call to reestablish position to where you want it.
- Issue a GN or GNP call with the U command code. Including a U command code on an SSA tells IMS to use the first position it established at that level as qualification for the call. This is like supplying an equal-to relational operator for the segment occurrence that IMS has positioned on at that level.

For example, suppose that you first issue the GU call with the greater-than-or-equal-to relational operator in the SSA for segment B, and then you issue this GN call:

```
GN  Abbbbbbb*U
    Bbbbbbbb*U
    Cbbbbbbb
```

The U command code tells IMS to use segment A1 as the A parent, and segment B11 as the B parent. IMS returns segment C111. But if you issue the same call without the U command code, IMS starts searching from segment B13 and moves forward to the next database record until it encounters a B segment. IMS returns the first B segment it encounters.

Related concepts

[“Position after DLET” on page 243](#)

After a successful DLET call, position immediately follows the segment occurrence you deleted. This is true when you delete a segment occurrence with or without dependents.

Multiple processing

The order in which an application program accesses segments in a hierarchy depends on the purpose of the application program. Some programs access segments directly, others sequentially. Some application programs require that the program process segments in different hierarchic paths, or in different database records, in parallel.

If your program must process segments from different hierarchic paths or from different database records in parallel, using multiple positioning or multiple PCBs can simplify the program's processing. For example:

- Suppose your program must retrieve segments from **different hierarchic paths** alternately: for example, in the following figure, it might retrieve B11, then C11, then B12, then C12, and so on. If your program uses **multiple positioning**, IMS maintains positions in both hierarchic paths. Then the program is not required to issue GU calls to reset position each time it needs to retrieve a segment from a different path.
- Suppose your program must retrieve segments from **different database records** alternately: for example, it might retrieve a B segment under A1, and then a B segment under another A root segment. If your program uses **multiple PCBs**, IMS maintains positions in both database records. Then the program does not have to issue GU calls to reset position each time it needs to access a different database record.




Figure 54. Multiple processing

Multiple positioning

When you define the PSB for your application program, you have a choice about the kind of positioning you want to use: single or multiple. All of the examples used so far, and the explanations about current position, have used single positioning.

Specify the kind of position you want to use for each PCB on the PCB statement when you define the PSB. The POS operand for a DEDB is disregarded. DEDBs support multiple positioning only.

Single positioning

IMS maintains position in **one** hierarchic path for the hierarchy that is defined by that PCB. When you retrieve a segment, IMS clears position for all dependents and all segments on the same level.

Multiple positioning

IMS maintains position in **each** hierarchic path in the database record that is being accessed. When you retrieve a segment, IMS clears position for all dependents but keeps position for segments at the same level. You can process different segment types under the same parent in parallel.

For example, suppose you issue these two calls using the hierarchy shown in the following figure:

```

GU  Abbbbbbb(AKEYbbbb=bA1)
    Bbbbbbbb(BKEYYbbbb=bB11)
    Cbbbbbbb(CKEYYbbbb=bC111)
GN  Ebbbbbbb(EKEYYbbbb=bE11)
  
```




Figure 55. Multiple positioning hierarchy

After issuing the first call with single positioning, IMS has three positions established: one on A1, one on B11, and one on C111. After issuing the second call, the positions on B11 and C111 are canceled. Then IMS establishes positions on A1 and E11.

After issuing the first call with single and multiple positioning, IMS has three positions established: one on A1, one on B11, and one on C111. However, after issuing the second call, single positioning cancels positions on B11 and C111 while multiple positioning retains positions on B11 and C111. IMS then establishes positions on segments A1 and E11 for both single and multiple positioning.

After issuing the first call with multiple positioning, IMS has three positions established (just as with single positioning): one on A1, one on B11, and one on C111. But after issuing the second call, the positions on B11 and C111 are retained. In addition to these positions, IMS establishes position on segments A1 and E11.




Figure 56. Single and multiple positioning hierarchy

The examples that follow compare the results of single and multiple positioning using the hierarchy in the following figure.

Table 41. Results of single and multiple positioning with DL/I calls

Sequence	Result of Single Positioning	Result of Multiple Positioning
Example 1		
GU (where AKEY equals A1)	A1	A1
GNP B	B11	B11
GNP C	C11	C11
GNP B	Not found	B12
GNP C	C12	C12
GNP B	Not found	B13
GNP C	C13	C13
GNP B	Not found	Not found
GNP C	Not found	Not found
Example 2		
GU A (where AKEY equals A1)	A1	A1
GN B	B11	B11
GN C	C11	C11
GN B	B21	B12
GN C	C21	C12
Example 3		
GU A (where AKEY equals A1)	A1	A1
GN C	C11	C11
GN B	B21	B11
GN B	B22	B12
GN C	C21	C12
Example 4		
GU A (where AKEY equals A1)	A1	A1
GN B	B11	B11
GN C	C11	C11
GN D	D111	D111
GN E	E111	E111
GN B	B21	B12
GN D	D221	D112
GN C	C under next A	C12
GN E	E under next A	E121

Multiple positioning is useful when you want to examine or compare segments in two hierarchic paths. It lets you process different segment types under the same parent in parallel. Without multiple positioning, you would have to issue GU calls to reestablish position in each path.

Advantages of using multiple positioning

The advantages of using multiple positioning include the following:

- You might be able to design your program with greater data independence than you would using single positioning. You can write application programs that use GN and GNP calls, and GU and ISRT calls with missing levels in their SSAs, independent of the relative order of the segment types being processed. If you improve your program's performance by changing the relative order of segment types and all of the application programs that access those segment types use multiple positioning, you could make the change without affecting existing application programs. To do this without multiple positioning, the program would have to use GN and GNP calls, and GU and ISRT calls with incompletely specified SSAs.
- Your program can process dependent segment types in parallel (it can switch back and forth between hierarchic paths without reissuing GU calls to reset position) more efficiently than is possible with single positioning. You indicate to IMS the hierarchic path that contains the segments you want in your SSAs in the call. IMS uses the position established in that hierarchic path to satisfy your call. The control blocks that IMS builds for each kind of positioning are the same. Multiple positioning does not require more storage, nor does it have a big impact on performance.

Remember: Multiple positioning might use more processor time than single positioning, and that multiple positioning cannot be used with HSAM databases.

How multiple positioning affects your program

Multiple positioning affects the order and structure of your DL/I calls.

GU and ISRT

The only time multiple positioning affects GU and ISRT calls is when you issue these calls with missing SSAs in the hierarchic path. When you issue a GU or ISRT call that does not contain an SSA for each level in the hierarchic path, IMS builds the SSA for the missing levels according to the current position:

- If IMS has a position established at the missing level, the qualification IMS uses is derived from that position, as reflected in the DB PCB.
- If no position is established at the missing level, IMS assumes a segment type for that level.
- If IMS moves forward from a position that is established at a higher level, it assumes a segment type for that level.

Because IMS builds the missing qualification based on current position, multiple positioning makes it possible for IMS to complete the qualification independent of current positions that are established for other segment types under the same parent occurrence.

DLET and REPL with multiple positioning

Multiple positioning does not affect DLET or REPL calls; it only affects the Get Hold calls that precede them.

Qualified GN and GNP calls

When your program issues a GN or GNP call, IMS tries to satisfy the call by moving forward from current position. When you use multiple positioning, more than one current position exist: IMS maintains a position at each level in **all** hierarchic paths, instead of at each level in **one** hierarchic path. To satisfy GN and GNP calls with multiple positioning, IMS moves forward from the current position in the path that is referred to in the SSA.

Mixing qualified and unqualified GN and GNP calls

Although multiple positioning is intended to be used with qualified calls for parallel processing and data independence, you may occasionally want to use unqualified calls with multiple positioning. For example, you may want to sequentially retrieve all of the segment occurrences in a hierarchy, regardless of segment type.

Tip: Limit unqualified calls to GNP calls in order to avoid inconsistent results. Mixing qualified and unqualified SSAs may be valid for parallel processing, but doing so might also decrease the program's data independence.

The following rules apply to mixing qualified and unqualified GN and GNP calls:

1. When you issue an unqualified GN or GNP, IMS uses the position that is established by the preceding call to satisfy the GN or GNP call. For example:

Your program issues these calls:	DL/I returns these segments:
GU A (where AKEY = A1)	A1
GN B	B11
GN E	E11
GN	F111

When your program issues the unqualified GN call, IMS uses the position that is established by the last call, the call for the E segment, to satisfy the unqualified call.

2. After you successfully retrieve a segment with an unqualified GN or GNP, IMS establishes position in only one hierarchic path: the path containing the segment just retrieved. IMS cancels positions in other hierarchic paths. IMS establishes current position on the segment that is retrieved and sets parentage on the parent of the segment that is retrieved. If you issue a qualified call for a segment in a different hierarchic path after issuing an unqualified call, the results are unpredictable. For example:

Your program issues these calls:	DL/I returns these segments:
GU A (where AKEY = A1)	A1
GN B	B11
GN E	E11
GN	F111
GN B	unpredictable

When you issue the unqualified GN call, IMS no longer maintains a position in the other hierarchic path, so the results of the GN call for the B segment are unpredictable.

3. If you issue an unqualified GN or GNP call and IMS has a position established on a segment that the unqualified call might encounter, the results of the call are unpredictable. Also, when you issue an unqualified call and you have established position on the segment that the call "should" retrieve, the results are unpredictable.

For example:

Your program issues these calls:	DL/I returns these segments:
GU A (where AKEY = A1)	A1
GN E	E11
GN D	D111
GN B	B12
GN B	B13

Your program issues these calls:	DL/I returns these segments:
GN	E11 (The only position IMS has is the one established by the GN call.)

In this example, IMS has a position established on E11. An unqualified GN call moves forward from the position that is established by the previous call. Multiple positions are lost; the only position IMS has is the position that is established by the GN call.

To summarize these rules:

1. To satisfy an unqualified GN or GNP call, IMS uses the position established in the last call for that PCB.
2. If an unqualified GN or GNP call is successful, IMS cancels positions in all other hierarchic paths. Position is maintained only within the path of the segment retrieved.

Resetting position with multiple positioning

To reset position, your program issues a GU call for a root segment. If you want to reset position in the database record you are currently processing, you can issue a GU call for that root segment, but the GU call cannot be a path call.

Example: Suppose you have positions established on segments B11 and E11. Your program can issue one of the calls below to reset position on the next database record.

Issuing this call causes IMS to cancel all positions in database record A1:

```
GU  AbbbbbbbAKEYbbbb=bA2)
```

Or, if you wanted to continue processing segments in record A1, you issue this call to cancel all positions in record A1:

```
GU  AbbbbbbbAKEYbbbb=bA1)
```

Issuing this call as a path call does not cancel position.

Multiple DB PCBs

When a program has multiple PCBs, it usually means that you are defining views of several databases, but this also can mean that you need several positions in one database record. Defining multiple PCBs for the same hierarchic view of a database is another way to maintain more than one position in a database record.

Using multiple PCBs also extends what multiple positioning does, because with multiple PCBs you can maintain positions in two or more database records and within two or more hierarchic paths in the same record.

For example, suppose you were processing the database record for Patient A. Then you wanted to look at the record for Patient B and also be able to come back to your position for Patient A. If your program uses multiple PCBs for the medical hierarchy, you issue the first call for Patient A using PCB1 and then issue the next call, for Patient B, using PCB2. To return to Patient A's record, you issue the next call using PCB1, and you are back where you left off in that database record.

Using multiple PCBs can decrease the number of Get calls required to maintain position and can sometimes improve performance. Multiple PCBs are particularly useful when you want to compare information from segments in two or more database records. However, the internal control block requirements increase with each PCB that you define.

You can use the AIBTDLI interface with multiple PCBs by assigning different PCBNAMEs to the PCBs during PSB generation. Just as multiple PCBs must have different addresses in the PSB PCBLIST, multiple PCBs must have different PCBNAMEs when using the AIBTDLI interface. For example, if your application program issues DL/I calls against two different PCBs in a list that identifies the same database, you

achieve the same effect with the AIBTDLI interface by using different PCBNAMEs on the two PCBs at PSB generation time.

Chapter 16. Using IMS application program sync points

IMS application programs can (and should) take checkpoints. These checkpoints and system sync points can affect IMS operations.

Commit process

During the synchronization point (sync point) processing for an application, IMS creates a log record to establish commitment of database changes and availability of output messages. The commit process is not complete until IMS physically writes this log record to the OLDS because an incomplete set of database change and message records exist on the log for system restart.

The commit processes work differently for full-function and Fast Path applications. For full-function, IMS makes database changes in the buffer pool at the time of a DL/I call, and can write the changes to disk before the commit point. If you restart the system, IMS backs out these uncommitted changes by using the log. IMS stores inserted message segments in the message queue and must similarly discard them.

For Fast Path, IMS keeps all changes in memory until it physically logs the commit record. Only then does IMS write database changes to DASD and send output messages. Because no changes appear on external storage (except for the log) until the commit record is written, IMS does not perform backout processing for the database. IMS discards the updates in memory. With Fast Path, system restart ensures that IMS writes committed updates to DASD and sends output messages.

Relationship between checkpoints and sync points

IMS tracks all checkpoints and sync points. IMS usually uses a sync point during recovery, but returns to the checkpoint in the following situations: In the following figure, for example, if a system-wide failure occurs in the DB/DC environment just after the MTO takes a system checkpoint but just before program B commits (assuming that program A has not made any updates since its last commit), IMS must return to the system checkpoint before Beta started.

- For a full recovery in the DB/DC environment, IMS returns to the earliest of either the checkpoint before the current checkpoint or the checkpoint before the first uncommitted application program update.
- For a full recovery in the DBCTL environment, IMS always returns to the checkpoint before the first uncommitted application program update.
- For a full recovery in the DCCTL environment, IMS always returns to the checkpoint before the latest system checkpoint.
- In the DB/DC or DCCTL environments, if a BUILDQ is requested on the restart, IMS returns to the last SNAPQ or DUMPQ checkpoint. IMS returns to this checkpoint even if it is older than the checkpoint normally needed for the restart.




Figure 57. Independence of system checkpoints and application sync points

Synchronization point processing in CPI Communications-driven programs

For CPI Communications-driven programs running under Advanced Program-to-Program Communications for IMS (APPC/IMS), the application programs control their own sync point processing. An application program can issue certain CPI Resource Recovery calls: SRRCMIT calls to commit data and SRRBACK calls to back out data. The protected resources managed by IMS (local) include:

- IMS TM message-queue messages
- IMS DB databases
- Db2 for z/OS databases

The highest level of synchronization supported for a conversation is SYNCPT, so CPI Communications-driven applications can have protected conversations.

Sync point and resource manager

IMS can be either the sync point manager or the resource manager, depending on the setting of the sync point level. For SYNCLVL=NONE or CONFIRM and AOS=B, S, or X, IMS is the sync point manager and the resource manager, but for RRS=Y and SYNCLVL=SYNCPT, z/OS Resource Recovery Services (RRS) is the sync point manager and IMS is the resource manager. For RRS=N, IMS is the sync point manager.

Two-phase commit in the synchronization process

Application programs in a DBCTL, DCCTL, DB/DC, APPC/IMS, or OTMA environment can be involved in a two-phase commit process to record a sync point. At the completion of a two-phase commit, the resource manager commits database and message changes.

The two phases are:

1. Phase 1, in which the sync-point coordinator directs sync point preparation and asks the connected resource managers whether updates to connected databases can be committed.

The sync-point coordinator can be:


- An IMS DB/DC subsystem for its resource managers and attached databases.
- An IMS DCCTL subsystem for attached databases.
- A Coordinator Controller (CCTL) subsystem for units of work associated with the CCTL region. IMS DB acts as a resource manager when connected to a CCTL and also when accessed by ODBA application programs through the Open Database Access (ODBA) interface.

- z/OS Resource Recovery Services (RRS) for its protected conversations with APPC/IMS applications programs or OTMA clients. IMS acts as a resource manager when connected to RRS.
2. Phase 2, in which the sync-point coordinator directs commit or abort processing and states that the resources must either be committed or aborted.

In the DBCTL environment, if an application program makes no update DL/I calls or makes only inquiry-type DL/I calls, the CCTL requests a "forget" response to Phase 1 (if forget processing has been enabled). This means that only a limited Phase 2 occurs for that application program because no database resources have been altered. See *IMS Version 15.3 Exit Routines* for details on how to enable forget processing.

The sync-point coordinator can request an abort without a Phase 1.

The following figure shows the two phases of the sync-point cycle for an IMS DBCTL environment and describes the activities taking place.



Notes:

1. If the resource manager indicates that it cannot commit the updates, the sync-point coordinator should abort the unit of recovery, and the rest of this figure does not apply.
2. If the sync-point coordinator tells the resource manager to commit the updates, then it **must** commit.

Figure 58. Two-phase commit process

Unit of recovery

A *unit of recovery* (UOR) is the work done by a thread (connection between a resource-manager control region and a sync-point coordinator) during a sync-point interval, that is between two sync points.

In-flight unit of recovery

The unit of recovery is said to be *in-flight* from its creation or its last sync point until the resource manager logs the end of Phase 1. If a resource manager fails before or during Phase 1 and is subsequently restarted, IMS aborts all database updates.

In-doubt unit of recovery for DBCTL connected to CCTL

From the time that the resource manager issues its response to the PREPARE request (the completion of Phase 1), to the time it receives a COMMIT or ABORT request from the CCTL, units of recovery are said to be *in-doubt*. When the resource manager is restarted after a failure, it tells the CCTL which in-doubt UORs exist, if any. The CCTL then takes action to resolve these in-doubt UORs. This is called resolve in-doubt processing, or resynchronization. If a CCTL cannot resolve all in-doubt UORs, you can use IMS or CCTL commands to display the units of recovery and take appropriate actions for committing or aborting them.

Recovery tokens for DBCTL connected to CCTL

A *recovery token* is a 16-byte identifier for each unit of recovery. The resource manager validates the recovery token to protect against duplication of units of recovery. In the DBCTL environment, you can display the recovery token using the IMS /DISPLAY CCTL command. The recovery token is the primary identifier used by DBRC, which performs unit-of-recovery management. DBRC keeps track of backouts that are appropriate for the Batch Backout utility to perform.

Recoverable in-doubt structure

An IMS DBCTL subsystem builds a recoverable in-doubt structure (RIS) for each in-doubt UOR when any of the following occurs:

- A CCTL fails
- A CCTL thread fails
- A resource manager fails

The resource manager uses a recoverable in-doubt structure during reconnecting to the CCTL if in-doubt UORs existed when either the CCTL or the resource manager failed. IMS logs all recoverable in-doubt structures during system checkpoints.

A recoverable in-doubt structure contains the following information:

- The recovery token in a residual recovery element (RRE)
- Changed data records in an in-doubt extended error queue element (IEEQE)
- An indication of data that is inaccessible because of unresolved in-doubt UORs
- Links to other recoverable in-doubt structures using extended error queue element (EEQE) queue elements (EQELs)

DBCTL single-phase commit

A CCTL communicating with just one resource manager (IMS DBCTL subsystem) can request a sync point using just a single phase. If the CCTL communicates with more than one resource manager, it must use the two-phase commit process.

When the CCTL decides to commit a UOR, it can request a single-phase sync point. Single-phase commit can affect the recoverability of in-doubt data. A transaction is only in-doubt for the short time between the sync-point request and DBCTL's commit. IMS can recover in-doubt data after a thread failure during single-phase commit, but cannot recover in-doubt data after a subsystem failure.

Sync-point log records

During the two-phase commit process, IMS creates log records to establish the commitment of database changes. All these log records can be used by the IMS Change Accumulation and recovery utilities.

All online log records involving the sync-point cycle contain a recovery token. This token ensures that IMS can recover and restart each unit of recovery. The sequence of log records for a unit of recovery reveals the sync-point cycle that it followed.

IMS logs the following records during the sync-point process:

Log record	Description

X'08'	Schedule record

X'07'	Unschedule (terminate) record

X'0A08'	CPI Communications-driven application program schedule record

X'0A07'	CPI Communications-driven application program unschedule (terminate) record

X'5945'	Fast Path 64-bit buffer usage

X'5937'	Fast Path start commit

X'5938'	Fast Path start abort

X'5610'	Start of Phase 1

X'5611'	End of Phase 1

X'3730'	Start of Phase 2 Commit

X'5612'	End of Phase 2 Commit

X'3801'	Start of abort

X'4C01'	End of abort

X'5607'	Start unit of recovery

X'5613'	Recoverable in-doubt structure created

X'5614'	Recoverable in-doubt structure deleted

Sync points with a data-propagation manager

When using a data-propagation manager (such as the IMS DataPropagator) to update Db2 for z/OS databases synchronously with IMS DL/I databases, the updates to the Db2 for z/OS databases are committed (or aborted) at the same time as the IMS updates. This provides consistency between the database management subsystems. IMS DB/DC, DCCTL, and DBCTL (BMP regions only) support the IMS Data Capture exit routine.

Restriction: In an IMS DBCTL environment, the data-propagation manager is available only for BMP regions.

For more information about the IMS DataPropagator, go to the following web URL: <http://www.ibm.com/software/data/db2imstools/imstools/imspdprop.html>

Chapter 17. Recovering databases and maintaining database integrity

You can issue checkpoints, restart programs, and maintain database integrity in your application programs.

Java applications running in Java batch processing (JBP) regions can issue symbolic checkpoint and restart calls by using the IMS Java dependent region resource adapter.

Related concepts

[“Developing JBP applications with the IMS Java dependent region resource adapter” on page 683](#)
JBP applications are similar to JMP applications, except that JBP applications do not receive input messages from the IMS message queue. Unlike batch message processing (BMP) applications, JBP applications must be non-message-driven applications.

Issuing checkpoints

Two kinds of checkpoint (CHKP) calls exist: the basic CHKP and the symbolic CHKP. All IMS programs and CICS shared database programs can issue the basic CHKP call; only BMPs and batch programs can use either call.

IMS Version 15.3 Application Programming APIs explains when and why you should issue checkpoints in your program. Both checkpoint calls cause a loss of database position when the call is issued, so you must reestablish position with a GU call or some other method. You cannot reestablish position in the middle of non-unique keys or nonkeyed segments.

Restriction: You must not specify CHKPT=EOV on any DD statement to take an IMS checkpoint.

Some differences exist if you issue the same call sequence against a full-function database or a DEDB, and an MSDB.

Depending on the database organization, a CHKP call can result in the database position for the PCB being reset. When the CHKP call is issued, the locks held by the program are released. Therefore, if locks are necessary for maintaining your database position, the position is reset by the CHKP call. Position is reset in all cases except those in which the organization is either GSAM (locks are not used) or DEDB, and the CHKP call is issued after a GC status code. For a DEDB, the position is maintained at the unit-of-work boundary.

Issuing a CHKP resets the destination of the modifiable alternate PCB.

Related Reading: For more information on CHKP calls, see the topic "CHKP (Basic) Call" and the topic "CHKP (Symbolic) Call" in *IMS Version 15.3 Application Programming APIs*.

Related concepts

[“Commit-point processing in MSDBs and DEDBs” on page 302](#)

Your existing application programs can use either the MSDB commit view or the default DEDB commit view.

Restarting your program from the latest checkpoint

If you use basic checkpoints instead of symbolic checkpoints, provide the necessary code to restart the program from the latest checkpoint if the program terminates abnormally.

One way to restart the program from the latest checkpoint is to store repositioning information in a HDAM or PHDAM database. With this method, your program writes a database record containing repositioning information to the database each time a checkpoint is issued. Before your program terminates, it should delete the database record.

For more information on the XRST call, see the topic "XRST Call" in *IMS Version 15.3 Application Programming APIs*.

Maintaining database integrity (IMS batch, BMP, and IMS online regions)

IMS uses these DL/I calls to back out database updates: ROLB, ROLL, ROLS, SETS, and SETU.

The ROLB and ROLS calls can back out the database updates or cancel the output messages that the program has created since the program's most recent commit point. A ROLL call backs out the database updates and cancels any non-express output messages the program has created since the last commit point. It also deletes the current input message. SETS allows multiple intermediate backout points to be noted during application program processing. SETU operates like SETS except that it is not rejected by unsupported PCBs in the PSB. If your program issues a subsequent ROLS call specifying one of these points, database updates and message activity performed since that point are backed out.

CICS online programs with DBCTL can use the ROLS and SETS or SETU DL/I calls to back out database changes to a previous commit point or to an intermediate backout point.

Backing out to a prior commit point: ROLL, ROLB, and ROLS

When a program determines that some of its processing is invalid, some calls enable the program to remove the effects of its incorrect processing. These are the Roll Back calls: ROLL, ROLS using a DB PCB (or ROLS without an I/O area or token), and ROLB.

When you issue one of these calls, IMS:

- Backs out the database updates that the program has made since the program's most recent commit point.
- Cancels the non-express output messages that the program has created since the program's most recent commit point.

The main difference between these calls is that ROLB returns control to the application program after backing out updates and canceling output messages, ROLS does not return control to the application program, and ROLL terminates the program with an abend code of U0778. ROLB can return the first message segment to the program since the most recent commit point, but ROLL and ROLS cannot.

The ROLL and ROLB calls, and the ROLS call without a specified token, are valid when the PSB contains PCBs for GSAM data sets. However, segments inserted in the GSAM data sets since the last commit point are not backed out by these calls. An extended checkpoint-restart can be used to reposition the GSAM data sets when restarting.

You can use a ROLS call either to back out to the prior commit point or to back out to an intermediate backout point that was established by a prior SETS call. This section refers only to the form of the ROLS call that backs out to the prior commit point. For information about the other form of ROLS, see 'Backing out to an intermediate backout point: SETS, SETU, and ROLS'.

The table below summarizes the similarities and the differences between the ROLB, ROLL, and ROLS calls.

Table 42. Comparison of ROLB, ROLL, and ROLS

Actions Taken:	ROLB	ROLL	ROLS
Back out database updates since the last commit point.	X	X	X
Cancel output messages created since the last commit point.	X ¹	X ¹	X ¹
Delete from the queue the message in process. Previous messages (if any) processed since the last commit point are returned to the queue to be reprocessed.		X	

Table 42. Comparison of ROLB, ROLL, and ROLS (continued)

Actions Taken:	ROLB	ROLL	ROLS
Return the first segment of the first input message issued since the most recent commit point.	X ²		
U3303 abnormal termination. Returns the processed input messages to the message queue.			X ³
U0778 abnormal termination. No dump.		X	
No abend. Program continues processing.	X		

Notes:

1. ROLB, ROLL, or ROLS calls cancel output messages that are sent with an express PCB unless the program issued a PURG. For example, if the program issues the call sequence that follows, MSG1 would be sent to its destination because PURG tells IMS that MSG1 is complete and the I/O area now contains the first segment of the next message (which in this example is MSG2). MSG2, however, would be canceled.

```
ISRT    EXPRESS PCB, MSG1
PURG    EXPRESS PCB, MSG2
ROLB    I/O PCB
```

Because IMS has the complete message (MSG1) and because an express PCB is being used, the message can be sent before a commit point.

2. Returned only if you supply the address of an I/O area as one of the call parameters.
3. The transaction is suspended and requeued for subsequent processing.

ROLL call

A ROLL call backs out the database updates and cancels any non-express output messages the program has created since the last commit point. It also deletes the current input message. Any other input messages that were processed since the last commit point are returned to the queue to be reprocessed. IMS then terminates the program with an abend code U0778. This type of abnormal termination terminates the program without a storage dump.

When you issue a ROLL call, the only parameter you supply is the call function, ROLL.

You can use the ROLL call in a batch program. If your system log is on DASD, and if dynamic backout has been specified through the use of the BKO execution parameter, database changes made since the last commit point will be backed out; otherwise they will not. One reason for issuing ROLL in a batch program is for compatibility.

After backout is complete, the original transaction is discarded if it can be, and it is not re-executed. IMS issues the APPC/MVS verb, ATBCMTP TYPE(ABEND), specifying the TPI to notify remote transaction programs. Issuing the APPC/MVS verb causes all active conversations (including any that are spawned by the application program) to be DEALLOCATED TYP(ABEND_SVC).

ROLB call

The advantage of using a ROLB call is that IMS returns control to the program after executing a ROLB call, so the program can continue processing. The parameters for the ROLB call are:

- The call function, ROLB
- The name of the I/O PCB or AIB

The total effect of the ROLB call depends on the type of IMS application program that issued it.

- For current IMS application programs:

After IMS backout is complete, the original transaction is represented to the IMS application program. Any resources that cannot be rolled back by IMS are ignored; for example, output that is sent to an express alternate PCB and a PURG call that is issued before the ROLB call.

- For modified IMS application programs:

The same consideration for the current IMS application program applies. The application program must notify any spawned conversations that a ROLB was issued.

- For CPI-C driven IMS application programs:

Only IMS resources are affected. All database changes are backed out. Any messages that are inserted to non-express alternate PCBs are discarded. Also, any messages that are inserted to express PCBs that have not had a PURG call are discarded. The application program must notify the originating remote program and any spawned conversations that a ROLB call was issued.

MPPs and transaction-oriented BMPs

If the program supplies the address of an I/O area as one of the ROLB parameters, the ROLB call acts as a message retrieval call and returns the first segment of the first input message issued since the most recent commit point. This is true only if the program has issued a GU call to the message queue since the last commit point; if it has not, it was not processing a message when it issued the ROLB call.

If the program issues GN call to the message queue after issuing a ROLB call, IMS returns the next segment of the message that was being processed when the ROLB call was issued. If no more segments exist for that message, IMS returns a QD status code.

If the program issues a GU call to the message queue after the ROLB call, IMS returns the first segment of the next message to the application program. If no more messages exist on the message queue for the program to process, IMS returns a QC status code.

If you include the I/O area parameter, but you have not issued a successful GU call to the message queue since the last commit point, IMS returns a QE status code to your program.

If you do not include the address of an I/O area in the ROLB call, IMS does the same thing for you. If the program has issued a successful GU call in the commit interval and then issues a GN call, IMS returns a QD status code. If the program issues a GU call after the ROLB call, IMS returns the first segment of the next message or a QC status code, if no more messages exist for the program.

If you have not issued a successful GU call since the last commit point, and you do not include an I/O area parameter on the ROLB call, IMS backs out the database updates and cancels the output messages that were created since the last commit point.

Batch programs

If your system log is on DASD, and if dynamic backout has been specified through the use of the BKO execution parameter, you can use the ROLB call in a batch program. The ROLB call does not process messages as it does for MPPs; it backs out the database updates made since the last commit point and returns control to your program. You cannot specify the address of an I/O area as one of the parameters on the call; if you do, an AD status code is returned to your program. You must, however, have an I/O PCB for your program. Specify CMPAT=YES on the CMPAT keyword in the PSBGEN statement for your program's PSB.

ROLS call

You can use the ROLS call in two ways to back out to the prior commit point and return the processed input messages to IMS for later reprocessing:

- Have your program issue the ROLS call using the I/O PCB but without an I/O area or token in the call. The parameters for this form of the ROLS call are:

The call function, ROLS

The name of the I/O PCB or AIB

- Have your program issue the ROLS call using a database PCB that has received one of the data-unavailable status codes. This has the same result as if unavailable data were encountered and the INIT call was not issued. A ROLS call must be the next call for that PCB. Intervening calls using other PCBs are permitted.

On a ROLS call with a TOKEN, message queue repositioning can occur for all non-express messages, including all messages processed by IMS. The processing uses APPC/MVS calls, and includes the initial message segments. The original input transaction can be represented to the IMS application program. Input and output positioning is determined by the SETS call. This positioning applies to current and modified IMS application programs but does not apply to CPI-C driven IMS programs. The IMS application program must notify all remote transaction programs of the ROLS.

On a ROLS call without a TOKEN, IMS issues the APPC/MVS verb, ATBCMTP TYPE(ABEND), specifying the TPI. Issuing this verb causes all conversations associated with the application program to be DEALLOCATED TYPE(ABEND_SVC). If the original transaction is entered from an LU 6.2 device and IMS receives the message from APPC/MVS, a discardable transaction is discarded rather than being placed on the suspend queue like a non-discardable transaction.

The parameters for this form of the ROLS call are:

- The call function, ROLS
- The name of the DB PCB that received the BA or BB status code

In both of these parameters, the ROLS call causes a U3303 abnormal termination and does not return control to the application program. IMS keeps the input message for future processing.

Related concepts

[Administering APPC/IMS and LU 6.2 devices \(Communications and Connections\)](#)

Related reference

[Program Specification Block \(PSB\) Generation utility \(System Utilities\)](#)

[ROLB call \(Application Programming APIs\)](#)

Backing out to an intermediate backout point: SETS, SETU, and ROLS

You can use a ROLS call either to back out to an intermediate backout point that was established by a prior SETS or SETU call, or to back out to the prior commit point.

The ROLS call that backs out to an intermediate point backs out only DL/I changes. This version of the ROLS call does not affect CICS changes that use CICS file control or CICS transient data.

The SETS and ROLS calls set intermediate backout points within the call processing of the application program and then backout database changes to any of these points. Up to nine intermediate backout points can be set. The SETS call specifies a token for each point. IMS then associates this token with the current processing point. A subsequent ROLS call using the same token backs out all database changes and discards all non-express messages that were performed after the SETS call with the same token. The following figure shows how the SETS and ROLS calls work together.

In addition, to assist the application program in managing other variables that it may want to reestablish after a ROLS call, user data can be included in the I/O area of the SETS call. This data is then returned when the ROLS call is issued with the same token.




Figure 59. SETS and ROLS calls working together

SETS and SETU calls

The SETS call sets up to nine intermediate backout points or cancels all existing backout points. With the SETS call, you can back out pieces of work. If the necessary data to complete one piece of work is unavailable, you can complete a different piece of work and then return to the former piece.

To set an intermediate backout point, issue the call using the I/O PCB, and include an I/O area and a token. The I/O area has the format *LLZZuser-data*, where *LL* is the length of the data in the I/O area including the length of the *LLZZ* portion. The *ZZ* field must contain binary zeros. The data in the I/O area is returned to the application program on the related ROLS call. If you do not want to save some of the data that is to be returned on the ROLS call, set the *LL* that defines the length of the I/O area to 4.

For PLITDLI, you must define the *LL* field as a fullword rather than a halfword, as it is for the other languages. The content of the *LL* field for PLITDLI is consistent with the I/O area for other calls using the *LLZZ* format. The content is the total length of the area, including the length of the 4-byte *LL* field, minus 2.

A 4-byte token associated with the current processing point is also required. This token can be a new token for this program execution, or it can match a token that was issued by a preceding SETS call. If the token is new, no preceding SETS calls are canceled. If the token matches the token of a preceding SETS call, the current SETS call assumes that position. In this case, all SETS calls that were issued subsequent to the SETS call with the matching token are canceled.

The parameters for this form of the SETS call are:

- The call function, SETS
- The name of the I/O PCB or AIB
- The name of the I/O area containing the user data
- The name of an area containing the token

For the SETS call format, see the topic "SETS/SETU Call" in *IMS Version 15.3 Application Programming APIs*.

To cancel all previous backout points, the call is issued using the I/O PCB but does not include an I/O area or a token. When an I/O area is not included in the call, all intermediate backout points that were set by prior SETS calls are canceled.

The parameters for this form of the SETS call are:

- The call function, SETS
- The name of the I/O PCB or AIB

Because it is not possible to back out committed data, commit-point processing causes all outstanding SETS to be canceled.

If PCBs for DEDB, MSDB, and GSAM organizations are in the PSB, or if the program accesses an attached subsystem, a partial backout is not possible. In that case, the SETS call is rejected with an SC status code. If the SETU call is used instead, it is not rejected because of unsupported PCBs, but will return an SC status code as a warning that the PSB contains unsupported PCBs and that the function is not applicable to these unsupported PCBs.

Related reading: For status codes that are returned after the SETS call and the explanations of those status codes and the response required, see *IMS Version 15.3 Application Programming APIs*.

ROLS

The ROLS call backs out database changes to a processing point set by a previous SETS or SETU call, or to the prior commit point. The ROLS call then returns the processed input messages to the message queue.

To back out database changes and message activity that have occurred since a prior SETS call, issue the ROLS call using the I/O PCB, and specify an I/O area and token in the call. If the token does not match a token that was set by a preceding SETS call, an error status is returned. If the token matches the token of a preceding SETS call, the database updates made since this corresponding SETS call are backed out, and all non-express messages that were inserted since the corresponding SETS are discarded. SETS that are issued as part of processing that was backed out are canceled. The existing database positions for all supported PCBs are reset.

If a ROLS call is in response to a SETU call, and if there are unsupported PCBs (DEDB, MSDB, or GSAM) in the PSB, the position of the PCBs is not affected. The token specified by the ROLS call can be set by either a SETS or SETU call. If no unsupported PCBs exist in the PSB, and if the program has not used an attached subsystem, the function of the ROLS call is the same regardless of whether the token was set by a SETS or SETU call.

If the ROLS call is in response to a SETS call, and if unsupported PCBs exist in the PSB or the program used an attached subsystem when the preceding SETS call was issued, the SETS call is rejected with an SC status code. The subsequent ROLS call is either rejected with an RC status code, indicating unsupported options, or it is rejected with an RA status code, indicating that a matching token that was set by a preceding successful SETS call does not exist.

If the ROLS call is in response to a SETU call, the call is not rejected because of unsupported options. If unsupported PCBs exist in the PSB, this is not reflected with a status code on the ROLS call. If the program is using an attached subsystem, the ROLS call is processed, but an RC status is returned as a warning indicating that if changes were made using the attached subsystem, those changes were not backed out.

The parameters for this form of the ROLS call are:

- The call function, ROLS
- The name of the I/O PCB or AIB
- The name of the I/O area to receive the user data
- The name of an area containing the 4-byte token

Related reading: For status codes that are returned after the ROLS call and the explanations of those status codes and the response require, see *IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes*.

Related concepts

[“Backing out to a prior commit point: ROLL, ROLB, and ROLS calls” on page 410](#)

When a program determines that some of its processing is invalid, you can use these calls to remove the effects of its incorrect processing: Roll Back calls ROLL, ROLS using a database PCB, ROLS with no I/O area or token, and ROLB.

Reserving segments for the exclusive use of your program

You may want to reserve a segment and prohibit other programs from updating the segment while you are using it. To some extent, IMS does this for you through resource lock management. The Q command code lets you reserve segments in a different way.

Restriction: The Q command code is not supported for MSDB organizations or for a secondary index that is processed as a database.

Resource lock management and the Q command code both reserve segments for your program's use, but they work differently and are independent of each other. To understand how and when to use the Q command code and the DEQ call, you must understand resource lock management.

The function of resource lock management is to prevent one program from accessing data that another program has altered until the altering program reaches a commit point. Therefore, you know that if you have altered a segment, no other program (except those using the GO processing option) can access that segment until your program reaches a commit point. For database organizations that support the Q command code, if the PCB processing option allows updates and the PCB holds position in a database record, no other program can access the database record.

The Q command code allows you to prevent other programs from updating a segment that you have accessed, even when the PCB that accessed the segment moves to another database record.

Related reading: For more information on the Q command code, see the topic "Q command code" in *IMS Version 15.3 Application Programming APIs*.

Chapter 18. Secondary indexing and logical relationships

Secondary indexing and logical relationships are techniques that can change your application program's view of the data. The DBA makes the decision about whether to use these options.

Examples of when you use these techniques are:

- If an application program must access a segment type in a sequence other than the sequence specified by the key field, secondary indexing can be used. Secondary indexing also can change the application program's access to or view of the data based on a condition in a dependent segment.
- If an application program requires a logical structure that contains segments from different databases, logical relationships are used.

Related concepts

[“SSA guidelines” on page 170](#)

Using SSAs can simplify your programming, because the more information you can give IMS to do the searching for you, the less program logic you need to analyze and compare segments in your program.

How secondary indexing affects your program

One instance of using a secondary index occurs when an application program needs to select database records in a sequence other than that defined by the root key.

IMS stores root segments in the sequence of their key fields. A program that accesses root segments out of the order of their key fields cannot operate efficiently.

You can index any field in a segment by defining an XDFLD statement for the field in the DBD for the database. If the Get call is not qualified on the key but uses some other field, IMS must search all the database records to find the correct record. With secondary indexing, IMS can go directly to a record based on a field value that is not in the key field.

For more information about secondary indexes and examples, see *IMS Version 15.3 Database Administration*.

SSAs with secondary indexes

If your program uses a secondary index, you can use the name of an indexed field in your SSAs. When you do this, IMS goes directly to the secondary index and finds the pointer segment with the value you specify. Then IMS locates the segment that the index segment points to in the primary database and returns the segment to your program.

To use an indexed field name in the SSA, follow these guidelines:

- Define the indexed field, using the XDFLD statement, in the DBD for the primary database during DBD generation.
- Use the name that was given on the XDFLD statement as the field name in the qualification statement.
- Specify the secondary index as the processing sequence during PSB generation. Do this by specifying the name of the secondary index database on the PROCSEQ parameter for a full-function secondary index database or the PROCSEQD parameter for a Fast Path secondary index database on the PCB during PSB generation.

If you modify the XDFLD of the indexed segment (using the REPL call), you lose any parentage that you had established before issuing the REPL call. The key feedback area is no longer valid after a successful REPL call.

For example, to index the PATIENT segment on the NAME field, the segment must have been defined on the XDFLD statement in the DBD for the medical database. If the name of the secondary index database is

INDEX, you specify PROCSEQ=INDEX in the PCB. To issue a qualification that identifies a PATIENT by the NAME field instead of by PATNO, use the name that you specified on the XDFLD statement. If the name of the XDFLD is XNAME, use XNAME in the SSA, as follows:

In the DBD:

XDFLD NAME=XNAME

In the PSB:

PROCSEQ=INDEX for full-function secondary index databases or PROCSEQD=INDEX for Fast Path secondary index databases

In the program:

GU PATIENTb(XNAMEbbb=bJBBROKEbbb)

A qualified GU/GN segment name with SSA using the primary key field for target=root segment is supported when a primary DEDB database is accessed through its secondary index using a PCB with the PROCSEQD= parameter.

A qualified GU/GN segment name with SSA using the primary key field for target=dependent segment is not supported. An AC status code is returned for the qualified Get call when a primary DEDB database is accessed through its secondary index using a PCB with the PROCSEQD= parameter.

Multiple qualification statements with secondary indexes

When you qualify a call using the name of an indexed field, you can include multiple qualification statements.

You can use two AND operators to connect the qualification statements:

*** or &**

When used with secondary indexing, this AND is called the dependent AND. To satisfy the call, IMS scans the index once and searches for one pointer segment in the index that satisfies both qualification statements.

#

This is called the independent AND. You use it only with secondary indexing. When you use the independent AND to satisfy the call, IMS scans the index twice and searches for two or more different pointer segments in the index that point to the same target segment.

The distinction between the two ANDs applies only when the indexed field (the one defined as XDFLD in the DBD) is used in all qualifications. If one of the qualification statements uses another field, both ANDs work like the dependent AND.

The next two sections give examples of the dependent and independent AND. Although the examples show only two qualification statements in the SSA, you can use more than two. No set limit exists for the number of qualification statements you can include in an SSA, but a limit on the maximum size of the SSA does exist. You specify this size on the SSASIZE parameter of the PSBGEN statement. For information on this parameter, see *IMS Version 15.3 System Utilities*.

The dependent AND

When you use the dependent AND, IMS scans the index only once. To satisfy the call, it must find **one** pointer segment that satisfies both qualification statements.

For example, suppose you want to list patients whose bills are between \$500 and \$1000. To do this, you index the PATIENT segment on the BILLING segment, and specify that you want IMS to use the secondary index as the processing sequence. The following figure shows the three secondary indexing segments.

XDFLD=XBILLING




Figure 60. Example of using the dependent AND

You then use this call:

```
GU PATIENT (XBILLING>=00500*XBILLING<=01000)
```

To satisfy this call, IMS searches for one pointer segment with a value between 500 and 1000. IMS returns the PATIENT segment that is pointed to by that segment.

The independent AND

For example, suppose you want a list of the patients who have had both tonsillitis and strep throat. To get this information, you index the PATIENT segment on the ILLNAME field in the ILLNESS segment, and specify that you want IMS to use the secondary index as the processing sequence. In this example, you retrieve the PARENT segments based on a dependent's (the ILLNESS segment's) qualification. The following figure shows the four secondary indexing segments.




Figure 61. Example of using the independent AND

You want IMS to find two pointer segments in the index that point to the same PATIENT segment, one with ILLNAME equal to TONSILLITIS and one with ILLNAME equal to STREPTHRT. Use this call:

```
GU PATIENTb(XILLNAME=TONSILLITIS#XILLNAME=bSTREPTHRT)
```

This call retrieves the first PATIENT segment with ILLNESS segments of strep throat and tonsillitis. When you issue the call, IMS searches for an index entry for tonsillitis. Then it searches for an index entry for strep throat that points to the same PATIENT segment.

When you use the independent AND with GN and GNP calls, a special situation can occur. If you repeat a GN or a GNP call using the same qualification, it is possible for IMS to return the same segment to your program more than once. You can check to find out whether IMS has already returned a segment to you by checking the key feedback area.

If you continue issuing a GN call until you receive a not-found (GE) status code, IMS returns a segment occurrence once for each independent AND group. When IMS returns a segment that is identical to one that was already returned, the PCB key feedback area is different.

Related concepts

[“Multiple qualification statements” on page 171](#)

When you use a qualification statement, you can do more than give IMS a field value with which to compare the fields of segments in the database. You can give several field values to establish limits for the fields you want IMS to compare.

DL/I returns with secondary indexes

The term “key of the pointer segment” refers to the key as perceived by the application program. That is, the key does not include subsequent fields. IMS places this key in the position where the root key would be located if you had not used a secondary index—in the left-most bytes of the key feedback area.

The PATIENT segment that IMS returns to the application program's I/O area looks just as it would if you had not used secondary indexing. The key feedback area, however, contains something different. The concatenated key that IMS returns is the same, except that, instead of giving you the key for the segment you requested (the key for the PATIENT segment), IMS gives you the search portion of the key of the secondary index (the key for the segment in the INDEX database).

If you try to insert or replace a segment that contains a secondary index source field that is a duplicate of one that is already reflected in the secondary index, IMS returns an NI status code. An NI status code is returned only for batch programs that log to direct-access storage. Otherwise, the application program is abnormally terminated. You can avoid having your program terminated by making sure a duplicate index source field does not exist. Before inserting a segment, try to retrieve the segment using the secondary index source field as qualification.

Status codes for secondary indexes

If a secondary index is defined for a segment and if the definition specifies a unique key for the secondary index (most secondary indexes allow duplicate keys), your application program might receive the NI status code in addition to regular status codes.

This status code can be received for a PCB that either uses or does not use the secondary index as a processing sequence. See *IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes* for additional information about the NI status code.

Processing segments in logical relationships

Sometimes an application program needs to process a hierarchy that is made up of segments that already exist in two or more separate database hierarchies. Logical relationships make it possible to establish hierarchic relationships between these segments. When you use logical relationships, the result is a new hierarchy—one that does not exist in physical storage but that can be processed by application programs as though it does exist. This type of hierarchy is called a logical structure.

One advantage of using logical relationships is that programs can access the data as though it exists in more than one hierarchy, even though it is only stored in one place. When two application programs need to access the same segment through different paths, an alternative to using logical relationships is to store the segment in both hierarchies. The problem with this approach is that you must update the data in two places to keep it current.

Processing segments in logical relationships is not very different from processing other segments. The following examples are taken from a scenario for an inventory application program that processes data in a purchasing database, but which also needs access to a segment in a patient database.

For example, the hierarchy that an inventory application program needs to process contains four segment types:

- An ITEM segment containing the name and an identification number of a medication that is used at a medical clinic
- A VENDOR segment that contains the name and address of the vendor who supplies the item
- A SHIPMENT segment that contains information such as quantity and date for each shipment of the item that the clinic receives
- A DISBURSE segment that contains information about the disbursement of the item at the clinic, such as the quantity, the date, and the doctor who prescribed it

The TREATMNT segment in the medical database contains the same information that the inventory application program needs to process in the DISBURSE segment. Rather than store this information in both hierarchies, you can store the information in the TREATMNT segment, and define a logical relationship between the DISBURSE segment in the item hierarchy and the TREATMNT segment in the patient hierarchy. Doing this makes it possible to process the TREATMNT segment through the item hierarchy as though it is a child of SHIPMENT. DISBURSE then has two parents: SHIPMENT is DISBURSE's physical parent, and TREATMNT is DISBURSE's logical parent.

Three segments are involved in this logical relationship: DISBURSE, SHIPMENT, and TREATMNT. The following figure shows the item hierarchy on the right. The DISBURSE segment points to the TREATMNT segment in the patient hierarchy shown on the left. (The patient hierarchy is part of the medical database.)




Figure 62. Patient and item hierarchies

Three types of segments are found in a logical relationship:

- TREATMNT is called the logical parent segment. It is a physical dependent of ILLNESS, but it can be processed through the item hierarchy because a path is established by the logical child segment DISBURSE. The logical parent segment can be accessed through both hierarchies, but it is stored in only one place.
- SHIPMENT is called a physical parent segment. The physical parent is the parent of the logical child in the physical database hierarchy.
- DISBURSE is called a logical child segment. It establishes a path to the TREATMNT segment in the PATIENT hierarchy from the SHIPMENT segment in the ITEM hierarchy.

Because a logical child segment points to its logical parent, two paths exist through which a program can access the logical parent segment:

- When a program accesses the logical parent segment through the physical path, it reaches this logical parent segment through the segment's physical parent. Accessing the TREATMNT segment through ILLNESS is accessing the logical parent segment through its physical path.
- When a program accesses the logical parent segment through the logical path, it reaches this logical parent segment through the segment's logical child. Accessing the TREATMNT segment through SHIPMENT is accessing the logical parent segment through its logical path.

When a logical parent segment is accessed through the logical child, the logical child is concatenated with both the data from its logical parent segment and any data the user has chosen to associate with this pairing (intersection data) in a single segment I/O area, like this:




Figure 63. Concatenated segment

LL is the length field of the logical parent if this segment is a variable-length segment.

How logical relationships affect your programming

The calls you issue to process segments in logical relationships are the same calls that you use to process other segments. However, the processing is different depending on how the logical segment looks in your I/O area, what the DB PCB mask contains after a retrieve call, and how you can replace, delete, and insert physical and logical parent segments.

Because it is possible to access segments in logical relationships through the logical path or the physical path, the segments must be protected from being updated by unauthorized programs.

When DBAs define logical relationships, they define a set of rules that determine how the segments can be deleted, replaced, and inserted. Defining these rules is a database design decision. If your program processes segments in logical relationships, the DBA (or the person at your installation responsible for database design) should tell you:

- What segments look like in your I/O area when you retrieve them
- Whether your program is allowed to update and insert segments
- What to do if you receive a DX, IX, or RX status code

The requirements for inserting a logical child segment are:

- In load mode, the logical child can be inserted only under its physical parent. You do not supply the logical parent in the I/O area.
- In update mode, the format of the logical child is different, depending on whether it is accessed from its physical parent or from its logical parent.
 - If accessed from its physical parent, the logical child's format is the concatenated key of the logical parent followed by intersection data.
 - If accessed from its logical parent, the logical child's format is the concatenated key of the physical parent, followed by intersection data.
- The logical child can be inserted or replaced, depending on the insert rule for the logical or physical parent. Unless the insert rule of the logical or physical parent is PHYSICAL, the logical or physical parent must be supplied in the I/O area following the logical child.

Related concepts

[“Multiple qualification statements” on page 171](#)

When you use a qualification statement, you can do more than give IMS a field value with which to compare the fields of segments in the database. You can give several field values to establish limits for the fields you want IMS to compare.

Status codes for logical relationships

These status codes apply specifically to segments that are involved in logical relationships.

These are not all of the status codes that you can receive when processing a logical child segment or a physical or logical parent. If you receive one of these status codes, it means that you are trying to update the database in a way that you are not allowed to. Check with the DBA or person responsible for implementing logical relationships at your installation to find out what the problem is.

DX

IMS did not delete the segment because the physical delete rule was violated. If the segment is a logical parent, it still has active logical children. If the segment is a logical child, it has not been deleted through its logical path.

IX

You tried to insert either a logical child segment or a concatenated segment. If it was a logical child segment, the corresponding logical or physical parent segment does not exist. If it was a concatenated segment, either the insert rule was physical and the logical or physical parent does not exist, or the insert rule is virtual and the key of the logical or physical parent in the I/O area does not match the concatenated key of the logical or physical parent.

RX

The physical replace rule has been violated. The physical replace rule was specified for the destination parent, and an attempt was made to change its data. When a destination parent has the physical replace rule, it can be replaced only through the physical path.

Chapter 19. HALDB selective partition processing

You can restrict the processing of DL/I calls to a single HALDB partition or a range of HALDB partitions by using a DD statement with the ddname DFSHALDB to pass control statements. DFS HALDB must be provided in the JCL of the batch job, the BMP (Batch Message Processing dependent online region), or the JBP (Java Batch Processing dependent online region).

Control Statements for HALDB selective partition processing

➔ HALDB PCB= (nnnn , ppppppp NUM=yyy) ➔
ddddddd

Each HALDB control statement must have a PCB keyword that contains the required parameters. The required parameters for an individual control statement must be on one line; no continuation is allowed. The input can consist of multiple HALDB control statements. There should be no duplication of DB PCB numbers. In the event of a duplication, the control statement that has been read the most recently overrides the previous statement.

Any HALDB control statement that is syntactically correct results in an entry within a table. The maximum number of entries in the table is 20. All subsequent statements that are read, even though syntactically correct, are ignored and result in a U0201 abend, unless a statement is a duplicate of an entry that is already in the table.

Parameter descriptions for HALDB selective partition processing

nnnn

The DB PCB number as the relative number of the DB PCB defined in the PSB.

ddddddd

The DB PCB label or name.

ppppppp

The partition name. This parameter is required.

NUM=yyy

The range of consecutive partitions that this PCB is restricted to using, starting with the named partition. The range of consecutive partitions is defined as the partition selection order, which is the next partition selected starting from the target partition named in the DFSHALDB statement. The next partition is determined using either the high keys defined for the HALDB or the processing order defined by the partition selection exit. This parameter is optional.

The following examples show how to use HALDB selective partition processing statements.

DFSHALDB for single partition restriction

```
HALDB PCB=(4,POHIDKA)
HALDB PCB=(PCBNUM2,POHIDJA)
```

DFSHALDB for range partition restriction

```
HALDB PCB=(3,PVHDJ5A,NUM=4)
HALDB PCB=(PCBNUM7,PVHDJ5B,NUM=3)
```

Report generated for HALDB selective partition processing

When you use HALDB selective partition processing, a report called “HALDB Selective Partition Processing” is generated in the SYSHALDB data set. This report shows the control statements that have been issued and the reason for accepting or rejecting each statement. Control statements that have

been validated and accepted are shown as “Syntactically correct.” Other messages that might appear for syntactically correct statements, and their accompanying messages, are shown in the following table:

<i>Table 43. Messages provided in the report generated for HALDB selective partition processing</i>	
Message	Explanation
Duplicate, overrides previous statement	A HALDB statement for the same PCB was already found. The current statement overrides the previous HALDB statement.
Ignored, number of valid statements exceeds 20	More than 20 HALDB statements were provided, but only 20 statements are allowed. Reduce the number of HALDB statements to 20 or fewer, and run the job again. This message results in an abend U0201.
NUM parameter must be non-zero numeric	The partition range specified in the NUM keyword must be a non-zero value from 1 to 999.
NUM value exceeds three digits	The partition range specified in the NUM keyword must be a non-zero value from 1 to 999.
An equal sign must follow NUM keyword	An equal sign must follow the NUM keyword in the HALDB statement. Add an equal sign to the HALDB statement.
The NUM keyword is missing	A comma was found after the partition name, but the NUM keyword was not present. Either verify the syntax of the positional parameters in the HALDB statement, or add the NUM keyword and the range of partitions for the restriction.
NUM parameter is missing	The NUM keyword was found, but the NUM parameter value was not present. Either verify the syntax of the positional parameters in the HALDB statement, or add the NUM keyword and the range of partitions for the restriction.

For HALDB control statements that are not syntactically correct (statements that are processed and rejected), the messages and explanations that are issued are shown in the following table:

<i>Table 44. Messages provided in the report generated for syntactically incorrect HALDB statements</i>	
Message	Explanation
No HALDB statement type	The DFSHALDB data set did not contain a HALDB statement. Add a HALDB statement to prevent this error.
A space must follow HALDB statement type	The HALDB statement requires a space after HALDB and before the PCB keyword.
PCB keyword missing	The required keyword PCB was not found. The PCB keyword must be present to process the HALDB statement successfully.
Equal sign must follow PCB keyword	An equal sign did not follow the PCB keyword. The equal sign must follow the PCB keyword to process the HALDB statement successfully.

Table 44. Messages provided in the report generated for syntactically incorrect HALDB statements (continued)

Message	Explanation
Open parenthesis must follow equal sign	An open parenthesis did not follow PCB=. The open parenthesis must follow the PCB= to process the HALDB statement successfully.
Second parameter may be missing	The HALDB partition must be provided. Either add the partition name, or verify that the syntax of the positional parameters is correct.
First parameter exceeds four digits	The DB PCB number cannot exceed a four-digit value. Change the DB PCB number to the correct DB PCB number.
Delimiter is not a comma	A comma is missing between parameter values. The comma is used as a delimiter for the positional parameters. Either add the comma, or verify that the syntax of the positional parameters is correct.
Partition name must start with an alpha	The HALDB partition name must begin with a alphabetic character. Add the partition name or verify the syntax of the positional parameters is correct.
Delimiter is not a close parenthesis	A closing parenthesis is missing from the HALDB statement. Add a closing parenthesis around the PCB parameters.
Partition name exceeds seven characters	The HALDB partition name must be seven or fewer characters. Either add the partition name, or verify that the syntax of the positional parameters is correct.
Invalid character in partition name	The HALDB partition name contains an invalid character. Either add the partition name, or verify that the syntax of the positional parameters is correct.
Statement contains all spaces	The HALDB statement is missing. Add a valid HALDB statement.
Invalid statement input	A HALDB statement was found, but it does not appear to be complete. Verify the syntax of the HALDB statement and the positional parameters specified.
Space must follow close parenthesis	A space must follow the closing parenthesis. Add a space after the closing parenthesis.
First parameter missing	The PCB number or label is missing. Either add the PCB name or label, or verify that the syntax of the positional parameters is correct.
Comma and part name missing	Only the PCB number or label was provided in the HALDB statement. Either add the partition name, or verify that the syntax of the positional parameters is correct.

Table 44. Messages provided in the report generated for syntactically incorrect HALDB statements (continued)

Message	Explanation
Partition name is missing	The HALDB partition name must be provided in the HALDB statement. Either add the partition name, or verify that the syntax of the positional parameters is correct.
Partition name starts with numeric	The HALDB partition name must begin with an alphabetic character. Either add the partition name, or verify that the syntax of the positional parameters is correct.
First parameter must not be zero	The PCB number must be a non-zero number. Add a non-zero number for the DB PCB number.
Comment statement	An asterisk was found in column one of the HALDB statement. This statement was skipped and considered a comment.

After all of the statements are validated, the job abnormally terminates with an abend code of U0201.

Chapter 20. Processing GSAM databases

GSAM databases are available to application programs that can run online in IMS batch message processing (BMP) regions (message-driven or non-message-driven) or Java batch processing (JBP) regions or standalone in DLIBATCH regions.

If your application program accesses GSAM databases, as you design your program consider that:

- An IMS program can retrieve records and add records to the end of the GSAM database, but the program cannot delete or replace records in the database.
- You use separate calls to access GSAM databases. (Additional checkpoint and restart considerations are involved in using GSAM.)
- Your program must use symbolic CHKP and XRST calls if it uses GSAM. Basic CHKP calls cannot checkpoint GSAM databases.
- When an IMS program uses a GSAM database, the program treats it like a sequential file. The physical z/OS access methods that GSAM databases support are BSAM on direct access and tape devices and VSAM on direct-access storage devices (DASD). VSAM data sets must be non-keyed and non-indexed entry-sequenced (ESDS). GSAM does not support temporary, SYSIN, or SYSOUT files.
- Because GSAM is a sequential non-hierarchic database, it has no segments, keys, or parentage.

Java application programs running in JBP regions can access GSAM databases by using the IMS Java dependent region resource adapter.

Related concepts

[“Data areas in GSAM databases” on page 226](#)

Generalized Sequential Access Method (GSAM) databases are available only to application programs that can run as batch programs, batch-oriented BMPs, transaction-oriented BMPS, or JBPs. The program communication block (PCB) mask and the record search argument (RSA) that you use in a GSAM database call have special formats.

Related reference

[“Accessing GSAM data from a JBP application” on page 686](#)

GSAM data are frequently referred to as z/OS data sets or as flat files. This kind of data is non-hierarchical in structure. You can access data from GSAM databases from a JBP application.

Accessing GSAM databases

The calls you use to access Generalized Sequential Access Method (GSAM) databases are different from those you use to access other IMS databases, and you can use GSAM databases for input and output.

For example, your program can read input from a GSAM database sequentially and then load another GSAM database with the output data. Programs that retrieve input from a GSAM database usually retrieve GSAM records sequentially and then process them. Applications that send output to a GSAM database must add output records to the end of the database as the program processes the records. You cannot delete or replace records in a GSAM database, and any records that you add must go at the end of the database.

PCB masks for GSAM databases

For the most part, you process GSAM databases in the same way that you process other IMS databases. You use calls that are very similar to DL/I calls to communicate your requests. GSAM describes the results of those calls in a GSAM DB PCB.

Calls to GSAM databases can use either the AIBTDLI or the PCB interface.

The DB PCB mask for a GSAM database serves the same purpose as it does for other IMS databases. The program references the fields of the DB PCB through the GSAM DB PCB mask. The GSAM DB PCB mask must contain the same fields as the GSAM DB PCB and must be of the same length.

Some differences exist between a DB PCB for a GSAM database and one for other IMS databases. Some of the fields are different, and the GSAM DB PCB has one field that the other PCBs do not. Because GSAM is not a hierarchical database, some fields in a PCB mask for other IMS databases do not have meanings in a GSAM PCB mask. The fields that are not used when you access GSAM databases are:

- The second field: segment level number
- The sixth field: segment name
- The eighth field: number of sensitive segments

Even though GSAM does not use these fields, you must define them in the order and length shown in the following table in the GSAM DB PCB mask.

When you code the fields in a DB PCB mask, name the area that contains all the fields as you do for a DB PCB. The entry statement associates each DB PCB mask in your program with a DB PCB in your program's PSB based on the order of the PCBs in the PSB. The entry statement refers to the DB PCB mask in your program by the name of the mask or by a pointer.

When you code the entry statement in:

- COBOL, Java, Pascal, C, and assembler language programs, the entry statement must list the names of the DB PCB masks in your program.
- PL/I programs, the entry statement must list the pointers to the DB PCB masks in your program.

The first PCB name or pointer in the entry statement corresponds to the first PCB. The second name or pointer in the entry statement corresponds to the second PCB, and so on.

Table 45. GSAM DB PCB mask

Descriptor	Byte length	DB/DC	DBCTL	DCCTL	DB batch	TM batch
Database name¹	8	X	X	X	X	X
Segment level number²	2	N/A	N/A	N/A	N/A	N/A
Status code³	2	X	X	X	X	X
Processing options⁴	4	X	X	X	X	X
Reserved for IMS⁵	4	X	X	X	X	X
Segment name⁶	8	N/A	N/A	N/A	N/A	N/A
Length of key feedback area and undefined-length records area⁷	4	X	X	X	X	X
Number of sensitive segments⁸	4	N/A	N/A	N/A	N/A	N/A
Key feedback area⁹	8 or 12 for large data sets.	X	X	X	X	X
Length of undefined-length records¹⁰	4	X	X	X	X	X

Note:

1. **Database Name.** The name of the GSAM DBD. This field is 8 bytes and contains character data.
2. **Segment Level Number.** Not used by GSAM, but you must code it. It is 2 bytes.
3. **Status Code.** IMS places a two-character status code in this field after each call to a GSAM database. This code describes the results of the call. IMS updates this field after each call and does not clear it between calls. The application program should test this field after each call to find out whether the call was successful. If the call was completed successfully, this field contains blanks.

4. **Processing Options.** This is a 4-byte field containing a code that tells IMS the types of calls this program can issue. It is a security mechanism in that it can prevent a particular program from updating the database, even though the program can read the database. This value is coded in the PROCOPT parameter of the PCB statement when generating the PSB for the application program. The value does not change. For GSAM, the values are G, GS, L, or LS.
5. **Reserved for IMS.** This 4-byte field is used by IMS for internal linkage. It is not used by the application program.
6. **Segment Name.** This field is not used by GSAM, but it must be coded as part of the GSAM DB PCB mask. It is 8 bytes.
7. **Length of Key Feedback Area and Undefined-Length Records Area.** This is a 4-byte field that contains the decimal value of 12 (or 16 for large format data sets). This is the sum of the lengths of the Key Feedback Area and Undefined-Length Records Area.
8. **Number of Sensitive Segments.** This field is not used by GSAM, but it should be coded as part of the GSAM DB PCB mask. This field is 4 bytes.
9. **Key Feedback Area.** After a successful retrieval call, GSAM places the address of the record that is returned to your program in this field. This is called a record search argument (RSA). You can use it later if you want to retrieve that record directly by including it as one of the parameters on a GU call. This field is 8 bytes for basic format data sets or 12 bytes for large format data sets.
10. **Undefined-Length Records Area.** If you use undefined-length records (RECFM=U), the length in binary of the record you are processing is passed between your program and GSAM in this field. This field is 4 bytes long. When you issue a GU or GN call, GSAM places the binary length of the retrieved record in this field. When you issue an ISRT call, put the binary length of the record you are inserting in this field before issuing the ISRT call.

Related concepts

[“AIBTDLI interface” on page 226](#)

Use AIBTDLI as the interface between your application program and IMS.

[“GSAM record formats” on page 287](#)

GSAM records are nonkeyed. For variable-length records you must include the record length as the first 2 bytes of the record. Undefined-length records, like fixed-length records, contain only data (and control characters, if needed).

Retrieving and inserting GSAM records

GSAM records can be retrieved sequentially or directly. You can also add GSAM records to a new data set or add new records to the end of an existing data set in the database.

To retrieve GSAM records sequentially, use the GN call. The only required parameters are the GSAM PCB and the I/O area for the segment. To process the whole database, issue the GN call until you get a GB status code in the GSAM PCB. This status code means that you have reached the end of the database. GSAM automatically closes the database when you reach the end of it. To add records to a new data set or to add new records to the end of an existing data set in the database, use the ISRT call. GSAM adds the records sequentially in the order in which you supply them.

You can retrieve records directly from a GSAM database by supplying a record search argument (RSA) to the GSAM database. An RSA is like a segment search argument (SSA), but it contains the exact address of the record that you want to retrieve. The specific contents and format of the RSA depend on the access method that GSAM is using. For BSAM tape data sets and VSAM data sets, the RSA contains the relative byte address (RBA). For BSAM disk data sets, the RSA contains the disk address and uses the relative track and record format.

You can change your application programs to accommodate for extra 4 bytes when retrieving a record for a large format data set by using the INIT call with an I/O area containing the character string of RSA12. The INIT RSA12 call is coded in a GSAM application program before any calls to the GSAM database are coded. When a GSAM application issues the INIT RSA12 call, it tells IMS that the program can accept a 12-byte RSA when retrieving a record for large format data sets. The INIT RSA12 call must be issued by any application that uses large format data sets. Failure to issue the INIT RSA12 call for large format data

sets might cause an unexpected result. In the absence of an INIT RSA12 call, IMS continues to pass back an 8-byte RSA when retrieving a record for a basic format data set.

The following table provides more details about the format of the RSA for basic format data sets:

Table 46. Format of the RSA for basic format data sets

Position	Address
Positions 1-4	• BSAM (DASD) relative track and record (TTRZ) for the block in the buffer. • BSAM RBA. • VSAM RBA.
Position 5	Relative data set of the concatenated data set. The first data set number is 1.
Position 6	Relative volume of the data set. The first volume of data set is 1.
Positions 7 and 8	The current displacement.

The following table provides more details about the format of the RSA for large format data sets:

Table 47. Format of the RSA for large format data sets

Position	Address
Positions 1-4	• BSAM (DASD) relative track and record (TTTR) for the block in the buffer. • BSAM RBA.
Position 5	Zone byte
Position 6	Relative data set of the concatenated data set. The first data set number is 1.
Position 7	Relative volume of the data set. The first volume of data set is 1.
Positions 8-10	Null bytes. Not used.
Positions 11-12	The current displacement.

Before you can supply an RSA in a GU call to a GSAM database, that RSA must have previously been returned to you as a result of a GN or ISRT call. For GSAM to return an RSA, the GN or ISRT call must be issued with a fourth parameter that points to an 8-byte (basic format data set) or 12-byte (large format data set) RSA save area in your program. Save this RSA until you want to retrieve that particular record.

To retrieve that particular record, issue a GU call for the record and specify the address of its RSA as a fourth parameter of the GU call. GSAM returns the record to the I/O area that you named as one of the call parameters.

Restriction: Retrieve records directly from a GSAM database on DASD only. When using buffered I/O, buffer definitions for the output PCB may affect performance.

Resetting the position in a GSAM Database

You can use the GU call to reset the position in the GSAM database.

You can reset the position to the start of the GSAM database or to a specific record in the GSAM database:

- To reset the position to the start of the GSAM database using basic format data sets, issue a GU call with an RSA that consists of a fullword with a binary value of 1, followed by a fullword with a binary value of 0.
- To reset the position to the start of the GSAM database using large format data sets, issue a GU call with an RSA that consists of a fullword with a binary value of 1, followed by two fullwords with a binary value of 0.
- To reset the position to a specific record in the GSAM database, issue a GU call with an RSA that contains the saved RSA value from a prior ISRT or GN call for that record.

Related reference

[“GSAM coding considerations” on page 289](#)

The calls your program uses to access GSAM databases are not the same as the DL/I calls. The system service calls that you use with GSAM are symbolic CHKP and XRST.

[INIT call \(Application Programming APIs\)](#)

Explicit open and close calls to GSAM

IMS opens the GSAM data set when the first call is made and closes the data set when the application program terminates. Therefore, the application program does not usually need to make explicit open or close calls to GSAM.

However, explicit OPEN and CLSE calls are useful if:

- the application program loads a GSAM data set, and then in the same step reads the data set using GSAM (for example, to sort the data set). The application program should issue the GSAM CLSE call after the load is complete.
- the GSAM data set is an output data set, and it is possible that when the program executes it does not make GSAM ISRT calls. A data set is not created. Subsequent attempts to read the nonexistent data set (using GSAM or not) will likely result in an error. To avoid this situation, explicitly open the data set. DL/I closes the data set when the step terminates. Closing the data set prevents the possibility of attempting to read an empty data set.
- the GSAM data set is an output data set, and data exists beyond the EOF address in the dataset control block (DSCB). The previous job/step may have ended abnormally before the DSCB could be updated. If the program is restarted, but does not make GSAM ISRT calls, the EOF will not be updated at job/step termination when DL/I closes the data set. This could strand any data that exists past the EOF address. To avoid this situation, explicitly open the data set so that the DSCB can be updated with the correct EOF address.

The explicit OPEN or CLSE call need not include an I/O area parameter. Depending on the processing option of the PCB, the data set is opened for input or output. You can specify that an output data set contain either ASA or machine control characters. Including an I/O area parameter in the call and specifying OUTA in the I/O area indicates ASA control characters. Specifying OUTM specifies machine control characters.

GSAM record formats

GSAM records are nonkeyed. For variable-length records you must include the record length as the first 2 bytes of the record. Undefined-length records, like fixed-length records, contain only data (and control characters, if needed).

If you use undefined-length records, record length is passed between your program and GSAM in the 4-byte field that follows the key feedback area of the GSAM DB PCB. It is called the undefined-length records area. When you issue an ISRT call, supply the length. When you issue a GN or GU call, GSAM places the length of the returned record in this field. The advantage of using undefined-length records is that you do not need to include the record length at the beginning of the record, and records do not need to be of fixed length. The length of any record must be less than or equal to the block size (BLKSIZE) and greater than 11 bytes (an z/OS convention).

If you are using VSAM, you can use blocked or unblocked fixed-length or variable-length records. If you are using BSAM, you can use blocked or unblocked fixed-length, variable-length, or undefined-length records. Whichever you use, be sure to specify this on the RECFM keyword in the DATASET statement of the GSAM DBD. You can override this in the RECFM statement of the DCB parameter in the JCL. You can also include carriage control characters in the JCL for all formats.

Related concepts

“PCB masks for GSAM databases” on page 283

For the most part, you process GSAM databases in the same way that you process other IMS databases. You use calls that are very similar to DL/I calls to communicate your requests. GSAM describes the results of those calls in a GSAM DB PCB.

“Origin of GSAM data set characteristics” on page 290

For an input data set, the record format (RECFM), logical record length (LRECL), and block size (BLKSIZE) are based on the input data set label.

GSAM I/O areas

If you provide an optional I/O area, it must contain one of these values.

- INP for an input data set
- OUT for an output data set
- OUTA for an output data set with ASA control characters
- OUTM for an output data set with machine control characters

For GN, ISRT, and GU calls, the format of the I/O area depends on whether the record is fixed-length, undefined-length (valid only for BSAM), or variable-length. For each kind of record, you have the option of using control characters.

The formats of an I/O area for fixed-length or undefined-length records are:

- With no control characters, the I/O area contains only data. The data begins in byte 0.
- With control characters, the control characters are in byte 0 and the data begins in byte 1.

If you are using undefined-length records, the record length is passed between your program and GSAM in the PCB field that follows the key feedback area. When you are issuing an ISRT call, supply the length. When you are issuing a GN or GU call, GSAM places the length of the returned record in this field. This length field is 4 bytes long.

The formats for variable-length records differ because variable-length records include a length field, which other records do not have. The length field is 2 bytes. Variable-length I/O areas, like fixed-length and undefined-length I/O areas, can have control characters.

- Without control characters, bytes 0 and 1 contain the 2-byte length field, and the data begins in byte 2.
- With control characters, bytes 0 and 1 still contain the length field, but byte 2 contains the control characters, and the data starts in byte 3.

GSAM status codes

Your program should test for status codes after each GSAM call, just as it does after each DL/I or system service call.

If, you find that you have an error and terminate your program after checking the status codes, be sure to note the PCB in error before you terminate. The GSAM PCB address is helpful in determining problems. When a program that uses GSAM terminates abnormally, GSAM issues PURGE and CLSE calls internally, which changes the PCB information.

Status codes that have specific meanings for GSAM are:

AF

GSAM detected a BSAM variable-length record with an invalid format. Terminate your program.

AH

You have not supplied an RSA for a GU call.

AI

There has been a data management OPEN error.

AJ

One of the parameters on the RSA that you supplied is invalid.

AM

You have issued an invalid request against a GSAM database.

AO

An I/O error occurred when the data set was accessed or closed.

GB

You reached the end of the database, and GSAM has closed the database. The next position is the beginning of the database.

IX

You issued an ISRT call after receiving an AI or AO status code. Terminate your program.

Symbolic CHKP and XREST with GSAM

To checkpoint GSAM databases, use symbolic CHKP and XREST calls.

By using GSAM to read or write the data set, symbolic CHKP and XREST calls can be used to reposition the data set at the time of restart, enabling you to make your program restartable. When you use an XREST call, IMS repositions GSAM databases for processing. CHKP and XREST calls are available to application programs that can run as batch programs, batch-oriented BMPs, or transaction-oriented BMPs.

Restriction: When restarting GSAM databases:

- You cannot use temporary data sets with a symbolic CHKP or XREST call.
- A SYSOUT data set at restart time may give duplicate output data.
- You cannot restart a program that is loading a GSAM or VSAM database.
- The GSAM database data set must have the same data set format (BASIC or LARGE) as when the symbolic CHKP call was issued.

When IMS restores the data areas specified in the XREST call, it also repositions any GSAM databases that your program was using when it issued the symbolic CHKP call. If your program was loading GSAM databases when the symbolic CHKP call was issued, IMS repositions them (if they are accessed by BSAM). If you make a copy of the GSAM data set for use as input to the restart process, ensure that the short blocks are written to the new data set as short blocks, for example, using IEBGENER with RECFM=U for SYSUT1. You can also do the restart using the original GSAM data set.

During GSAM XREST processing, a check is made to determine if the GSAM output data set to be repositioned is empty, and if the abending job had previously inserted records into the data set.

GSAM coding considerations

The calls your program uses to access GSAM databases are not the same as the DL/I calls. The system service calls that you use with GSAM are symbolic CHKP and XREST.

The following table summarizes GSAM database calls. The five calls you can use to process GSAM databases are:

- CLSE
- GN
- GU
- ISRT
- OPEN

The COBOL, PL/I, Pascal, C, and assembler language call formats and parameters for these calls are the same and are described in the following table. GSAM calls do not differ significantly from DL/I calls, but GSAM calls must reference the GSAM PCB, and they do not use SSAs.

Java application programs running in Java batch processing (JBP) regions can access GSAM databases by using the IMS Java dependent region resource adapter.

Table 48. Summary of GSAM calls

Call Formats	Meaning	Use	Options	Parameters
CLSE	Close	Explicitly closes GSAM database	None	function, gsam pcb
GNbb	Get Next	Retrieves next sequential record	Can supply address for RSA to be returned	function, gsam pcb, i/o area [,rsa name]
GUBb	Get Unique	Establishes position in database or retrieves a unique record	None	function, gsam pcb, i/o area, rsa name
ISRT	Insert	Adds new record at end of database	Can supply address for RSA to be returned	function, gsam pcb, i/o area [,rsa name]
OPEN	Open	Explicitly opens GSAM database	Can specify printer or punch control characters	function, gsam pcb [, open option]

Related reference

“Accessing GSAM data from a JBP application” on page 686

GSAM data are frequently referred to as z/OS data sets or as flat files. This kind of data is non-hierarchical in structure. You can access data from GSAM databases from a JBP application.

Origin of GSAM data set characteristics

For an input data set, the record format (RECFM), logical record length (LRECL), and block size (BLKSIZE) are based on the input data set label.

If this information is not provided by a data set label, the DD statement or the DBD specifications are used. The DD statement has priority.

An output data set can have the following characteristics:

- Record format
- Logical record length
- Block size
- Other JCL DCB parameters
- DNS type

Specify the record format on the DATASET statement of the GSAM DBD. The options are:

- V for variable
- VB for variable blocked
- F for fixed
- FB for fixed blocked
- U for undefined

The V, F, or U definition applies and is not overridden by the DCB=RECFM= specification on the DD statement. However, if the DD RECFM indicates blocked and the DBD does not, RECFM is set to blocked. If the DD RECFM of A or M control character is specified, it applies as well.

Unless an undefined record format is used, specify the logical record using the RECORD= parameter of the DATASET statement of DBDGEN, or use DCB=LRECL=xxx on the DD statement. If the logical record is specified on both, the DD statement has priority. Refer to the following table for the maximum record length

<i>Table 49. BSAM and VSAM logical record lengths for GSAM data sets by record format</i>		
Record Format	BSAM logical record length	VSAM logical record length
Fixed/Fixed Block	32760 bytes	32760 bytes
Variable/Variable Blocked	32756 bytes	32756 bytes
Undefined	32760 bytes	not supported

Specify block size using the BLOCK= or SIZE= parameter of the DATASET statement of DBDGEN, or use DCB=BLKSIZE=xxx on the DD statement. If block size is specified on both, the DD statement has priority. If the block size is not specified by the DBD or the DD statement, the system determines the size based on the device type, unless the undefined record format is used.

The other JCL DCB parameters that can be used, include:

- CODE
- DEN
- DNSTYPE
- TRTCH
- MODE
- STACK
- PRTSP, which can be used if RECFM does not include A or M
- DCB=BUFNO=X, which, when used, causes GSAM to use X number of buffers

Restriction: Do not use BFALN, BUFL, BUFOFF, FUNC, NCP, and KEYLEN.

Related concepts

[“GSAM record formats” on page 287](#)

GSAM records are nonkeyed. For variable-length records you must include the record length as the first 2 bytes of the record. Undefined-length records, like fixed-length records, contain only data (and control characters, if needed).

DD statement DISP parameter for GSAM data sets

The DD statement DISP parameter varies, depending on whether you are creating input or output data sets and how you plan to use the data sets.



Attention: Specifying the DISP=OLD or DISP=SHR parameter for a normal start with non-empty data sets will overwrite the existing records from the beginning of the data set.

- For input data sets, use the DISP=OLD parameter.
- For output data sets, consider the following options:
 - To create an output data set allocated by the DD statement, set DISP=NEW.
 - To add new records to an empty data set when performing normal start or a restart after failure, set DISP=MOD, DISP=SHR, or DISP=OLD.
 - When restarting the step, set DISP=OLD for existing data sets and DISP=MOD for empty data sets.

- To add new records to an existing non-empty data set when performing a restart after failure, set DISP=MOD, DISP=SHR, or DISP=OLD. These parameters add new records from the restart point on the existing data set.
- To add new records to the end of an existing non-empty data set when performing normal start, set DISP=MOD.

Extended checkpoint restart for GSAM data sets

If you are using extended checkpoint restart for GSAM data sets, these recommendations may apply.

- Do not use passed data sets.
- Do not use backward references to data sets in previous steps.
- Do not use DISP=MOD to add records to an existing tape data set.
- Do not use DISP=DELETE or DISP=UNCATLG.
- Use DFSMS striped data sets under the following conditions:
 - When the data sets is managed by SMS.
 - When the data sets are likely to exceed the system extent limit for volumes.
- Additionally, keep in mind that:
 - No attempt is made to reposition a SYSIN, SYSOUT, or temporary data set.
 - No attempt is made to reposition any of the concatenated data sets for a concatenated DD statement if any of the data sets are a SYSIN or SYSOUT.
 - If you are using concatenated data sets, specify the same number and sequence of data sets at restart time and checkpoint time.
 - GSAM/VSAM load mode restrictions apply to both non-striped and striped data sets.
 - If the PSB contains an open GSAM VSAM output data set when the symbolic checkpoint call is issued, the system returns an AM status code in the database PCB as a warning. This means that the data set is not repositioned at restart and the checkpoint has completed normally.
 - If an ISRT call is issued after a CLSE call and the GSAM data set is defined as DISP=OLD, all CHKP calls made prior to the CLSE call will contain invalid reposition information. Ensure a CHKP call is issued after a CLSE all when using DISP=OLD to avoid an abend U0271 after an extended restart (XRST).

Copying GSAM data sets between checkpoint and restart

To position GSAM data sets when restarting non-striped GSAM DASD data sets, use the relative track and record format (TTRZ or TTTRZ for large format data sets).

GSAM uses the TTRZ or TTTRZ on the volume to position non-striped GSAM DASD data sets when restarting. For a tape data set, the relative record on the volume is used. The relative record on the tape volume cannot be changed.

To copy non-striped DASD data sets between checkpoint and restart:

- Copy the data set to the same device type.
- Avoid any reblocking by using the undefined record format (RECFM=U) for both the input and the output data set.

Each copied volume contains the same number of records as the original volumes.

Note: GSAM uses the relative block number (RBN) to reposition striped DASD data sets. When data sets that are managed by SMS are used with GSAM databases, you cannot control how each volume is copied. After the data set is copied, unlike with non-striped DASD data sets, you do not need to ensure that the TTRZ or the TTTRZ of the restart record is unchanged.

Converting data sets from non-striped data sets to striped data sets

Convert GSAM/BSAM non-striped data sets to striped data sets before you must perform an extended restart when a system allocation limit is exceeded or a system X'37' error condition occurs. Non-striped data sets that are not managed by SMS extend beyond their initial primary or secondary allocation only by volume, but with non-striped GSAM/BSAM multiple volume data sets that are managed by SMS, the resulting new space allocation takes effect for all of the volumes in the data set.

If you copy non-striped data sets that are managed by SMS after you change the space allocation values, the number of records in the new volumes will be different from the number of records in the old volume. The new primary and secondary allocation values are used with non-striped data sets. As the data is copied, all of the space that is allocated on the new volume is used before the data is copied to the next volume.

If an error condition (system X'37' or system allocation limit exceeded) occurs during the processing of a GSAM/BSAM non-striped data set, and the data set is converted to a striped data set after the error occurs, a restart after failure will not complete successfully. Because the issued checkpoint saved a TTRZ or a TTTRZ value in the log record for repositioning, the log record for striped data sets will be used by GSAM restart after failure, which requires a relative block number (RBN) to perform the repositioning.

Concatenated data sets used by GSAM

GSAM can use concatenated data sets, which may be on unlike device types, such as DASD and tape, or on different DASD devices. Logical record lengths and block sizes can differ, and it is not required that the data set with the largest block size be concatenated first.

The maximum number of concatenated data sets for a single DD statement is 255. The number of buffers determined for the first of the concatenated data sets is used for all succeeding data sets. Generation data groups can result in concatenated data sets.

Specifying GSAM data set attributes

When specifying GSAM data set attributes, the following settings are recommended.

- On the DBD, specify RECFM. (It is required.)
- On the DATASET statement, specify the logical record length using RECORD=. If the data set can become larger than 65535 tracks on a DASD volume and you want the data set to not span multiple volumes, specify the DSNTYPE=LARGE parameter.
- On the DD statement, do not specify LRECL, RECFM, or BLKSIZE. The system determines block size, with the exception of RECFM=U. The system determines logical record length from the DBD.
- For the PSB, specify PROCOPT=LS for output and GS for input. If you include S, GSAM uses multiple buffers instead of a single buffer for improved performance.

IMS will add 2 bytes to the record length value specified in the DBD in order to accommodate the ZZ field that is needed to make up the BSAM RDW. Whenever the database is GSAM or BSAM and the records are variable (V or VB), IMS will add 2 bytes to the record length value in the GSAM records passed by the application. Such addition allows IMS to accommodate the ZZ field that makes up the BSAM RDW (Record Descriptor Word).

Example of GSAM or BSAM where the records are variable

```
//IDASD DD DUMMY
//ODASD DD UNIT=SYSDA,VOL=SER=000000,DISP=(,KEEP),
// SPACE=(TRK,(5,1)),DSN=GSAM.VARIABLE1,
// DCB=(RECFM=VB,BLKSIZE=32760,LRECL=32756)
//SYSIN DD *,DCB=BLKSIZE=80
S 1 1 1 1 1 DBDNAME
L ISRT
L V8187 DATA 1ST RECORD LOADED TO GSAM <---RDW
L ISRT
L V8187 DATA 2ND RECORD LOADED TO GSAM
L ISRT
L V8187 DATA 3RD RECORD LOADED TO GSAM
```

```
L ISRT
L V8187 DATA 4TH RECORD LOADED TO GSAM
```

In the above example, four GSAM records (IMS segment) can be contained in one 32756 byte (MVS) record.

DLI, DBB, and BMP region types and GSAM

To access GSAM databases, IMS builds its DLI control blocks using PSB and DBD information from PSBLIB, DBDLIB and ACBLIB. The source of the PSB and DBD information depends on the region type.

For DLI offline batch regions, IMS obtains PSB and DBD information from PSBLIB and DBDLIB. For DBB offline batch regions, IMS database management obtains PSB and DBD information from ACBLIB. For online batch regions (BMPs), IMS builds its DLI control blocks with information from ACBLIB. If an application is scheduled in a BMP region and the PSB associated with the application contains one or more GSAM PCBs, IMS scheduling obtains PSB information from ACBLIB and PSBLIB. In this case, the PSB in ACBLIB and PSBLIB must be the same. GSAM database management does not obtain PSB and DBD information from ACBLIB. Instead, GSAM database management obtains PSB and DBD information from PSBLIB and DBDLIB.

When you initialize a DLI, DBB or BMP region using GSAM, you must include an //IMS DD and GSAM DD statements. When DBB or BMP regions are not using GSAM, //IMS DD statements do not need to be included. To load PSBs and DBDs and build GSAM control blocks, you must include an //IMS DD statement. In the following figure, an example of the //IMS DD statement with data sets that are larger than 65535 tracks is shown.

Figure 64. //IMS DD statement example

```
//STEP      EXEC      PGM=DFSRR00,PARM=[BMP|DBB|DLI],...
//STEPLIB   DD        DSN=executionlibrary-name,DISP=SHR
//          DD        DSN=pgmlib-name,DISP=SHR
//IMS       DD        DSN=psblib-name,DISP=SHR
//          DD        DSN=dbdlib-name,DISP=SHR
//IMSACB    DD        DSN=acblib-name,disp=shr (required for DBB)
//SYSPRINT  DD        SYSOUT=A
//SYSUDUMP  DD        SYSOUT=A
//ddnameex  DD        (add DD statements for required GSAM databases)
//ddnameex  DD        (add DD statements for non-GSAM IMS databases
//ddnameex  DD        for DLI/DBB)
//ddnameex  DD        DSNTYPE=LARGE,...
.
.
.
/*
```

Chapter 21. Processing Fast Path databases

You can write application programs to access Fast Path databases, including main storage databases and data entry databases.

The two kinds of Fast Path databases are:

- Main storage databases (MSDBs), which are available in a DB/DC environment, and contain only root segments in which you store data that you access most frequently.
- Data entry databases (DEDBs) are hierarchic databases that can have as many as 15 hierarchic levels and as many as 127 segment types. DEDBs are available to both IMS users and CICS users with DBCTL.

Restriction: This DEDB information applies to CICS users with DBCTL. CICS users can access MSDBs in DBCTL in read mode, but update mode is not supported.

VSO considerations

VSO is transparent to the processing of an application. Where the data resides is immaterial to the application.

Data locking for MSDBs and DEDBs

All MSDB calls, including the FLD call, can lock the data at the segment level. The lock is acquired at the time the call is processed and is released at the end of the call. All DEDB calls, with the exception of HSSP calls, are locked at the VSAM CI level. For single-segment, root-only, fixed-length VSO areas, if you specify PROCOPT R or G, the application program can obtain segment-level locks for all calls. If you specify any other PROCOPT, the application program obtains VSAM CI locks.

Segment-level locking (SLL) provides a two-tier locking scheme. First, a share (SHR) lock is obtained for the entire CI. Then, an exclusive (EXCL) segment lock is obtained for the requested segment. This scheme allows for contention detection between SLL users of the CI and EXCL requestors of the CI. When contention occurs between an existing EXCL CI lock user and a SHR CI lock requestor, the SHR CI lock is upgraded to an EXCL CI lock. During the time that this EXCL CI lock is held, subsequent SHR CI lock requests must wait until the EXCL CI is released at the next commit point.

DEDB FLD calls are not locked at call time. Instead, the lock is acquired at a commit point.

During sync-point processing, the lock is re-acquired (if not already held), and the changes are verified. Verification failure results in the message being reprocessed (for message-driven applications) or an FE status code (for non-message-driven applications). Verification can fail if the segment used by the FLD call has been deleted or replaced before a sync-point.

Segment retrieval for a FLD call is the same as for a GU call. An unqualified FLD call returns the first segment in the current area, just as an unqualified GU call does. After the FLD call is processed, all locks for the current CI are released if the current CI is unmodified by any previous call.

When a compression routine is defined on the root segment of a DEDB with a root-only structure, and when that root segment is a fixed-length segment, its length becomes variable after being compressed. To replace a compressed segment, you must perform a delete and an insert. In this case, segment level control and locking will not be available.

Related concepts

[Data entry databases \(Database Administration\)](#)

[Main storage databases \(MSDBs\) \(Database Administration\)](#)

[High-speed sequential processing \(HSSP\) \(Database Administration\)](#)

Fast Path database calls

Use Fast Path database calls in your application programs to access Fast Path databases.

The following table summarizes the database calls you can use with Fast Path databases.

Table 50. Summary of Fast Path database calls

Function Code	Types of MSDBs:			DEDBs
	Nonterminal-Related	Terminal-Related Fixed	Terminal-Related Dynamic	
DEQ				X
FLD	X	X	X	X
GU, GHU	X	X	X	X
GN, GHN	X	X	X	X
GNP, GHNP DLET			X	X
ISRT			X	X
POS				X
REPL	X	X	X	X
RLSE				X

DL/I calls to DEDBs can include the same number of SSAs as existing levels in the hierarchy (a maximum of 15). They can also include command codes and multiple qualification statements.

Restriction:

- Fast Path ignores command codes that are used with sequential dependent segments.
- If you use a command code that does not apply to the call you are using, Fast Path ignores the command code.
- If you use F or L in an SSA for a level greater than the established parent, Fast Path ignores the F or L command code.
- DL/I calls to DEDBs cannot include the independent AND, which is used only with secondary indexing.

Calls to DEDBs can use all command codes. Only calls to DEDBs that use subset pointers can use the R, S, Z, W, and M command codes. The following table shows which calls you can use with these command codes.

Table 51. Subset pointer command codes and calls

Command Code	DLET	GU GHU	GN GHN	GNP GHNP	ISRT	REPL
M		X	X	X	X	X
R		X	X	X	X	
S		X	X	X	X	X
W		X	X	X	X	X
X	X	X	X	X	X	X

Main storage databases (MSDBs)

MSDBs contain only root segments. Each segment is like a database record, because the segment contains all of the information about a particular subject.

In a DL/I hierarchy, a database record is made up of a root segment and all its dependents. For example, in the medical hierarchy, a particular PATIENT segment and all the segments underneath that PATIENT segment comprise the database record for that patient. In an MSDB, the segment is the whole database record. The database record contains only the fields that the segment contains. MSDB segments are fixed length.

The two kinds of MSDBs are terminal related and non-terminal related. In terminal-related MSDBs, each segment is owned by one logical terminal. The segment that is owned can be updated only by that terminal. Related MSDBs can be fixed or dynamic. You can add segments to and delete segments from dynamic related MSDBs. You cannot add segments to or delete segments from fixed related MSDBs.

In the second kind of MSDB, called non-terminal related (or nonrelated) MSDBs, the segments are not owned by logical terminals.

Restrictions on using calls for MSDBs

To retrieve segments from an MSDB, you can issue Get calls just as you do to retrieve segments from other IMS databases. Because MSDBs contain only root segments, you only use GU and GN calls (and GHU and GHN calls when you plan to update a segment). If the segment name field in the SSA contains *MYLTERM, the GU, GHU, and FLD calls return the LTERM-owned segment, and the remainder of the SSA is ignored.

When you are processing MSDBs, you should consider the following differences between calls to MSDBs and to other IMS databases:

- You can use only one SSA in a call to an MSDB.
- MSDB calls cannot use command codes.
- MSDB calls cannot use multiple qualification statements (Boolean operators).
- The maximum length for an MSDB segment key is 240 bytes (not 255 bytes, as in other IMS databases).
- If the SSA names an arithmetic field (types P, H, or F) as specified in the database description (DBD), the database search is performed using arithmetic comparisons (rather than the logical comparisons that are used for DL/I calls).
- If a hexadecimal field is specified, each byte in the database field is represented in the SSA by its two-character hexadecimal representation. This representation makes the search argument twice as long as the database field.

Characters in hexadecimal-type SSA qualification statements are tested for validity before translation to the database format. Only numerals 0 through 9 and letters A through F are accepted.

- Terminal-related and non-terminal-related LTERM-keyed MSDBs are not supported for ETO or LU 6.2 terminals. Attempted access results in no data being retrieved and an AM status code. See *IMS Version 15.3 Communications and Connections* for more information on ETO and LU 6.2.
- MSDBs cannot be shared among IMS subsystems in a sysplex group. When using the Fast Path Expedited Message Handler (EMH), terminal related and non-terminal related with terminal key MSDBs can only be accessed by static terminals. These static terminals run transactions with Sysplex Processing Code (SPC) of Locals Only as specified in DBFHAGU0 (Input Edit Router exit routine).

The restrictions above do not apply to CICS users.

Data entry databases (DEDBs)

A DEDB contains a root segment and as many as 127 dependent segment types. One of these can be a sequential dependent; the other 126 are direct dependents. Sequential dependent segments are stored in chronological order. Direct dependent segments are stored hierarchically.

Restriction: IMS does not support inserting SDEP segments into the same DEDB area from multiple distributed IMS PSTs or UORs, as a global UOR, instanced from a non-IMS resource manager.

This restriction applies to ODBA users, ODBM users, CCTL DRA users such as CICS, and Db2 stored procedures that access IMS using ODBA.

DEDBs can provide high data availability. Each DEDB can be partitioned, or divided into multiple areas. Each area contains a different collection of database records. In addition, you can make as many as seven copies of each area data set. If an error exists in one copy of an area, application programs continue to access the data by using another copy of that area. Use of the copy of an area is transparent to the application program. When an error occurs to data in a DEDB, IMS does not stop the database. IMS makes the data in error unavailable but continues to schedule and process application programs. Programs that do not need the data in error are unaffected.

DEDBs can be shared among application programs in separate IMS systems. Sharing DEDBs is virtually the same as sharing full-function databases, and most of the same rules apply. IMS systems can share DEDBs at the area level (instead of at the database level as with full-function databases), or at the block level.

Restriction: IMS does not support inserting SDEP segments into the same DEDB area from multiple distributed IMS PST or UORs as a global UR, instanced from a non-IMS resource manager.

This restriction applies to users of:

- ODBA
- ODBM
- CCTL DR, such as CICS
- Db2 stored procedures that access IMS using ODBA

Related reading: For more information on DEDB data sharing, see the explanation of administering IMS systems that share data in *IMS Version 15.3 System Administration*.

Updating segments: REPL, DLET, ISRT, and FLD

You can use REPL, DLET, ISRT, and Field (FLD) calls to update a Data entry database (DEDB) or a Main storage database (MSDB). A single FLD call does what a Get Hold call and a REPL call do in two calls. FLD calls are unique to DEDBs and MSDBs.

You can issue REPL, DLET, and ISRT calls to non-terminal-related MSDBs (without terminal-related keys) or DEDBs. You can issue REPL calls to related or nonrelated MSDBs. Before issuing these calls, you must issue a Get Hold call for the segment that you want to update.

Alternatively, you can use a **FLD** call to access and change field contents within a segment.

The FLD call has two types. **FLD/VERIFY** compares a target segment field value to the value that you supply in a field search argument (FSA). **FLD/CHANGE** changes the target segment field value to your specifications in the FSA. A FLD/CHANGE call succeeds only if the previous FLD/VERIFY call succeeds.

FLD calls can simplify some operations. For example, consider how to address the following scenario, which is based on the ACCOUNT segment that is described in [Database hierarchy examples \(Application Programming\)](#). A bank must perform the following process to determine whether a customer can withdraw a requested amount from their account:

1. Retrieve the customer's account segment.
2. Verify that the customer's account balance is greater than the requested withdrawal amount.

3. Update the balance to reflect the withdrawal if the starting balance is greater than the requested withdrawal amount.

If you do not use FLD calls, you must write a program to issue a GU call to retrieve the segment, verify its contents with program logic, and then issue a REPL call to update the balance to reflect the withdrawal.

If you issue FLD calls with a root segment search argument (SSA), you can retrieve, compare, and change the customer segment. Use a FLD/VERIFY call to compare the BALANCE field with the requested withdrawal amount. If the comparison is satisfactory, use a FLD/CHANGE call to update the BALANCE field.

The FLD call SSA has the same format as SSAs for other calls. If no SSA exists, the first segment in the DEDB or MSDb is retrieved.

The segment retrieved by a FLD call is the same as the segment that can be retrieved by a GHU call. After the FLD call, the position is lost. An unqualified GN call after a FLD call returns the next segment in the current area.

Note: A FLD call cannot change the length of a compressed segment, even if the segment is variable length. To make changes, you must issue a REPL call.

Sync point processing (SYNC call) will fail if a FLD call that changes segment length is issued to a DEDB segment type that has a compression routine specified. If the FLD call is issued from a BMP region, the SYNC call fails with status code FV. If the FLD call is issued from an MPP or IFP region, the SYNC call will ABEND U799 with RC 3.

Restriction: Do not concurrently delete or replace a call with length change in the following circumstances:

- FLD call to any DEDB AREA
- GU call with VIEW=MSDB or VIEW=MSDBL to Root Only DEDB AREA

The application will ABEND U1026 or some segments may be skipped on a GN call.

Related reference

[FLD call \(Application Programming APIs\)](#)

Checking the contents of a field: FLD/VERIFY

A FLD/VERIFY call compares the contents of a specified field in a segment to the value that you supply. The way that a FLD/VERIFY call compares the two depends on the operator you supply.

When you supply the name of a field and a value for comparison, you can determine if the value in the field is:

- Equal to the value you have supplied
- Greater than the value you have supplied
- Greater than or equal to the value you have supplied
- Less than the value you have supplied
- Less than or equal to the value you have supplied
- Not equal to the value you have supplied

After IMS performs the comparison that you have asked for, it returns a status code (in addition to the status code in the PCB) to tell you the results of the comparison.

If failures occur during processing of the FLD call, IMS reprocesses the call only if the failures are caused by deadlocks or verification errors. For other types of failures, abend U0819 is issued.

You specify the name of the field and the value that you want its value compared to in a field search argument, or FSA. The FSA is also where IMS returns the status code. You place the FSA in an I/O area before you issue a FLD call, and then you reference that I/O area in the call—just as you do for an SSA in a DL/I call. An FSA is similar to an SSA in that you use it to give information to IMS about the information

you want to retrieve from the database. An FSA, however, contains more information than an SSA. The table below shows the structure and format of an FSA.

Table 52. FSA structure

FSA Component	Field Length
FLD NAME	8
SC	1
OP	1
FLD VALUE	Variable
CON	1

The five fields in an FSA are:

Field Name (FLD Name)

This is the name of the field that you want to update. The field must be defined in the DBD.

Status Code (SC)

This is where IMS returns the status code for this FSA. If IMS successfully processes the FSA, it returns a blank status code. If IMS fails to process the FSA, it returns a FE status code to the PCB to indicate a nonblank status code in the FSA and returns a nonblank FSA status code. The FSA status codes that IMS might return to you on a FLD/VERIFY call are:

B

The length of the data supplied in the field value is invalid, or the segment length of the data in the database is too small to contain the field length specified in the DBD.

D

The verify check is unsuccessful. In other words, the answer to your query is no.

E

The field value contains invalid data. The data you supplied in this field is not the same type of data that is defined for this field in the DBD.

H

The requested field is not found in the segment.

Operator (OP)

This tells IMS how you want the two values compared. For a FLD/VERIFY call, you can specify:

E

Verify that the value in the field is equal to the value you have supplied in the FSA.

G

Verify that the value in the field is greater than the value you have supplied in the FSA.

H

Verify that the value in the field is greater than or equal to the value you have supplied in the FSA.

L

Verify that the value in the field is less than the value you have supplied in the FSA.

M

Verify that the value in the field is less than or equal to the value you have supplied in the FSA.

N

Verify that the value in the field is not equal to the value you have supplied in the FSA.

Field Value (FLD Value)

This area contains the value that you want IMS to compare to the value in the segment field. The data that you supply in this area must be the same type of data in the field you have named in the first field of the FSA. The five types of data are: hexadecimal, packed decimal, alphanumeric (or a combination of data types), binary fullword, and binary halfword. The length of the data in this area must be the same as the length that is defined for this field in the DBD.

Exceptions:

- If you are processing hexadecimal data, the data in the FSA must be in hexadecimal. This means that the length of the data in the FSA is twice the length of the data in the field in the database. IMS checks the characters in hexadecimal fields for validity before that data is translated to database format. (Only 0 to 9 and A to F are valid characters.)
- For packed-decimal data, you do not need to supply the leading zeros in the field value. This means that the number of digits in the FSA might be less than the number of digits in the corresponding database field. The data that you supply in this field must be in a valid packed-decimal format and must end in a sign digit.

When IMS processes the FSA, it does logical comparisons for alphanumeric and hexadecimal fields; it does arithmetic comparisons for packed decimal and binary fields.

Connector (CON)

If this is the only or last FSA in this call, this area contains a blank. If another FSA follows this one, this area contains an asterisk (*). You can include several FSAs in one FLD call, if all the fields that the FSAs reference are in the same segment. If you get an error status code for a FLD call, check the status codes for each of the FSAs in the FLD call to determine where the error is.

When you have verified the contents of a field in the database, you can change the contents of that field in the same call. To do this, supply an FSA that specifies a change operation for that field.

Changing the contents of a field: FLD/CHANGE

To indicate to IMS that you want to change the contents of a particular field, use an FSA, just as you do in a FLD/VERIFY call.

The difference is in the operators that you can specify and the FSA status codes that IMS can return to you after the call. To use FLD/CHANGE:

- You specify the name of the field that you want to change in the first field of the FSA (Field Name).
- You specify an operator in the third field of the FSA (Operator), which indicates to IMS how you want to change that field.
- You specify the value that IMS must use to change the field in the last area of the FSA (Field Value).

By specifying different operators in a FLD/CHANGE call, you change the field in the database in these ways:

- Add the value supplied in the FSA to the value in the field.
- Subtract the value supplied in the FSA from the value in the field.
- Set the value in the database field to the value supplied in the FSA.

You code these operators in the FSA with these symbols:

- To add: +
- To subtract: -
- To set the field equal to the new value: =

You can add and subtract values only when the field in the database contains arithmetic (packed-decimal, binary-fullword, or binary-halfword) data.

The status codes you can receive in a FLD/CHANGE FSA are:

A

Invalid operation; for example, you specified the + operator for a field that contains character data.

B

Invalid data length. The data you supplied in the FSA is not the length that is defined for that field in the DBD.

C

You attempted to change the key field in the segment. Changing the key field is not allowed.

E

Invalid data in the FSA. The data that you supplied in the FSA is not the type of data that is defined for this field in the DBD.

F

You tried to change an unowned segment. This status code applies only to related MSDBs.

G

An arithmetic overflow occurred when you changed the data field.

H

The requested field was not found in the segment.

If failures occur during processing of the FLD call, IMS reprocesses the call only if the failures are caused by deadlocks or verification errors. For other types of failures, abend U0819 is issued.

Example of using FLD/VERIFY and FLD/CHANGE

Using the bank account segment from the "Bank Account Example" database, assume that a customer wants to withdraw \$100 from a checking account. The checking account number is 24056772. To find out whether the customer can withdraw this amount, you must check the current balance. If the current balance is greater than \$100, you want to subtract \$100 from the balance, and add 1 to the transaction count in the segment.

You can do all of this processing by using one FLD call and three FSAs. The three FSAs are described:

1. Verify that the value in the BALANCE field is greater than or equal to \$100. For this verification, you specify the BALANCE field, the H operator for greater than or equal to, and the amount. The amount is specified without a decimal point. Field names less than eight characters long must be padded with trailing blanks to equal eight characters. You also have to leave a blank between the field name and the operator for the FSA status code. This FSA looks like this:

```
BALANCEbb
H10000*
```

The last character in the FSA is an asterisk, because this FSA will be followed by other FSAs.

2. Subtract \$100 from the value in the BALANCE field if the first FSA is successful. If the first FSA is unsuccessful, IMS does not continue processing. To subtract the amount of the withdrawal from the amount of the balance, you use this FSA:

```
BALANCEbb
-10000*
```

Again, the last character in the FSA is an asterisk, because this FSA is followed by a third FSA.

3. Add 1 to the transaction count for the account. To do this, use this FSA:

```
TRANCNTbb
001b
```

In this FSA, the last character is a blank (b), because this is the last FSA for this call.

When you issue the FLD call, you do not reference each FSA individually; you reference the I/O area that contains all of them.

Commit-point processing in MSDBs and DEDBs

Your existing application programs can use either the MSDB commit view or the default DEDB commit view.

MSDB commit view

When you update a segment in an MSDB, IMS does not apply your updates immediately. Updates do not go into effect until your program reaches a commit point.

As a result of the way updates are handled, you can receive different results if you issue the same call sequence against a full-function database or a DEDB and an MSDB. For example, if you issue GHU and REPL calls for a segment in an MSDB, and then issue another Get call for the same segment in the same commit interval, the segment that IMS returns to you is the "old" value, not the updated one. If, however, you issue the same call sequence for a segment in a full-function database or DEDB, the second Get call returns the updated segment.

When the program reaches a commit point, IMS also reprocesses the FLD VERIFY/CHANGE call. If the VERIFY test passes, the change is applied to the database. If the VERIFY test fails, the changes made since the previous commit point are undone, and the transaction is reprocessed.

DEDBs with MSDB commit view

To use the MSDB commit view for DEDBs, specify VIEW=MSDB on the PCB statement; if you do not specify VIEW=MSDB, the DEDB uses the default DEDB commit view. So no changes to any existing application programs are required in order to migrate your MSDBs to DEDBs.

Assume that you specify VIEW=MSDB in the PCB and an application program issues GHU and REPL calls to a DEDB followed by another GHU call for the segment in the same commit interval. Then the application program receives the old value of the data and not the new value from the REPL call. If you do not specify VIEW=MSDB, your application program receives the new updated values of the data, just as you expect for a DEDB or other DL/I database.

You can specify VIEW=MSDB for any DEDB PCB. If it is specified for a non-DEDB database, you receive message DFS0904 during ACBGEN.

If you issue a REPL call with a PCB that specifies VIEW=MSDB, the segment must have a key. This requirement applies to any segment in a path if command code 'D' is specified. Otherwise, the AM status code is returned. See *IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes* for information about that status code.

The following code shows an example of a PCB that specifies the VIEW option.

Sample PCB specifying View=MSDB

PCB	,	*00000100
	TYPE=DB,	*00000200
	NAME=DEDBJN21,	*00000300
	PROCOPT=A,	*00000400
	KEYLEN=30,	*00000500
	VIEW=MSDB,	*00000600
	POS=M	00000700

Related reference

[“Issuing checkpoints” on page 263](#)

Two kinds of checkpoint (CHKP) calls exist: the basic CHKP and the symbolic CHKP. All IMS programs and CICS shared database programs can issue the basic CHKP call; only BMPs and batch programs can use either call.

[“Commit-point processing in a DEDB” on page 319](#)

IMS retains database updates in processor storage until the program reaches a commit point. IMS saves updates to a DEDB in Fast Path buffers. The database updates are not applied to the DEDB until after the program has successfully completed commit-point processing.

Processing DEDBs (IMS and CICS with DBCTL)

You can use subset pointers, secondary indexes, the POS call, data locking, and the P and H processing options in your application program to process DEDBs.

Processing Fast Path DEDBs with subset pointer command codes

Subset pointers and the command codes you use with them are optimization tools that significantly improve the efficiency of your program when you need to process long segment chains.

Subset pointers are a means of dividing a chain of segment occurrences under the same parent into two or more groups or subsets. You can define as many as eight subset pointers for any segment type. You then define the subset pointers from within an application program. Each subset pointer points to the start of a new subset. For example, in the following topic, suppose you define one subset pointer that divides the last three segment occurrences from the first four. Your program can then refer to that subset pointer through command codes and directly retrieve the last three segment occurrences.




Figure 65. Processing a long chain of segment occurrences with subset pointers

You can use subset pointers at any level of the database hierarchy, except at the root level. If you try to use subset pointers at the root level, they are ignored.

The following figures show some of the ways you can set subset pointers. Subset pointers are independent of one another, which means that you can set one or more pointers to any segment in the chain. For example, you can set more than one subset pointer to a segment, as shown in the following figure.




Figure 66. Examples of setting subset pointers

You can also define a one-to-one relationship between the pointers and the segments, as shown in the following figure.




Figure 67. Additional examples of setting subset pointers

The following figure shows how the use of subset pointers divides a chain of segment occurrences under the same parent into subsets. Each subset ends with the last segment in the entire chain. For example, the last segment in the subset that is defined by subset pointer 1 is B7.




Figure 68. How subset pointers divide a chain into subsets

Before you use subset pointers

For your program to use subset pointers, the pointers must be defined in the DBD for the DEDB and in your program's PSB:

- In the DBD, you specify the number of pointers for a segment chain. You can specify as many as eight pointers for any segment chain.
- In the PSB, you specify which pointers your program is to use. Define this on the SENSEG statement. (Each pointer is defined as an integer from 1 to 8.) Also, indicate on the SENSEG statement whether your program can set the pointers it uses. If your program has read sensitivity, it cannot set pointers but can only retrieve segments using subset pointers that are already set. If your program has update sensitivity, it can also update subset pointers by using the S, W, M, and Z command codes.

After the pointers are defined in the DBD and the PSB, an application program can set the pointers to segments in a chain. When an application program finishes executing, the subset pointers used by that program remain as they were set by the program; they are not reset.

Designating subset pointers

To use subset pointers in your program, you must know the numbers for the pointers as they were defined in the PSB. When you use the subset pointer command codes, specify the number of each subset pointer you want to use followed by the command code. For example, you use R3 to indicate that you want to retrieve the first segment occurrence in the subset defined by subset pointer 3. No default exists, so if you do not include a number between 1 and 8, IMS considers your SSA invalid and returns an AJ status code.

Subset pointer command codes

To take advantage of subsets, application programs use five command codes. The R command code retrieves the first segment in a subset. The following 4 command codes, which are explained in the topic "DEDB command codes for DL/I" in *IMS Version 15.3 Application Programming APIs*, redefine subsets by modifying the subset pointers:

Z

Sets a subset pointer to 0.

M

Sets a subset pointer to the segment following the current segment.

S

Unconditionally sets a subset pointer to the current segment.

W

Conditionally sets a subset pointer to the current segment.

Before your program can set a subset pointer, it must establish a position in the database. A call must be fully satisfied before a subset pointer is set. The segment a pointer is set to depends on your current position at the completion of the call. If a call to retrieve a segment is not completely satisfied and a position is not established, the subset pointers remain as they were before the call was made. You can use subset pointer command codes in either an unqualified SSA or a qualified SSA. To use a command code in a call with an unqualified SSA, use the command code along with the number of the subset pointer you want, after the segment name. This is shown in the following figure.

Table 53. Unqualified SSA with subset pointer command code

Seg Name	*	Cmd Code	Ssptr.	b
8	1	Variable	Variable	1

To use a subset pointer command code with a qualified SSA, use the command code and subset pointer number immediately before the left parenthesis of the qualification statement, as shown in the following figure.

Table 54. Qualified SSA with subset pointer command code

Seg Name	*	Cmd Code	Ssptr.	(	Fld Name	R.O.	Fld Value	)
8	1	Variable	Variable	1	8	2	Variable	1

Inserting segments in a subset

When you use the R command code to insert an unkeyed segment in a subset, the new segment is inserted before the first segment occurrence in the subset. However, the subset pointer is not automatically set to the new segment occurrence.

For example, the following call inserts a new B segment occurrence in front of segment B5, but does not set subset pointer 1 to point to the new B segment occurrence:

```
ISRT  Abbbbbbb  
(Akeybbbb=bA1)  
Bbbbbbbb  
*R1
```

To set subset pointer 1 to the new segment, you use the S command code along with the R command code, as shown in the following example:

```
ISRT  Abbbbbbb  
(Akeybbbb=bA1)  
Bbbbbbbb  
*R1S1
```

If the subset does not exist (subset pointer 1 is set to 0), the segment is added to the end of the segment chain.

Deleting the segment pointed to by a subset pointer

If you delete the segment pointed to by a subset pointer, the subset pointer points to the next segment occurrence in the chain. If the segment you delete is the last segment in the chain, the subset pointer is set to 0.

Combining command codes

You can use the S, M, and W command codes with other command codes, and you can combine subset pointer command codes with each other, as long as they do not conflict. For example, you can use R and S together, but you cannot use S and Z together because their functions conflict. If you combine command codes that conflict, IMS returns an AJ status code to your program.

You can use one R command code for each SSA and one update command code (Z, M, S, or W) for each subset pointer.

Related concepts

[“SSAs and command codes” on page 174](#)

SSAs can also include one or more command codes, which can change and extend the functions of DL/I calls.

[“Calls with dependent segments for DEDBs” on page 320](#)

You can issue DL/I calls against direct and sequential dependent segments for DEDBs.

Subset pointer status codes

If you make an error in an SSA that contains subset pointer command codes, IMS can return either of these status codes to your program.

AJ

The SSA used an R, S, Z, W, or M command code for a segment that does not have subset pointers defined in the DBD.

The subset command codes included in the SSA are in conflict. For example, if one SSA contains an S command code and a Z command code for the same subset pointer, IMS returns an AJ status code. S indicates that you want to set the pointer to current position; Z indicates that you want to set the pointer to 0. You cannot use these command codes in one SSA.

The SSA includes more than one R command code.

The pointer number following a subset pointer command code is invalid. You either did not include a number, or you included an invalid character. The number following the command code must be between 1 and 8.

AM

The subset pointer referenced in the SSA is not specified in the program's PSB. For example, if your program's PSB specifies that your program can use subset pointers 1 and 4, and your SSA references subset pointer 5, IMS returns an AM status code.

Your program tried to use a command code that updates the pointer (S, W, or M), but the program's PSB did not specify pointer-update sensitivity.

Your program attempted to open a GSAM database without specifying an IOAREA.

Processing DEDBs with a secondary index

Application programs can process a secondary index for DEDB databases of either HISAM or SHISAM database structures.

A HISAM secondary index database or a SHISAM secondary index database offers sequential key secondary index support.

A DEDB database with sequential dependent (SDEP) segments can have a secondary index database. SDEP segments cannot be used as an index field. Therefore, a SDEP segment cannot have LCHILD or XDFLD statements defined under its SEGM statement. Because SDEP segments are transient data and they are deleted using SDEP SCAN and SDEP DELETE utilities, Fast Path secondary index support for SDEP segments is restricted; that is, a SDEP segment cannot be a target segment or a source segment for a secondary index database. When the target segment is a root segment, SDEP segments can be returned for a DEDB database that is accessed through its alternate sequence.

Fast Path secondary indexing supports both unique and non-unique keys. A HISAM secondary index database offers unique and non-unique key support, and a SHISAM secondary index offers unique key support only.

A HISAM secondary index database supports both unique and non-unique keys. For a HISAM secondary index database, the non-unique key support is provided using an ESDS overflow data set. Duplicate keys

are stored in Last-In First-Out (LIFO) order. The first inserted duplicate key is stored in the KSDS data set and the remaining duplicate keys are stored in the ESDS overflow data set in LIFO order.

The target segment is in the primary DEDB database. The target segment is the segment that an application program needs to retrieve. The target segment can be at any one of the 15 levels in a primary DEDB database. SDEP segments cannot be a target segment or a source segment for a secondary index database.

There are a maximum of 32 secondary indexes per segment and 255 secondary indexes per DEDB database.

A Fast Path secondary index database can be accessed as:

- Its own database.
- A secondary index to its primary DEDB database, with an option to have single (the default) or multiple (by DEDB implementation only) secondary index segments.
- Fast Path secondary index user partitions.
- An option to suppress index maintenance for BMP applications.
- An option to access Fast Path secondary index user partition databases as one logical separate database.

Fast Path secondary indexes provide support for boolean qualification that are similar to full-function DL/I calls. The boolean operators supported are:

- Logical AND (coded * or &)
- Logical OR (coded + or |)

Restrictions

The restrictions that apply to processing DEDBs with a secondary index are:

- A DEDB database with sequential dependent (SDEP) segments can have a secondary index database, but SDEP segments cannot be used as an index field. Therefore, a SDEP segment cannot have LCHILD or XDFLD statements defined under its SEGM statement. Because SDEP segments are transient data and they are deleted using SDEP SCAN and SDEP DELETE utilities, Fast Path secondary index support for SDEP segments is restricted; that is, a SDEP segment cannot be a target segment or a source segment for a secondary index database.
- Fast Path secondary indexing does not support shared secondary indexes. Multiple secondary index segments support is only for DEDB implementation, and is not the same as shared secondary indexes support.
- A Fast Path secondary index database supports only symbolic pointers. There is no direct pointer support. Using symbolic pointers for indexed segments, a Fast Path secondary index database is not impacted when its primary DEDB database is reorganized.
- A qualified GU/GN segment name with SSA using the primary key field for target=root segment is supported when a primary DEDB database is accessed through its secondary index using a PCB with the PROCSEQD= parameter.
- A qualified GU/GN segment name with SSA using the primary key field for target=dependent segment is not supported. An AC status code is returned for the qualified Get call when a primary DEDB database is accessed through its secondary index using a PCB with the PROCSEQD= parameter.
- The independent AND (#) boolean operator is not supported.
- No boolean support is provided for SSAs with XDFLD and fields from the target segment. Boolean support is only for XDFLDs.

Example 1 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I calls to access COURSE segment on a primary DEDB database through its secondary index using GU and GN DL/I calls. For simplicity, assume that there is only one segment instance for each dependent segment.

PCB2NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, NAMESXDB. The COURSE segment is the target segment and it is a root segment. The source segment is the same as the target segment.

```
PCB2NDX
PCB TYPE=DB,DBDNAME=EDUCDB,PROCOPT=GR,KEYLEN=100,
  PROCSEQD=NAMESXDB
SENSEG NAME=COURSE,PARENT=0      <<- (target seg=root)
SENSEG NAME=CLASS,PARENT=COURSE
SENSEG NAME=INSTRUCT,PARENT=CLASS
SENSEG NAME=STUDENT,PARENT=CLASS
PSBGEN PSBNAME=NAMEXPSB,LANG=COBOL
END
PCB PCB2NDX
GU COURSE (NAMEINDX=CHEMISTRY)
GN COURSE
```

GU COURSE returns the COURSE segment for CHEMISTRY in the primary DEDB database using the secondary index key NAMEINDX=CHEMISTRY.

The key of the pointer segment, CHEMISTRY, is returned in the key feedback area.

GN COURSE returns the COURSE segment in the primary DEDB database that is pointed by the next pointer segment after segment CHEMISTRY in the Fast Path secondary index database, NAMESXDB.

The key of the next pointer segment after CHEMISTRY (the next sequential key after the secondary index key CHEMISTRY) in the Fast Path secondary index database, NAMESXDB, is returned in the key feedback area.

Example 2 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I calls to access COURSE segment on a primary DEDB database through its secondary index using GU and GN DL/I calls. For simplicity, assume that there is only one segment instance for each dependent segment.

PCB2NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, NAMESXDB. The COURSE segment is the target segment and it is a root segment. The source segment is the same as the target segment.

```
PCB PCB2NDX
GU COURSE (NAMEINDX=CHEMISTRY)
GN 1st GN
GN 2nd GN
GN 3rd GN
GN 4th GN
```

GU COURSE returns the COURSE segment for CHEMISTRY in the primary DEDB database using the secondary index key NAMEINDX=CHEMISTRY.

The key of the pointer segment, CHEMISTRY, is returned in the key feedback area.

The first GN call returns the segment of the DEDB inverted structure of the CHEMISTRY COURSE segment in the primary DEDB database. Because the COURSE segment is the target segment and it is a root segment, all segments in the physical structure are accessible as defined in PCB PCB2INDX. GN returns the CLASS segment in the database record under the CHEMISTRY COURSE segment that was retrieved by the GU call in the primary DEDB database.

The key of the pointer segment, CHEMISTRY, concatenated with the key of the CLASS segment under the CHEMISTRY COURSE segment, is returned in the key feedback area.

The second GN call returns the INSTRUCT segment in the database record under the CHEMISTRY COURSE segment

The key of the pointer segment, CHEMISTRY, concatenated with the key of the CLASS segment and the key of the INSTRUCT segment under the CHEMISTRY COURSE segment, is returned in the key feedback area.

The third GN call returns the STUDENT segment in the database record under the CHEMISTRY COURSE segment.

The key of the pointer segment, CHEMISTRY, concatenated with the key of the CLASS segment and the key of the STUDENT segment under the CHEMISTRY COURSE segment, is returned in the key feedback area.

Because the STUDENT segment is the last segment in the COURSE database record in the primary DEDB database, the fourth GN call returns the COURSE segment in the primary DEDB database using the next secondary index key after the CHEMISTRY segment in the secondary index database, NAMESXDB.

The key of the next pointer segment after CHEMISTRY (the next sequential key after the secondary index key CHEMISTRY) in the Fast Path secondary index database, NAMESXDB, is returned in the key feedback area.

Example 3 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I calls to access COURSE segment on a primary DEDB database through its secondary index using GU and GNP DL/I calls. For simplicity, assume that there is only one segment instance for each dependent segment.

PCB2NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, NAMESXDB. The COURSE segment is the target segment and it is a root segment. The source segment is the same as the target segment.

PCB	PCB2NDX	
GU	COURSE (NAMEINDX=CHEMISTRY)	
GNP		1st GNP
GNP		2nd GNP
GNP		3rd GNP
GNP		4th GNP

GU COURSE returns the COURSE segment for CHEMISTRY in the primary DEDB database using the secondary index key NAMEINDX=CHEMISTRY.

The key of the pointer segment, CHEMISTRY, is returned in the key feedback area.

The first GNP call returns the first segment under the DEDB inverted structure of the CHEMISTRY COURSE segment in the primary DEDB database. Because the COURSE segment is the target segment and it is a root segment, all segments in the physical structure are accessible as defined in PCB2INDX. GNP returns the CLASS segment in the database record under the CHEMISTRY COURSE segment that was retrieved by the GU call in the primary DEDB database.

The key of the pointer segment, CHEMISTRY, concatenated with the key of the CLASS segment under the CHEMISTRY COURSE segment, is returned in the key feedback area.

The second GNP call returns the INSTRUCT segment in the database record under the CHEMISTRY COURSE segment.

The key of the pointer segment, CHEMISTRY, concatenated with the key of the CLASS segment and the key of the INSTRUCT segment under the CHEMISTRY COURSE segment, is returned in the key feedback area.

The third GNP call returns the STUDENT segment in the database record under the CHEMISTRY COURSE segment.

The key of the pointer segment, CHEMISTRY, concatenated with the key of the CLASS segment and the key of the STUDENT segment under the CHEMISTRY COURSE segment, is returned in the key feedback area.

Since the STUDENT segment is the last segment in the COURSE database record in the primary DEDB database, the fourth GNP call returns a GE status code.

Example 4 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I calls to access INSTRUCT segment on a primary DEDB database through its secondary index using GU and GN DL/I calls. For simplicity, assume that there is only one segment instance for each dependent segment.

PCB3NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, INSTSXDB. The CLASS segment is the target segment and it is not a root segment. The source segment (INSTRUCT segment) is not the same as the target segment (CLASS segment).

Although there are multiple LCHILD/XDFLD pairs in the CLASS SEGM statement, only one is used in this example.

DBDGEN excerpt for CLASS and INSTRUCT SEGM statements in the EDUCDB DEDB DBD:

```
SEGM  NAME=CLASS,BYTES=50,PARENT=COURSE
FIELD NAME=(CLASSNO,SEQ,U),BYTES=4,START=7
FIELD NAME=CLASNAME,BYTES=10,START=15

LCHILD NAME=(CLASXSEG,CLASSCDB),PTR=SYMB
XDFLD  NAME=CLASINDX,SRCH=CLASNAME

LCHILD NAME=(INSTXSEG,INSTSXDB),PTR=SYMB
XDFLD  NAME=INSTINDX,SEGMENT=INSTRUCT,SRCH=INSTNAME
SEGM  NAME=INSTRUCT,BYTES=50,PARENT=CLASS
FIELD NAME=(INSTNO,SEQ,U),BYTES=6,START=1
FIELD NAME=INSTPHNO,BYTES=10,START=11
FIELD NAME=INSTNAME,BYTES=20,START=21
...
```

PSBGEN Definition for PCB3NDX:

When the target segment is not a root segment, all direct parents of the target segment from the root segment must be defined in the PCB with the PROCSEQD parameter. Only the direct parents segments along the physical path from the root segment to the target segment and all child segments of the target segment are accessible when the target segment is not a root segment. All sibling segments of CLASS are not accessible. The coding sequence of the mandatory SENSEGs must be in the sequence of the physical path of the segments (for example, from the physical root to the target) even though the segments are retrieved always in logical sequence (for example, from the target or logical root to the physical root).

```
PCB3NDX
PCB  TYPE=DB,DBDNAME=EDUCDB,PROCOPT=GR,KEYLEN=100,
     PROCSEQD=INSTSXDB
SENSEG NAME=COURSE,PARENT=0          <<- mandatory SENSEG
SENSEG NAME=CLASS,PARENT=COURSE      <<- mandatory SENSEG (target seg)
SENSEG NAME=INSTRUCT,PARENT=CLASS    <<- optional SENSEG
PSBGEN PSBNAME=NAMEXPSB,LANG=COBOL
END

PCB  PCB3NDX
GU   CLASS (INSTINDX=TOMJONES)
GN
GN   1st GN
GN   2nd GN
GN   3rd GN
```

GU CLASS returns the CLASS segment for the instructor teaching the class, in the primary DEDB database using the secondary index key INSTINDX=TOMJONES.

The key of the pointer segment, TOMJONES, is returned in the key feedback area.

The first GU call returns the target segment of the primary DEDB database. Because the target segment (CLASS) is not a root segment, the subsequent GN returns the next segment in the DEDB inverted structure of the CLASS segment retrieved by the GU call in the primary DEDB database. For example, GN returns the COURSE segment which is a direct physical parent and also a logical child of the CLASS segment teaching the CHEMISTRY COURSE segment in the DEDB inverted structure.

The key of the pointer segment, TOMJONES, concatenated with the key of the COURSE segment, is returned in the key feedback area.

The second GN call returns the INSTRUCT segment in the database record, which is a logical child of the CLASS segment and a logical sibling of the COURSE segment in the DEDB inverted structure.

The key of the pointer segment, TOMJONES, concatenated with the key of the INSTRUCT segment is returned in the key feedback area.

Because no child or sibling segment is defined for the INSTRUCT segment in PCB PCB3NDX, the third GN call returns the CLASS segment in the primary DEDB database using the next segment in the secondary index database after INSTINDX=TOMJONES.

The key of the next pointer segment after TOMJONES (the next sequential key after the secondary index key TOMJONES) in the Fast Path secondary index database, INSTSXDB, is returned in the key feedback area.

Example 5 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I calls to access INSTRUCT segment on a primary DEDB database through its secondary index using GU and GNP DL/I calls. For simplicity, assume that there is only one segment instance for each dependent segment.

PCB3NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, INSTSXDB. The CLASS segment is the target segment and it is not a root segment. The source segment (INSTRUCT segment) is not the same as the target segment (CLASS segment).

Although there are multiple LCHILD/XDFLD pairs in the CLASS SEGM statement, only one is used in this example.

DBDGEN excerpt for CLASS and INSTRUCT SEGM statements in the EDUCDB DEDB DBD:

```
...
SEGM   NAME=CLASS,BYTES=50,PARENT=COURSE
FIELD  NAME=(CLASSNO,SEQ,U),BYTES=4,START=7)
FIELD  NAME=CLASNAME,BYTES=10,START=15

LCHILD NAME=(CLASXSEG,CLASSXDB),PTR=SYMB
XDFLD  NAME=CLASINDX,SRCH=CLASNAME

LCHILD NAME=(INSTXSEG,INSTSXDB),PTR=SYMB
XDFLD  NAME=INSTINDX,SEGMENT=INSTRUCT,SRCH=INSTNAME

SEGM   NAME=INSTRUCT,BYTES=50,PARENT=CLASS
FIELD  NAME=(INSTNO,SEQ,U),BYTES=6,START=1
FIELD  NAME=INSTPHNO,BYTES=10,START=11
FIELD  NAME=INSTNAME,BYTES=20,START=21
...
```

PSBGEN Definition for PCB3NDX:

When the target segment is not a root segment, all direct parents of the target segment from the root segment must be defined in the PCB with the PROCSEQD parameter. Only the direct parents segments along the physical path from the root segment to the target segment and all child segments of the target segment are accessible when the target segment is not a root segment.

```
PCB3NDX
PCB  TYPE=DB,DBDNAME=EDUCDB,PROCOPT=GR,KEYLEN=100,
     PROCSEQD=INSTSXDB
SENSEG NAME=COURSE,PARENT=0
SENSEG NAME=CLASS,PARENT=COURSE      <<-- Target segment
SENSEG NAME=INSTRUCT,PARENT=CLASS
PSBGEN PSBNAME=NAMEXPSB,LANG=COBOL
END

PCB  PCB3NDX
GU   CLASS(INSTINDX=TOMJONES)
GNP                                     1st GNP
GNP                                     2nd GNP
GNP                                     3rd GNP
```

GU CLASS returns the CLASS segment for the instructor teaching the particular class in the primary DEDB database using the secondary index key INSTINDX=TOMJONES.

The key of the pointer segment, TOMJONES, is returned in the key feedback area.

The first GU call returns the CLASS segment of the DEDB. Because the target segment (CLASS) is not a root segment, the first GNP returns the next segment in the DEDB inverted structure. For example, the first GNP returns the COURSE segment which is a direct physical parent, but a direct logical child, of the CLASS segment in the DEDB inverted structure.

The key of the pointer segment, TOMJONES, concatenated with the key of the COURSE segment, is returned in the key feedback area.

The second GNP call returns the INSTRUCT segment in the database record, which is a logical child of CLASS and a logical sibling of the COURSE segment in the inverted DEDB structure hierarchy.

The key of the pointer segment, TOMJONES, concatenated with the key of the INSTRUCT segment is returned in the key feedback area.

Because there is no child or sibling segment defined for the INSTRUCT segment in PCB PCB3NDX, the third GNP call under the CLASS segment with the secondary index key of TOMJONES returns a GE status code.

Example 6 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I calls to access CLASS segment on a primary DEDB database through its secondary index using GU and GN DL/I calls. For simplicity, assume that there is only one segment instance for each dependent segment.

PCB4NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, CLASSXDB. The CLASS segment is the target segment and it is not a root segment. The source segment is the same as the target segment.

PSBGEN Definition for PCB4NDX:

When the target segment is not a root segment, all direct parents of the target segment from the root segment must be defined in the PCB with the PROCSEQD parameter. Only the direct parents segments along the physical path from the root segment to the target segment and all child segments of the target segment are accessible when the target segment is not a root segment.

```
PCB4NDX
PCB  TYPE=DB, DBDNAME=EDUCDB, PROCOPT=GR, KEYLEN=100,
      PROCSEQD=CLASSXDB
SENSEG  NAME=COURSE, PARENT=0
SENSEG  NAME=CLASS, PARENT=COURSE    <--- Target segment
SENSEG  NAME=INSTRUCT, PARENT=CLASS
SENSEG  NAME=STUDENT, PARENT=CLASS
PSBGEN  PSBNAME=NAMEXPSB, LANG=COBOL
END

PCB    PCB4NDX
GU     CLASS(CLASINDX=CHEM1A)
GN
GN      1st GN
GN      2nd GN
GN      3rd GN
GN      4th GN
```

GU CLASS returns the CLASS segment for the class name, CHEM1A, in the primary DEDB database using the secondary index key INSTINDX=CHEM1A.

The key of the pointer segment, CHEM1A, is returned in the key feedback area.

The first GN call returns the COURSE segment of the DEDB inverted structure of the CHEM1A CLASS segment in the primary DEDB database. Because the CLASS segment is the target segment and it is not a root segment, GN returns the next segment in the DEDB inverted structure of the CLASS segment retrieved by the GU call in the primary DEDB database. GN returns the COURSE segment for CHEMISTRY which is a direct parent of the CHEM1A CLASS segment in the DEDB inverted structure.

The key of the pointer segment, CHEM1A, concatenated with the primary key of the COURSE segment, is returned in the key feedback area.

The second GN call returns the INSTRUCT segment in the database record under CHEM1A CLASS segment.

The key of the pointer segment, CHEM1A, concatenated with the key of the INSTRUCT segment under the CHEM1A CLASS segment, is returned in the key feedback area.

The third GN call returns the STUDENT segment in the database record under CHEM1A CLASS segment.

The key of the pointer segment, CHEM1A, concatenated with the key of the STUDENT segment under the CHEM1A CLASS segment, is returned in the key feedback area.

Because the STUDENT segment is the last segment for the CLASS segment in the database record, the fourth GN call returns the CLASS segment in the primary DEDB database using the next secondary index key after the CHEM1A segment in the secondary index database, CLASINDEX.

The key of the next pointer segment after CHEM1A (the next sequential key after the secondary index key CHEM1A) in the Fast Path secondary index database, CLASSXDB, is returned in the key feedback area.

Example 7 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I calls to access class name on a primary DEDB database through its secondary index using GU and GN DL/I calls. For simplicity, assume that there is only one segment instance for each dependent segment.

PCB4NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, CLASSXDB. The CLASS segment is the target segment and it is a not root segment. The source segment is the same as the target segment.

PCB	PCB4NDX
GU	CLASS (CLASINDEX=CHEM1A)
GN	CLASS

GU CLASS returns the CLASS segment for CHEM1A in the primary DEDB database using the secondary index key NAMEINDEX=CHEM1A.

The key of the pointer segment, CHEM1A, is returned in the key feedback area.

GN CLASS returns the CLASS segment in the primary DEDB database that is pointed to by the next pointer segment after segment CHEM1A in the Fast Path secondary index database, CLASSXDB.

The key of the next pointer segment after CHEM1A (the next sequential key after the secondary index key CHEM1A) in the Fast Path secondary index database, CLASSXDB, is returned in the key feedback area.

Example 8 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I calls to access the class name on a primary DEDB database through its secondary index using GU and GN DL/I calls and C command code. For simplicity, assume that there is only one segment instance for each dependent segment.

PCB4NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, CLASSXDB. The CLASS segment is the target segment and it is a not root segment. The source segment is the same as the target segment.

PCB	PCB4NDX	
GU	INSTRUCT *C (CHEM1AI12345)	
GN		1st GN
GN		2nd GN

GU INSTRUCT returns the INSTRUCT segment for instructor number, I12345, under the CLASS segment for class name, CHEM1A.

The key of the pointer segment, CHEM1A, concatenated with the key of INSTRUCT segment, I12345, is returned in the key feedback area.

The first GN call returns the STUDENT segment of the DEDB inverted structure of the CHEM1A CLASS segment in the primary DEDB database. Because the CLASS segment is the target segment and it is not a root segment, GN returns the next segment in the DEDB inverted structure of the INSTRUCT segment retrieved by the GU call in the primary DEDB database. GN returns the STUDENT segment which is a child segment of the CHEM1A CLASS segment in the DEDB inverted structure.

The key of the pointer segment, CHEM1A, concatenated with the key of the STUDENT segment, is returned in the key feedback area.

Because the STUDENT segment is the last segment for the CLASS segment in the database record, the second GN call returns the CLASS segment in the primary DEDB database using the next secondary index key after the CHEM1A segment in the secondary index database, CLASINDEX.

The key of the next pointer segment after CHEM1A (the next sequential key after the secondary index key CHEM1A) in the Fast Path secondary index database, CLASSXDB, is returned in the key feedback area.

Example 9 of accessing a primary DEDB database that uses a Fast Path secondary index

DL/I call to insert an INSTRUCT segment on a primary DEDB database through its secondary index for the CLASS segment with a secondary index key of CHEM1A.

PCB4NDX is the PCB with the PROCSEQD= parameter defined to use the Fast Path secondary index database, CLASSXDB. The CLASS segment is the target segment and it is a not root segment. The source segment is the same as the target segment.

```
PCB4NDX
ISRT CLASS(CLASINDEX = CHEM1A)
      INSTRUCT I23456 JOHN SMITH
```

The ISRT call inserts an INSTRUCT segment with the key of I23456 under the CLASS segment with a secondary index key of CHEM1A

The key of the pointer segment, CHEM1A, concatenated with the key of the INSTRUCT segment, I23456, is returned in the key feedback area.

Related concepts

[Creating secondary indexes \(Database Administration\)](#)

Related tasks

[Adding a secondary index to a DEDB \(Database Administration\)](#)

Retrieving location with the POS call (for DEDB only)

Use the POS (Position) call to retrieve the location of a specific sequential dependent segment; retrieve the location of the last-inserted sequential dependent segment, its time stamp, and the IMS ID; or retrieve the time stamp of a sequential dependent or Logical Begin. You can also use the POS call to tell the amount of unused space within each DEDB area. For example, you can use the information that IMS returns for a POS call to scan or delete the sequential dependent segments for a particular time period.

The topic "POS Call" in *IMS Version 15.3 Application Programming APIs* explains how you code the POS call and what the I/O area for the POS call looks like. If the area that the POS call specifies is unavailable, the I/O area is unchanged, and the FH status code is returned.

Locating a specific sequential dependent

When you have position on a particular root segment, you can retrieve the position information and the area name of a specific sequential dependent of that root. If you have a position established on a sequential dependent segment, the search starts from that position. IMS returns the position information for the first sequential dependent segment that satisfies the call. To retrieve this information, issue a POS call with a qualified or unqualified SSA containing the segment name of the sequential dependent. Current position after this kind of POS call is the same place that it would be after a GNP call.

After a successful POS call, the I/O area contains:

LL

A 2-byte or 4 byte field giving the total length of the data in the I/O area. When the number of DEDB areas is 2048 or fewer, the LL is 2-byte field. When the number of DEDB areas is greater than 2048, the LL is 4-byte field.

Area Name

An 8-byte field giving the ddname from the AREA statement.

Position

An 8-byte field containing the position information for the requested segment.

Exception: If the sequential dependent segment that is the target of the POS call is inserted in the same synchronization interval, no position information is returned. Bytes 11-18 contain X'FF'. Other fields contain normal data.

Unused CIs

A 4-byte field containing the number of unused CIs in the sequential dependent part.

Unused CIs

A 4-byte field containing the number of unused CIs in the independent overflow part.

Locating the last inserted sequential dependent segment

You can also retrieve the position information for the most recently inserted sequential dependent segment of a given root segment. To do this, you issue a POS call with an unqualified or qualified SSA containing the root segment as the segment name. Current position after this type of call follows the same rules as position after a GU call.

You can also retrieve the position of the SDEP, its time stamp, and the ID of the IMS that owns the segment. To do this, you issue a POS call with a qualified SSA and provide the keyword PCSEGTSP in position one of the I/O area as input to the POS call. The keyword requests the POS call to return the position of the SDEP, its time stamp, and the ID of the IMS that owns the segment.

Requirement: The I/O area must be increased in size to 42 bytes to allow for the added data being returned. The I/O area includes a 2-byte LL field that is not shown in the following table. This LL field is described after the following table.

Table 55. Qualified POS call: keywords and map of I/O area returned

Keyword	word 0	word 1	word 2	word 3	word 4	word 5	word 6	word 7	word 8	word 9
<null>	Field 1		Field 2	Field 3	Field 4		N/A		N/A	
PCSEGTSP	Field 1		Field 2		Field 5		Field 6		Field 7	

Field 1

Area name

Field 2

Sequential dependent location from qualified SSA

Field 3

Unused CIs in sequential dependent part

Field 4

Unused CIs in independent overflow part

Field 5

Committed sequential dependent segment time stamp

Field 6

IMS ID

Field 7

Pad

After a successful POS call, the I/O area contains:

LL

(Not shown in table) A 2-byte field, in binary, containing the total length of the data in the I/O area.

(Field 1)

Area Name

An 8-byte field giving the ddname from the AREA statement.

(Field 2)

Position

An 8-byte field containing the position information for the most recently inserted sequential dependent segment. This field contains zeros if no sequential dependent exists for this root.

Sequential dependent location from qualified SSA

IMS places two pieces of data in this 8-byte field after a successful POS call. The first 4 bytes contain the cycle count, and the second 4 bytes contain the VSAM RBA.

If the sequential dependent segment that is the target of the POS call is inserted in the same synchronization interval, no position information is returned. Bytes 11-18 contain X'FF'. Other fields contain normal data.

(Field 3)

Unused CIs in sequential dependent part

A 4-byte field containing the number of unused control intervals in the sequential dependent part.

(Field 4)

Unused CIs in independent overflow part

A 4-byte field containing the number of unused control intervals in the independent overflow part.

(Field 5)

Committed Sequential Dependent Segment Time Stamp

An 8-byte field containing the time stamp that corresponds to the SDEP segment located by the qualified POS call.

(Field 6)

IMS ID

Identifies the IMS that owns the CI where the SDEP segment was located.

(Field 7)

Pad

An 8-byte pad area to align the I/O area on a double word boundary. No data is returned to this field.

Identifying free space

To retrieve the area name and the next available position within the sequential dependent part from all online areas, you can issue an unqualified POS call. This type of call also retrieves the unused space in the independent overflow and sequential dependent parts.

After a unsuccessful unqualified POS call, the I/O area contains the length (LL), followed by the same number of entries as existing areas within the database. Each entry contains the fields shown below:

Area Name

An 8-byte field giving the ddname from the AREA.

Position

An 8-byte field with binary zeros.

Unused SDEP CIs

A 4-byte field with binary zeros.

Unused IOV CIs

A 4-byte field with two binary zeros followed by a bad status code.

Commit-point processing in a DEDB

IMS retains database updates in processor storage until the program reaches a commit point. IMS saves updates to a DEDB in Fast Path buffers. The database updates are not applied to the DEDB until after the program has successfully completed commit-point processing.

Unlike Get calls to an MSDB, however, a Get call to an updated segment in a DEDB returns the updated value, even if a commit point has not occurred.

When a BMP is processing DEDBs, it must issue a CHKP or SYNC call to do commit-point processing before it terminates. Otherwise, the BMP abnormally terminates with abend U1008.

Related concepts

[“Commit-point processing in MSDBs and DEDBs” on page 302](#)

Your existing application programs can use either the MSDB commit view or the default DEDB commit view.

P processing option

If the P processing option is specified in the PCB for your program, a GC status code is returned to your program whenever a call to retrieve or insert a segment causes a unit of work (UOW) boundary to be crossed.

Related reading: For more information on the UOW for DEDBs, see *IMS Version 15.3 Database Administration*.

Although crossing the UOW boundary probably has no particular significance for your program, the GC status code indicates that this is a good time to issue either a SYNC or CHKP call. The advantages of issuing a SYNC or CHKP call after your program receives a GC status code are:

- Your position in the database is retained. Issuing a SYNC or CHKP call normally causes position in the database to be lost, and the application program must reestablish position before it can resume processing.
- Commit points occur at regular intervals.

When a GC status code is returned, no data is retrieved or inserted. In your program, you can either:

- Issue a SYNC or CHKP call, and resume database processing by reissuing the call that caused the GC status code.
- Ignore the GC status code, and resume database processing by reissuing the call that caused the status code.

Related concepts

[“Calls with dependent segments for DEDBs” on page 320](#)

You can issue DL/I calls against direct and sequential dependent segments for DEDBs.

H processing option

If the H processing option has been specified in the PCB for your call program, a GC status code is returned whenever a call to retrieve or insert a segment causes a unit of work (UOW) or an area boundary to be crossed. The program must cause a commit process before any other calls can be issued to that PCB.

If a commit process is not caused, an FR status code results (total buffer allocation exceeded), and all database changes for this synchronization interval are "washed" (sync-point failure).

A GC status code is returned when crossing the area boundary so that the application program can issue a SYNC or CHKP call to force cleanup of resources (such as buffers) that were obtained in processing the previous area. This cleanup might cause successive returns of a GC status code for a GN or GHN call, even if a SYNC or CHKP call is issued appropriately for the previous GC status code.

When an application is running HSSP and proceeding through the DEDB AREA sequentially, a buffer shortage condition may occur due to large IOV chains. In this case, a FW status code is returned to

the application. Usually, the application issues a commit request and position is set to the next UOW. However, this does not allow the previous UOW to finish processing. In order to finish processing the previous UOW, you can issue a commit request after the FW status code is received and set the position to remain in the same UOW. You must also reposition the application to the position that gave the FW status code. The following shows an example of the command sequence and corresponding application responses.

GN		root1	
GN		root2	
GN		root3	
GN		root4	/*FW status code received*/
CHKP			
GN	SSA=(root4)	root4	/*User reposition prior to CHKP*/
GN		root5	

Calls with dependent segments for DEDBs

You can issue DL/I calls against direct and sequential dependent segments for DEDBs.

The DL/I calls that you can issue against a root segment are: GU, GN (GNP has no meaning for a root segment), DLET, ISRT, and REPL. You can issue all DL/I calls against a direct dependent segment, and you can issue Get and ISRT calls against sequential dependents segments.

Direct dependent segments

DL/I calls to direct dependents include the same number of SSAs as existing levels in the hierarchy (a maximum of 15). They can also include command codes and multiple qualification statements. The same rules apply to using command codes on DL/I calls to DEDBs as to full-function databases.

If you use the D command code in a call to a DEDB, the P processing option need not be specified in the PCB for the program. The P processing option has a different meaning for DEDBs than for full-function databases.

Some special command codes can be used only with DEDBs that use subset pointers. Your program uses these command codes to read and update the subset pointers.

Sequential dependent segments

Because sequential dependents are stored in chronological order, they are useful in journaling, data collection, and auditing application programs. You can access sequential dependents directly. However, sequential dependents are normally retrieved sequentially using the Database Scan utility.

Restriction: When processing sequential dependent segments:

- You can only use the F command code with sequential dependents; IMS ignores all other command codes.
- You cannot use Boolean operators in calls to sequential dependents.

Related reading: For more information about the utility, see *IMS Version 15.3 Database Utilities*.

Related concepts

[“Processing Fast Path DEDBs with subset pointer command codes” on page 304](#)

Subset pointers and the command codes you use with them are optimization tools that significantly improve the efficiency of your program when you need to process long segment chains.

Related reference

[“P processing option” on page 319](#)

If the P processing option is specified in the PCB for your program, a GC status code is returned to your program whenever a call to retrieve or insert a segment causes a unit of work (UOW) boundary to be crossed.

DEDB DL/I calls to extract DEDB information

DL/I calls can be issued to obtain structural information about Data Entry Databases (DEDBs). Any application that can issue DL/I calls can take advantage of these DL/I calls.

There are two basic call types:

- The first type returns the minimum I/O area length required for a specific type '2' DL/I call.
- The second type returns specific information about the specified DEDB.

Each of these DL/I calls uses a call interface block called the Input Output Input Area (IOAI), a telecommunication program PCB (TP PCB), and specific calls that require an I/O area. Some required initialization of the IOAI is common for all calls and some initialization is specific to an individual call. The IOAI and the I/O area must be obtained in key 8 storage.

The following table describes the DL/I calls to extract DEDB information.

Table 56. DEDB DL/I Calls	
DL/I Call	Description
AL_LEN	Returns the minimum length of the I/O area that is required for an AREALIST call.
DI_LEN	Returns the minimum length of the I/O area that is required for a DEDBINFO call.
DS_LEN	Returns the minimum length of the I/O area required for a DEDBSTR call.
AREALIST	Returns a list of areas that are part of the specified DEDB, with each area mapped by DBFCDAL1.
DEDBINFO	Returns DEDB information from the DMCB, mapped by DBFCDDI1.
DEDBSTR	Returns a list of segments and a segment data for DEDB with each segment mapped by DBFCDDS1.

The DL/I call that use the standard interface with register 1 must point to IOAI_CA.

The following figure shows the IOAI structure.

			starting offset	note
IOAI_START	DS	0F		
IOAI_NAME	DC	CL4 'IOAI'	0	*1
IOAI_#FPU	DC	CL4 '#FPU'	4	*1
IOAI_#FPI	DC	CL8 '#FPUCDPI'	8	*1
IOAI_SUBC	DC	CL8 ' '	10	*1
*				
IOAI_BLEN	DC	A(0)	18	*1
IOAI_ILEN	DC	A(0)	1C	*1
IOAI_IOAREA	DC	A(0)	20	*1
*				
IOAI_CALL	DC	A(0)	24	*1
IOAI_PCBI	DC	A(0)	28	*1
IOAI_IOAI	DC	A(0)	2C	*1
*				
IOAI_DLEN	DC	A(0)	30	*2
IOAI_STATUS	DC	CL2 ' '	34	*2
IOAI_B_LEVEL	DC	XL2 '0'	36	*2
IOAI_STATUS_RC	DC	A(0)	38	*2
IOAI_USERVER	DC	A(0)	3C	*1
IOAI_IMSVER	DC	A(0)	40	*2

```

*
IOAI_IMSLEVEL  DC   A(0)           44      *2
*
IOAI_APPL_NAME DC   CL8' '         48      *1
IOAI_USERDATA  DC   CL8' '         50      *1
IOAI_TIMESTAMP DC   CL8' '         58      *2
*
input words.
IOAI_IN0       DC   A(0)           60      *3
IOAI_IN1       DC   A(0)           64      *3
IOAI_IN2       DC   A(0)           68      *3
IOAI_IN3       DC   A(0)           6C      *3
IOAI_IN4       DC   A(0)           70      *3
*
feedback words
IOAI_FDBK0     DC   A(0)           74      *2
IOAI_FDBK1     DC   A(0)           78      *2
IOAI_FDBK2     DC   A(0)           7C      *2
IOAI_FDBK3     DC   A(0)           80      *2
IOAI_FDBK4     DC   A(0)           84      *2
*
workareas.
IOAI_WA0       DC   A(0)           88      *4
IOAI_WA1       DC   A(0)           8C      *4
IOAI_WA2       DC   A(0)           90      *4
IOAI_WA3       DC   A(0)           94      *4
IOAI_WA4       DC   A(0)           98      *4
*
DS 20F'0'      9C      for future expansion
IOAI_END_CHAR  DC   CL4' IEND'      EC      *1
IOAI_LEN       len(DBFIOAI) = x'F0' bytes

```

Note:

1. The user is responsible for initializing these fields.
2. IMS uses these fields to return data to the caller. Which fields contain returned data depends on the DL/I call and are documented in the section on the specific call types.
3. May be used to pass additional data on the DL/I call, as documented under each DL/I call.
4. These fields are unchanged, and can be used as work areas by the application.

The fields in the following table must be initialized for all of the following DL/I calls.

Table 57. Field initialization for DEDB DL/I calls	
Field	Description
IOAI_NAME	The characters 'IOAI' identifying this block.
IOAI_#FPU	The characters '#FPU' indicating this is a #FPU call.
IOAI_#FPI	The characters '#FPUCDPI' indicating this is a subset call.
IOAI_SUBC	The DL/I call: AL_LEN, AREALIST, DS_LEN, DEDBSTR, DI_LEN or DEDBINFO.
IOAI_BLEN	The total length of the IOAI (x'F0').
IOAI_CALL	Address of IOAI_#FPU.
IOAI_PCBI	Address of the TPCB.
IOAI_IOAI	Address of this block. The user must set the high order bit on to indicate the end of the DL/I list.
IOAI_USERVER	Call version number. Defaults to one. This is the version number of a specific call. This field will be updated in the future if a specific call is altered such that the application must be sensitive to the changes.
IOAI_END_CHAR	The chars 'IEND' identifying the end of block.

The following fields are initialized for specific DL/I calls. If a specific call does not need an I/O area, these fields are ignored.

Table 58. Fields initialized for specific DEDB DL/I calls	
Field	Description
IOAI_ILEN	The total length of the I/O area, including prefix and suffix.
IOAI_IOAREA	Address of the I/O area.
I/O Area	1st word: The I/O area length (same as IOAI_ILEN). Last word: X'FFFFFFFF', which is an 'end of I/O area' marker.
IOAI_IN0 -> IOAI_IN4	Five input words that might be required.

The following fields are updated by IMS for all the DEDB DL/I call types.

Table 59. Fields updated by IMS for all DL/I call types	
Field	Description
IOAI_DLEN	The length of the output data that is returned by IMS. This field is informational only.
IOAI_STATUS	A 2-byte status code.
IOAI_STATUS_RC	A return code if needed.
IOAI_IMSVER	The maximum version of this call.
IOAI_IMSLEVEL	The IMS level.

The following fields might be updated by specific DL/I calls.

Table 60. Fields updated by specific DL/I calls	
Field	Description
I/O Area	1st word: unchanged. Data: see specific call types. Last word: potentially changed.
IOAI_FDBK0 -> IOAI_FDBK4	Five output words which may return data as documented by specific calls.

```

DBFCDAL1 mapping:      offset
CDAL_START    DS      0F
CDAL_ARNM     DS      CL8      00      Area name
CDAL_FLGS     DS      0XL4     08      Flag Bytes
CDAL_FLG1     DS      XL1      08      Flags for area status:
CDAL_F10P     EQU      X'01'    - Area is opened
CDAL_F1BK     EQU      X'02'    - Temporary bit for backout
CDAL_F1UT     EQU      X'04'    - Utility active on this area
CDAL_F1ER     EQU      X'08'    - Error recovery needed
CDAL_F1AF     EQU      X'80'    - Sequential dep. part full
CDAL_F1EP     EQU      X'40'    - I/O error
CDAL_F1ST     EQU      X'20'    - Area stop request
CDAL_F1RE     EQU      X'10'    - Area restart request
CDAL_FLG2     DS      XL1      09      Reserved for Flag Byte #2
CDAL_FLG3     DS      XL1      0A      Reserved for Flag Byte #3
CDAL_FLG4     DS      XL1      0B      Reserved for Flag Byte #4
              DS      1F      0C      for growth
CDAL_LEND     DS      0F      End of area list entry
CDAL_LEN      EQU      *-&AA._START; Len of area list entry

```

```

DBFCDDI0 mapping:      offset
CDDI_START    DS      0D

```

```

CDDI_DBNM   DS    CL8    00    Database name
CDDI_ANR    DS    H      08    Number of areas defined
CDDI_HSLV   DS    H      0A    Max SEGM level in the DB
CDDI_SGNR   DS    H      0C    Highest valid SEGM code
CDDI_SEGL   DS    H      0E    Maximum IOA length
CDDI_HBLK   DS    F      10    Number of anchor blocks
CDDI_RMMN   DS    CL8    14    Randomizing module name
CDDI_RMEP   DS    F      1C    Randomizing module entry point
            DS    8F      20    Reserved
            DS    0D      20    Align on double word boundary
CDDI_LEN    EQU    *-&AA._START; Length of this area (x'40')

```

```

DBFCDDS1 mapping:      offset
CDDS_START   DS    0F      00
CDDS_GNAM    DS    CL8     00 SEGMENT NAME
CDDS_GDOF    DS    H       08 OFFSET FROM START SEQ TO DATA
CDDS_MAX     DS    H       0A MAX SEG LEN
CDDS_MIN     DS    H       0C MIN SEG LEN
CDDS_DBOF    DS    H       0E OFFSET TO SEG ENTRIES
CDDS_NRFLD   DS    FL1     10 NUMBER OF FIELDS IN SEG
CDDS_SC      DS    FL1     11 SEGMENT CODE
CDDS_PREF    DS    H       12 POINTER OFFSET IN PARENT PREF
CDDS_FLG1    DS    X       14 FLAG BYTE
CDDS_FL1K    EQU    X'80'   KEY SEGMENT
CDDS_FL1S    EQU    X'40'   SEQUENTIAL DEP SEGMENT
CDDS_FL1P    EQU    X'20'   PCL POINTER TO PARENT
CDDS_FISRT   DS    X       15 INSERT RULES
CDDS_PARA    DS    H       16 OFFSET TO PARENT SEGMENT
CDDS_SBLP    DS    F       18 SIBLING POINTER
CDDS_LEVL    DS    XL1     1C SEGMENT LEVEL
CDDS_KEYL    DS    XL1     1D KEY LENGTH - 1
CDDS_KDOF    DS    H       1E OFFSET TO KEY FIELD IN SEGMENT
CDDS_RSRVE   DS    XL4     20 FOR USE IN UMDRO | RESERVED
CDDS_CMPC    DS    A       24 A(CMPC)
CDDS_FLG2    DS    XL1     28 FLAG BYTE 2 (fixed length)
            DS    XL3     29 FOR GROWTH
            DS    5F      2C for growth
CDDS_END     DS    0F      END
CDDS_LEN     EQU    *-&AA._START; len of SDB entry

```

The following status codes are specific to these new DL/I calls.

Table 61. Status codes for specific DEDB DL/I calls	
Status Code	Description
AA	Invalid #FPU/#FPUCDPI call.
AB	Getmain error.
AC	DEDB name not found.
AD	The I/O area was not long enough to contain the data.
AE	IOAI_LEN was zeros. It must be filled by the caller.
AF	The I/O area address was not passed in by IOAI_IOAREA.
AG	The IOAI does not point to itself, IOAI_IOAI.
AH	The IOAI did not contain 'IOAI'.
AI	The I/O area length in the I/O area does not match IOAI.
AJ	The I/O area did not contain the end-of-list marker.
AK	The IOAI did not have end-of-block marker 'IEND'.
AL	IOAI_BLEN is not correct.

<i>Table 61. Status codes for specific DEDB DL/I calls (continued)</i>	
Status Code	Description
AM	DEDB not passed in via the IOAI on the #FPUCDPI call.
FX	The IOAREA of a POS call has a 4-byte LL field.

AL_LEN Call

The AL_LEN call returns the minimum length of the I/O area required for an AREALIST call.

Input

IOAI

Formatted and filled out as documented above.

IOAI_INO

Points to storage containing the DEB name.

Output

IOAI_STATUS

Call status, ' ' means successful.

IOAI_FDBK0

The minimum length of the I/O area.

IOAI_FDBK1

The number of AREAS in this DEDB.

DI_LEN Call

Return the minimum length of the I/O area required for an DEDBINFO call.

Input

IOAI

Formatted and filled out as documented above.

IOAI_INO

Points to storage containing the DEB name.

Output

IOAI_STATUS

Call status, ' ' means successful.

IOAI_FDBK0

The minimum length of the I/O area.

DS_LEN Call

Return the minimum length of the I/O area required for a DEDBSTR call.

Input

IOAI

Formatted and filled out as documented above.

IOAI_INO

Points to storage containing the DEB name.

Output

IOAI_STATUS

Call status, '' means successful.

IOAI_FDBK0

The minimum length of the I/O area.

IOAI_FDBK1

The number of SEGMENTS in this DEDB.

AREALIST Call

The AREALIST call returns a list of areas that are part of the specified DEDB, with each area mapped by DBFCDAL1.

Input

IOAI

Formatted and filled out as documented above.

IOAI_INO

Points to storage containing the DEB name.

I/O Area

Formatted as documented above.

Output

IOAI_STATUS

Call status, '' means successful.

IOAI_FDBK0

The minimum length of the I/O area.

IOAI_FDBK1

The number of AREAS in this DEDB.

The I/O Area

0	4	8	C	14	len-4
//					
I/O	offset	data	DEDB	area list	end of data
area	to	length	name	using DBFCDAL1	marker
len	data			control blocks	x'EEEEEEEE'
//					
len:4	4	4	8	variable	4

DEDBINFO Call

Return DEDB information from the DMCB, mapped by DBFCDDI1.

Input

IOAI

Formatted and filled out as documented above.

IOAI_INO

Points to storage containing the DEB name.

I/O Area

Formatted with length in the first word, and 'FFFFFFFF' as an end of I/O area marker.

Output

IOAI_FDBK0

The minimum length of the I/O area.

IOAI_FDBK1

The minimum I/O area for the DEDBSTR call.

IOAI_FDBK2

The minimum I/O area for the AREALIST call.

The I/O Area

0	4	8	C	14	len-4
I/O area len	offset to data	data length	DEDB name	the DEDB info using DBFCDDI1 control block	end of data marker x'EEEEEEEE'
len:4	4	4	8	len(DBFCDDI1)	4

DEDSTR Call

Return a list of segments and segment data for a DEDB with each segment mapped by DBFCDDI1.

Input

IOAI

Formatted and filled out as documented above.

IOAI_INO

Points to storage containing the DEB name.

I/O Area

Formatted with length in the first word, and 'FFFFFFFF' as an end of I/O area marker.

Output

IOAI_STATUS

The minimum length of the I/O area.

IOAI_FDBK0

The minimum I/O area for the DEDBSTR call.

IOAI_FDBK1

The minimum I/O area for the SEGMENTS call.

The I/O Area

0	4	8	C	14	len-4
I/O area len	offset to data	data length	DEDB name	segments in DBFCDDI1 control blocks	end of data marker x'EEEEEEEE'
len:4	4	4	8	variable	4

Fast Path coding considerations

You can use DL/I calls to access Fast Path databases. You can also use two additional calls: FLD and POS. The type of Fast Path database that you are processing determines when you can use each of these calls.

To process MSDBs, you can use these calls:

- For nonterminal-related MSDBs:

FLD

GU and GHU
GN and GHN
REPL

- For terminal-related, fixed MSDBs:

FLD
GU and GHU
GN and GHN
REPL

- For terminal-related, dynamic MSDBs:

DLET
FLD
GU and GHU
GN and GHN
ISRT
REPL

You can use these calls to process a DEDB:

- DEQ
- DLET
- FLD
- GU and GHU
- GN and GHN
- GNP and GHNP
- ISRT
- POS
- REPL
- RLSE

Chapter 22. Writing ODBA application programs

By using the ODBA interface, IMS DB databases can be accessed from environments that are outside the scope of control for IMS, such as Db2 for z/OS stored procedures.

The ODBA interface is not needed within IMS-controlled regions, such as MPRs, BMPs, or IFPs, for calls to locally controlled databases.

The z/OS application programs (hereafter called the ODBA application programs) run in a separate z/OS address space that IMS regards as a separate region from the control region. The separate z/OS address space hereafter is called the z/OS application region.

The ODBA interface gains access to IMS DB through the Database Resource Adapter (DRA). The ODBA application programs (which can access any address space within the z/OS they are running in) gain access to IMS DB databases through the ODBA interface. The following figure illustrates this concept and shows the relationship between the components of this environment.




Figure 69. z/OS application region's connection to IMS DB

One z/OS application region can connect to multiple IMS DBs and multiple z/OS application regions can connect to a single IMS DB. The connection is similar to that of CICS to DBCTL.

General application program flow of ODBA application programs

An ODBA application program issues calls, including DL/I calls, using the AERTDLI interface of an application interface block (AIB).

Several conditions must be met for the AIB call to succeed:

1. If an AIB is not passed in the call, a U261 abend is issued.
2. If the AIB that is passed is not valid, a U476 abend is issued.
3. If the AIB that is passed is not large enough (264 bytes), the AIB return and reason codes are set to X'104' and X'228'.
4. If the AIB that is passed is not on a fullword boundary, the z/OS system will return an abend S201.
5. If there are other internal problems with the call, other return and reason codes are passed back to the z/OS application program. See *IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes* for a complete list of these return and reason codes.

The ODBA application program must link edit with a language module (DFSCDLIO) or this module can be loaded into the z/OS application region. The entry point for DFSCDLIO is AERTDLI.

Restriction: The ODBA interface does not support calls into batch DL/I regions.

The basic program flow for an ODBA application program is:

1. Establish the application execution environment.

The application execution environment must be initialized in the z/OS application region.

To initialize the environment, use either the CIMS INIT call or, if you need to establish connections to multiple IMS systems, the CIMS CONNECT call.

If you use the INIT subfunction of CIMS you can include a startup table ID in the optional AIBRSNM2 field of the AIB to connect to the IMS DB system listed in the startup table. If the AIBRSNM2 field is blank, connect to the IMS DB when you allocate a PSB.

The form of the connection call is:

```
CALL AERTDLI parmcount, CIMS, AIB
```

Where:

CIMS

Is the required call function.

AIB

Has the following fields:

AIBSFUNC

The subfunction is either INIT or CONNECT. This field is mandatory.

AIBRSNM1

An optional field that provides an eye catcher identifier of the application server that is associated with the AIB. This field is 8 bytes.

AIBRSNM2

An optional field for the INIT subfunction in which you can specify an optional 4-byte startup table ID. The ID is optional if the call is issued as preconditioning only. If the ID is given, the z/OS application region connects to the IMS DB specified in the DBCTLID parameter of the selected startup table. The AIBRSNM2 field is not supported with the CONNECT subfunction.

The characteristics of the connection are determined from the DRA startup table. The startup table name is DFSxxxx0, where xxxx is the startup table ID that is used in the CIMS and APSB calls. Each startup table defines a combination of connection attributes, one of which is a subsystem ID of the IMS DB.

If you use the CONNECT subfunction, the calling application program can optionally supply its own connection properties table by specifying the address of the table in an entry in the ODBA data store connection table used by the CONNECT subfunction. The connection properties table is mapped by the DFSPRP macro.

Related Reading: For more information about building a DRA startup table, see *IMS Version 15.3 System Definition*.

2. Allocate a PSB.

The APSB call, introduced for CPIC-driven programs, is used with the ODBA interface to allocate a PSB for the z/OS application region.

Security is checked before the call can succeed. For more information, see "Accessing IMS databases through the ODBA interface" in *IMS Version 15.3 Communications and Connections*.

The APSB call is in the following form:

```
CALL AERTDLI parmcount, APSB, AIB
```

Where:

APSB

Is the required call function.

AIB

Is the name of the application interface block. The fields in the AIB must be filled in:

AIBRSNM1

Is the 8-character PSB name.

AIBRSNM2

Is the 4-byte IMS alias name that is used as the startup table ID.

Several conditions must be met for the allocation request to succeed.

- The PSB must exist and security checking through RACF must succeed.
- An ODBA environment must have been established by either a CIMS INIT call or a CIMS CONNECT call.
- z/OS Resource Recovery Services (RRS) must be active when the APSB call is made.

Multiple PSBs can be active at the same time, which is typical for server environments. No token is specifically provided to identify which PSB is to be used for a given call to a given IMS DB, so the same AIB *must* be used for all calls to the same PSB instance (APSB, DB calls, DPSB). This enables multiple instances of the same PSB to be in use for the same IMS DB at the same time. The parallelism is controlled by the thread count specified in the startup table. The maximum number of threads and dependent regions supported by an IMS DB instance is 999.

3. Perform DB calls.

All DL/I calls, with a few exceptions, are supported through the AIB. The unsupported calls entail message handling (the IOPCB is available only for system calls), CKPT, ROLL, ROLB, and INQY PROGRAM. Alternate destination PCBs cannot be used. Both full-function databases and DEDBs are available.

4. Commit the changes.

Synchronization is performed by issuing the distributed commit calls, SRRCMIT or ATRCMIT, or possibly their rollback forms of SRRBACK or ATRBACK. IMS sync-point calls are not allowed. Commit is effective for all RRS controlled resources in the z/OS task.

5. Deallocate the PSB.

The DPSB call is used when the work unit is complete. In the default case, a commit call must be issued before a DPSB call can be issued. No DL/I call, including system service calls, can be made between the commit and the DPSB call.

The DPSB call is in the following form:

```
CALL AERTDLI parmcount, DPSB, AIB
```

Where:

DPSB

Is the required call function.

AIB

Is the name of the application interface block. The following fields in the AIB must be filled in:

AIBRSNM1

Is the 8-character PSB name.

AIBSFUNC

Is an optional field. Set it to 'PREPbbbb' when you want to deallocate the PSB before initialization of commit processing and when the commit processing is provided from outside the application.

IMS performs phase 1 commit processing and returns control to the requestor, but holds the in-doubt work until RRS (the commit manager) requests full commit processing. An example is in DB2 UDB for z/OS Stored Procedures, where Db2 for z/OS initializes commit processing on behalf of the procedure. See Db2 for z/OS for a discussion of this scenario.

6. Terminate the connection.

The termination call is in the following form:

CALL AERTDLI parmcount, CIMS, AIB

Where:

CIMS

Is the required call function.

AIB

Is the name of the application interface block. The following fields in the AIB must be filled in:

AIBSFUNC

Is a mandatory field whose value is TERM or TALL. Use TERM to sever a single IMS DB connection. Use TALL to sever all connections for this z/OS application region and remove the DRA from the address space.

AIBRSNM1

Is an optional field that provides an eye catcher identifier of the application server associated with the AIB. This field is 8 bytes in length.

AIBRSNM2

When subfunction equals TERM, provides the 4-byte startup table ID used in a previous APSB call. This field is not needed when the subfunction equals TALL.

Server program structure

The commit scope within the z/OS application environment is all the work under the TCB from which the commit request is made to z/OS Resource Recovery Services (RRS). Server environments, therefore, need a separate TCB under which the individual client requests will be managed. Each TCB will map to a PST for thread handling.

The following figure shows an example TCB structure for a server environment.




Figure 70. DRA uses one TCB per thread

Each connection to an IMS DB uses a thread under the TCB. When the APSB call is processed, a context is established and tied to the TCB. At commit time, all contexts for this TCB are committed or aborted by RRS.

Loading DFSCDLIO rather than link editing is attractive when the z/OS application region is a server supporting many clients with many instances of threads connected with the IMS DBs.

Db2 for z/OS stored procedures use of ODBA

Db2 for z/OS stored procedures connecting to ODBA must run in a z/OS Workload Manager-managed (WLM-managed) stored procedures address space.

Db2 for z/OS establishes the ODBA environment by specifying either the INIT subfunction or the CONNECT subfunction of the CIMS call for the stored procedure address space. If the CIMS INIT call is issued, the connection to a specific IMS DB occurs when the APSB call is issued. If the CIMS CONNECT call is used, the connection to one or more IMS DB systems can optionally occur either when the CIMS CONNECT call is issued or when the APSB call is issued.

Each stored procedure running in the stored procedure address space runs under its own TCB that is established by Db2 for z/OS when the stored procedure is initialized. Db2 for z/OS issues the commit call on behalf of the stored procedure when control is returned to Db2 for z/OS. Only the PREP subfunction of the DPSB call should be issued by the stored procedures.

Restriction: If stored procedures are nested under a single WLM stored procedure address space and call IMS ODBA, the ODBA threads will hang.

The following figure illustrates the connection from a Db2 for z/OS stored procedures address space to an IMS DB subsystem. This connection allows DL/I data to be presented through an SQL interface, either locally to this Db2 for z/OS or to DRDA connected Db2 for z/OS databases.




Figure 71. Db2 for z/OS stored procedures connection to IMS DB

The following figure illustrates the general relationships involved with using Db2 for z/OS stored procedures and IMS DB together.




Figure 72. Db2 for z/OS stored procedures relationships

Best practices for Db2 for z/OS stored procedures with ODBA

In a z/OS application runtime environments for Db2 for z/OS stored procedures that use the ODBA subsystem, you can make certain key changes to improve performance, reduce hung threads, cut down on the overhead caused by recycling addressed spaces, and increase stability.

In addition to following existing guidance for designing stored procedures, consider the following recommendations:

- Use the Open Database Manager (ODBM) component of the IMS Common Service Layer (CSL) to protect IMS from unexpected terminations of Db2 for z/OS stored procedures.
- Do not design stored procedures that use parameters that can cause problems with ODBA, such as the Db2 for z/OS ASUTIME value.
- Fast-running stored procedures that do not hold unnecessary database locks or issue more than one APSB call at a time within a DB2 thread improve performance and reduce the frequency of hung threads.
- Adjust the values of IMS system parameters, such as the MINTHRD= and MAXTHRD=, in the IMS DFSPRP macro.
- Configure your system to improve serviceability and activate key diagnostic data collection facilities to quickly collect all of the relevant diagnostic information and reduce the amount of time you need to troubleshoot ODBA problems.

Design best practices for ODBA Db2 for z/OS stored procedures

If you follow certain recommendations when you design Db2 for z/OS stored procedures that use ODBA, the stored procedures are less likely to encounter problems.

Use the Open Database Manager (ODBM) to protect IMS from unexpected termination of Db2 for z/OS stored procedures

ODBM uses the ODBA interface to communicate with IMS. Many of the parameters in the ODBM CSLDCxxx configuration PROCLIB member are the same as those used by ODBA in the DFSxxxx0 startup table that is built from the DFSPRP macro parameters.

Because ODBM uses ODBA, ODBA application servers, such as Db2 for z/OS or WebSphere Application Server for z/OS can be configured to connect to IMS through ODBM instead of through ODBA. Connecting through ODBM prevents a U0113 abend from occurring if a DB2 stored procedure or WebSphere Application Server application program terminates unexpectedly during DL/I processing.

Configuring an ODBA application server to use ODBM does not require any changes to existing application programs that run under the ODBA application server.

To use ODBM, you must also enable the Structured Call Interface (SCI) and Operations Manager (OM) components of CSL. For more information, see [Configuring ODBA application servers to use ODBM \(System Administration\)](#).

Design short-running stored procedures

When you write a stored procedure for ODBA, design it so that the total running time between each APSB call and DPSB call is relatively short – approximately one second or less. This recommendation is similar to the recommendation for MPP stored procedures.

An ODBA stored procedure holds database locks while it is running. Until the stored procedure ends and releases the locks, IMS **/DBR** commands for the database cannot complete and other IMS functionality can also be lost as a consequence. By minimizing the running time between APSB and DPSB calls, you also minimize the possibility of delaying a **/DBR** command and losing IMS function.

Another way to prevent this problem is to stop all ODBA stored procedures before you issue the **/DBR** command.

Do not use the ASUTIME statement in Db2 for z/OS stored procedures that interact with ODBA

The Db2 for z/OS ASUTIME parameter defines a limit on the amount of processor time that a Db2 for z/OS stored procedure is allowed to run before Db2 for z/OS cancels the procedure.

Recommendation: Do not use the ASUTIME statement in Db2 for z/OS stored procedures that interact with ODBA.

When the ASUTIME limit is reached, Db2 for z/OS issues a CIMS FTHD call to ODBA, giving ODBA six seconds before Db2 for z/OS ends the TCB. Depending on the environment, if ODBA is unable to complete preparations before the TCB goes away, threads can hang or other problems can occur.

Use only one APSB call at a time per Db2 for z/OS thread

Use only one APSB at a time per Db2 for z/OS thread. If you have more than one APSB active at a time, various timing problems can occur if the application abends.

You can avoid the timing problems if, after every APSB call, you make a DPSB call before the next APSB call, as shown in the following example:

```
APSB
(Other application calls)
DPSB
APSB
```

Use the same value on the IMS MINTHRD= and MAXTHRD= parameters in the DRA startup table

Frequently creating and tearing down threads while your application program is running can cause timing-related errors.

When the MINTHRD= and MAXTHRD= parameters are equal, ODBA creates the requested number of threads during initialization and does not tear them down until termination. In between, threads are not created or destroyed.

Recommendation: Set MINTHRD= equal to MAXTHRD= in your DRA startup parameters. These parameters are different from the Db2 for z/OS NUMTCB parameter. NUMTCB controls the number of threads that can run simultaneously under a single Db2 for z/OS address space, but there can be multiple Db2 for z/OS address spaces. Therefore, do not assume that the total number of Db2 for z/OS tasks in the system is limited by the NUMTCB setting.

Related tasks

[Configuring ODBA application servers to use ODBM \(System Administration\)](#)

Related reference

[CSLDCxxx member of the IMS PROCLIB data set \(System Definition\)](#)

[DRA startup table \(System Programming APIs\)](#)

Writing Db2 for z/OS stored procedures that use ODBA

When you are writing application programs to run in an ODBA environment, you can avoid errors by following a general pattern of ODBA calls.

Recommendation: Use only one APSB call at a time per Db2 for z/OS thread. If you have more than one APSB active at a time and the application abends, various timing problems can occur.

If your ODBA application includes multiple APSB calls, avoid undetectable deadlocks at the PST level by using one of the following methods:

- Before issuing each subsequent APSB call, issue DPSB and SRRCMIT calls after each set of updates.
- Use single threads in ODBA regions.
- Impose a standard sequence for accessing and updating databases across all ODBA applications.

The call pattern you follow differs depending on when the RRS-sync call is initiated.

- If the RRS-sync call is initiated from within the ODBA application, use the following call pattern:
 - a) Begin with an Allocate PSB (APSB) call to allocate a PSB for the ODBA application.
 - b) After the APSB call, write the main body of your application program, which will include statements such as DLI calls.
 - c) Code the RRS-sync.
 - d) Code the Deallocate PSB (DPSB) call. This deallocates the PSB which was allocated by the APSB call.
- If the RRS-sync call is initiated from outside of the ODBA application, use the following call pattern:
 - a) Begin with an Allocate PSB (APSB) call to allocate a PSB for the ODBA application.
 - b) After the APSB call, write the main body of your application program, which will include statements such as DLI calls.
 - c) Code a DPSB PREP call. This action ends the ability of the application to access the ODBA thread. IMS resources, including database buffers and database locks obtained as part of the DL/I calls, are not released until RRS Sync is performed.
 - d) Return to the application deployer.
 - e) Deploy the DB2 application to update DB2 resources.
 - f) RRS-sync.

Stopping Db2 for z/OS stored procedure threads

You have several options for stopping an ODBA stored procedure thread. Some are better than others due to the potential for hung threads or other availability issues.

The methods in the following list are presented in the order in which you should try them.

1. [“Db2 for z/OS -STOP PROCEDURE command” on page 337](#)
2. [“Cancel the thread in Db2 for z/OS” on page 337](#)
3. [“IMS STOP REGION commands” on page 337](#)
4. [“The REFRESH option of the z/OS MVS command VARY WLM” on page 337](#)
5. [“Recycle IMS” on page 337](#)

Db2 for z/OS -STOP PROCEDURE command

The Db2 for z/OS **-STOP PROCEDURE** command is the best way to stop a thread because it has the fewest potential negative side effects. The **-STOP PROCEDURE** command prevents Db2 for z/OS from accepting SQL CALL statements for one or more stored procedures.

For example, the following command stops stored procedures USERPRC1 and USERPRC3. While the **-STOP PROCEDURE** command is in effect, attempts to execute these stored procedure are queued.

```
-STOP PROCEDURE(USERPRC1,USERPRC3)
```

Cancel the thread in Db2 for z/OS

Use Db2 for z/OS **-DISPLAY THREAD** and **-CANCEL THREAD** commands to display threads and locks held, and then cancel the threads from Db2 for z/OS.

```
-DISPLAY THREAD(*) TYPE(PROC)
```

```
-CANCEL THREAD(token)
```

IMS STOP REGION commands

Use IMS commands to display the hung threads and then use one of the following commands to stop them.

- **/STOP REGION** *reg#* | *job*
- **/STOP REGION** *reg#* | *job* ABDUMP *tran*

The REFRESH option of the z/OS MVS command VARY WLM

You can use the REFRESH option of the z/OS MVS **VARY WLM** command to refresh the WLM environment; that is, it recycles the WLM environment.

Refreshing the WLM environment is problematic for ODBA because it terminates the stored procedures without using the IMS AIB DPSB or TALL calls to terminate the ODBA threads. Normally, when ODBA receives a TERM or TALL request, connections to IMS wait for active threads to return to the DRA before terminating. VARY REFRESH does not allow this, which can result in hung threads.

The following example shows the VARY command syntax.

```
VARY WLM,APPLENV=xxxxx,REFRESH
```

Recycle IMS

As a last resort, you can recycle the IMS address space to stop an ODBA thread if the previously described methods do not work.

Testing an ODBA application program

You should perform a program unit test on your ODBA application program to ensure that the program correctly handles its input data, processing, and output data. The amount and type of testing you do depends on the individual program.

Be aware of your established test procedures before you start to test your program. To begin testing, you need the following items:

- A test JCL statement
- A test database

Always begin testing programs against test-only databases. Do not test programs against production databases. If the program is faulty it might damage or delete critical data.

- Test input data

The input data that you use need not be current, but it should be valid data. You cannot be sure that your output data is valid unless you use valid input data.

The purpose of testing the program is to make sure that the program can correctly handle all the situations that it might encounter. To thoroughly test the program, try to test as many of the paths that the program can take as possible. For example:

Test each path in the program by using input data that forces the program to execute each of its branches. Be sure that your program tests its error routines. Again, use input data that will force the program to test as many error conditions as possible. Test the editing routines your program uses. Give the program as many different data combinations as possible to make sure it correctly edits its input data. The following table lists the tools you can use to test Online (IMSDB), Batch, and BMP programs.

Table 62. Tools you can use for testing your program

Tool	Online (IMS DB)	Batch	BMP
DFSDDLTO	No	Yes ¹	Yes
DL/I image capture program	Yes	Yes	Yes

Note: 1. For call-level programs only. (For a command-level batch program, you can use DL/I image capture program first, to produce calls for DFSDDLTO).

Tracing DL/I calls with image capture to test your ODBA program

The DL/I image capture program (DFSDLTRO) is a trace program that can trace and record DL/I calls issued by batch, BMP, and online (IMS DB environment) programs. You can produce calls for use as input to DFSDDLTO.

You can use the image capture program to:

- Test your program

If the image capture program detects an error in a call it traces, it reproduces as much of the call as possible, although it cannot document where the error occurred, and cannot always reproduce the full SSA.

- Produce input for DFSDDLTO (DL/I test program)

You can use the output produced by the image capture program as input to DFSDDLTO. The image capture program produces status statements, comment statements, call statements, and compare statements for DFSDDLTO. For example, you can use the image capture program with a ODBA application, to produce calls for DFSDDLTO.

- Debug your program

When your program terminates abnormally, you can rerun the program using the image capture program. The image capture program can then reproduce and document the conditions that led to the program failure. You can use the information in the report produced by the image capture program to find and fix the problem.

Using image capture with DFSDDLTO to test your ODBA program

The image capture program produces the following control statements that you can use as input to DFSDDLTO.

- Status statements

When you invoke the image capture program, it produces the status statement. The status statement it produces:

- Sets print options so that DFSDDLTO prints all call trace comments, all DL/I calls, and the results of all comparisons
- Determines the new relative PCB number each time a PCB change occurs while the application program is running
- Comments statement

The image capture program also produces a comments statement when you run it. The comments statements give:

The time and date IMS started the trace
The name of the PSB being traced

The image capture program also produces a comments statement preceding any call in which IMS finds an error.

- Call statements

The image capture program produces a call statement for each DL/I call.

- Compare statements


If you specify COMP on the DLITRACE control statement, the image capture program produces data and PCB comparison statements.

Running image capture online

When you run the image capture program online, the trace output goes to the IMS log data set. To run the image capture program online, you issue the IMS TRACE command from the z/OS console.

If you trace a BMP and you want to use the trace results with DFSDDLTO, the BMP must have exclusive write access to the databases it processes. If the application program does not have exclusive access, the results of DFSDDLTO may differ from the results of the application program.

The following diagram shows TRACE command format:



SET ON/OFF

Turns the trace on or off.

PSB *psbname*

Specifies the name of the PSB you want to trace. You can trace more than one PSB at the same time by issuing a separate TRACE command for each PSB.

COMP/NOCOMP

Specifies whether you want the image capture program to produce data and PCB compare statements to be used with DFSDDLTO.

Retrieving image capture data from the log data set

If the trace output is sent to the IMS log data set, you can retrieve it by using utility DFSERA10 and a DL/I call trace exit routine, DFSERA50.

DFSEREA50 deblocks, formats, and numbers the image capture program records to be retrieved. To use DFSERA50, you must insert a DD statement defining a sequential output data set in the DFSERA10 input stream. The default ddname for this DD statement is TRCPUNCH. The card must specify BLKSIZE=80.

Examples: You can use the following examples of DFSERA10 input control statements in the SYSIN data set to retrieve the image capture program data from the log data set:

- Print all image capture program records:

Column 1	Column 10
OPTION	PRINT OFFSET=5,VALUE=5F,FLDTYP=X

- Print selected image capture program records by PSB name:

Column 1	Column 10
OPTION	PRINT OFFSET=5,VALUE=5F,COND=M
OPTION	PRINT OFFSET=25,FLDTYP=C,FLDLEN=8, VALUE=psbname, COND=E

- Format image capture program records (in a format that can be used as input to DFSDDLTO):

Column 1	Column 10
OPTION	PRINT OFFSET=5,VALUE=5F,COND=M
OPTION	PRINT EXITR=DFSERA50,OFFSET=25,FLDTYP=C VALUE=psbname,FLDLEN=8,DDNAME=OUTDDN,COND=E

The DDNAME= parameter is used to name the DD statement used by DFSERA50. The data set defined on the OUTDDN DD statement is used instead of the default TRCPUNCH DD statement. For this example, the DD appears as:

```
//OUTDDN DD ...,DCB=(BLKSIZE=80),...
```

Requests for monitoring and debugging your ODBA program

To debug your ODBA program, you can issue the statistics (STAT) or log (LOG) request.

You can use the following two requests to help you in debugging your program:

- The statistics (STAT) request retrieves database statistics. STAT can be issued from both call- and command-level programs.
- The log (LOG) request makes it possible for the application program to write a record on the system log. You can issue LOG as a command or call in a batch program; in this case, the record is written to the IMS log. You can issue LOG as a call or command in an online program in the IMS DB environment; in this case, the record is written to the IMS DB log.

What to do when your ODBA program terminates abnormally

Whenever your program terminates abnormally, you can take some actions to simplify the task of finding and fixing the problem. ODBA does not issue any return or reason codes. Most non-terminating errors for ODBA application programs are communicated in AIB return and reason codes. You can record as much information as possible about the circumstances under which the program terminated abnormally. In addition, you can check for certain initialization and execution errors.

Recommended actions after an abnormal termination of an ODBA program

The suggestions given here are some common guidelines on what you should do if your program terminates abnormally.

- Document the error situation to help in investigating and correcting it. Some of the information that can be helpful include:
 - The program's PSB name
 - The call function
 - The terminal ID (online programs only)
 - The contents of the AIB or the PCB
 - The contents of the I/O area when the problem occurred

- If a database request was executing, the SSAs or SEGMENT and WHERE options, if any, the request used
- The date and time of day
- When your program encounters an error, it can pass all the required error information to a standard error routine.
- You can send a message to the system log by issuing a LOG request.

Diagnosing an abnormal termination of an ODBA program

If your program does not run correctly when you are testing it or when it is running, you need to isolate the problem. The problem might be anything from a programming error (for example, an error in the way you coded one of your requests) to a system problem.

You can check for the following errors when your program fails to run, terminates abnormally, or gives incorrect results.

ODBA initialization errors

Before your program receives control, IMS must have correctly loaded and initialized the PSB and DBDs used by your application program. Often, when the problem is in this area, you need a system programmer or DBA (or your equivalent specialist) to fix the problem. One thing you can do is to find out if there have been any recent changes to the DBDs, PSB, and the control blocks that they generate.

ODBA running errors

If you do not have any initialization errors, check the following in your program:

1. The output from the compiler. Make sure that all error messages have been resolved.
2. The output from the binder:
 - Are all external references resolved?
 - Have all necessary modules been included?
 - Was the language interface module correctly included?
3. Your JCL. Is the information that described the files that contain the databases correct? If not, check with your DBA.

Chapter 23. Programming with the IMS support for DRDA

IMS provides an implementation of the Distributed Relational Database Architecture (DRDA) protocol that you can use to write your own IMS Connect TCP/IP client applications.

DRDA is an open architecture that enables communication between applications and database systems on disparate platforms. Details about using the DRDA protocol to perform database access operations are in the open specifications for DRDA. The following information describes only the IMS-specific extensions provided by the IMS support for DRDA.

To use the IMS support for DRDA, you must create the DRDA client driver (DRDA source server). No additional software needs to be installed or configured on the client system. The DRDA target server consists of IMS Connect and the Open Database Manager (ODBM) running with IMS in z/OS.

The IMS support for DRDA includes support for both application-directed transaction demarcation (local) and XA-enabled (global) transactions.

IMS does not support the following DRDA functions:

- Multi-row input
- Client reroute
- Security plugin

The IMS support for DRDA is based on the DRDA Version 4 technical standard. The DRDA specification is documented by the Open Group Consortium at www.opengroup.org.

Server compatibility checking

All communication between a source and target DRDA server begins with initialization and security. In the initialization flow, the DRDA client issues the **EXCSAT** command and an **EXCSATRD** data object is sent back from the DRDA target server.

In the IMS support for DRDA implementation, the **EXCSATRD** reply data object includes a Server Release Level (SRVRLSLV) parameter. The SRVRLSLV parameter is a string that specifies the version number of the distributed database management (DDM) language recognized by the IMS Connect and ODBM server components. This string is used by the client to perform server compatibility checking to ensure that both IMS Connect and ODBM understand any codepoints that the client sends. The DDM version numbering is specific to the IMS support for DRDA. All compatibility checking for the IMS support for DRDA is done based on the SRVRLSLVL parameter.

Important: The SRVRLSLV parameter value sent back from the target server in response to the **EXCSAT** command is OD-ICON 1 OD-ODBM 1.

Updating the source server with the latest maintenance release of IMS without applying the same maintenance release to all your IMS Connect or ODBM installations may cause the source server to be out of synchronization with the target server. To prevent this possibility, the server compatibility check allows the connection to be made only if the IMS support for DRDA target server recognizes the DDM version level that is used by the source server.

How IMS data is mapped to the DRDA protocol

In a database query operation with the IMS support for DRDA, a *row* is defined as the concatenation of an instance of the **aibdbpcbStream** data structure plus all of the requested fields within an IMS hierarchic path. An **aibdbpcbStream** instance is a concatenation of an instance of the **aibStream** data structure followed by an instance of the **dbpcbStream** data structure. The requested fields are represented by the **RTRVFLD** objects sent with an **OPNQRY** command. The concatenation of the **aibdbpcbStream** instance and data fields represents a single row in a query row set.

The IMS support for DRDA supports only flexible blocking, where each query block can be a different size, depending on the size of the row or result set being returned. The specified query block size is used as an initial size, and the query block can expand beyond that size, if necessary, to complete the fetch operation.

In the IMS support for DRDA implementation, data is returned from the DRDA target server in byte stream format, and the client is responsible for data type processing.

Related concepts

[Overview of the CSL Open Database Manager \(System Administration\)](#)

[IMS Connect support for access to IMS DB \(Communications and Connections\)](#)

Related reference

[DRDA DDM command architecture reference \(Application Programming APIs\)](#)

DDM commands for data operations with the IMS support for DRDA

Use the distributed database management (DDM) commands provided by the IMS support for DRDA for singleton and batch data operations.

Before accessing the database, you need to first establish a database connection by issuing an **ACCRDB** command from your DRDA client application and successfully receive an **ACCRDBRM** data object back from the DRDA target server.

After the connection is established, you can issue DDM commands to access data from your DRDA client application.

- To retrieve data, issue an **OPNQRY** command.
- To insert, update, or delete data, issue an **EXCSQLIMM** command.

Data operations can be in singleton or batch operations. Specify the type of data operation by setting the Byte String Data Representation (BYTSTRDR) parameter in the **DLIFUNC** command object that is chained to the DDM command.

The following table shows the DDM commands that the DRDA client issues for data operations with the IMS support for DRDA.

Table 63. DDM commands for data operations with the IMS support for DRDA

Data operation	DDM command	BYTSTRDR parameter value for DLIFUNC command object
Insert data	EXCSQLIMM	ISRT
Retrieve data - DL/I Get Hold Unique	OPNQRY	GHU
Retrieve data - DL/I Get Unique	OPNQRY	GU
Retrieve data - DL/I Get Hold Next	OPNQRY	GHN
Retrieve data - DL/I Get Next	OPNQRY	GN
Retrieve data - DL/I Get Hold Next Within Parent	OPNQRY	GHNP
Retrieve data - DL/I Get Next Within Parent	OPNQRY	GNP
Update data	EXCSQLIMM	REPL
Delete data	EXCSQLIMM	DLET

The following table shows the DDM commands that the DRDA client issues for batch data operations with the IMS support for DRDA.

Table 64. DDM commands for batch data operations with the IMS support for DRDA

Batch data operation	DDM command	BYTSTRDR parameter value for DLIFUNC command object
Retrieve data	OPNQRY	RETRIEVE
Update data	EXCSQLIMM	UPDATE
Delete data	EXCSQLIMM	DELETE

Related reference

[DRDA DDM command architecture reference \(Application Programming APIs\)](#)

Part 3. Application programming for IMS TM

IMS provides support for writing application programs to access IMS transactions.

Chapter 24. Defining application program elements for IMS TM

You can write application programs to communicate with the IMS Transaction Manager using DL/I calls in assembler language, C, COBOL, Java, Pascal, or PL/I.

Formatting DL/I calls for language interfaces

When you use DL/I calls in assembler language, C language, COBOL, Pascal, or PL/I, you must call the DL/I language interface to initiate the functions specified with the DL/I calls.

IMS offers several interfaces for DL/I calls:

- A language-independent interface for any programs that are Language Environment conforming (CEETDLI)
- Language-specific interfaces for all supported languages (xxxTDLI)
- A non-language-specific interface for all supported languages (AIBTDLI)

Java makes use of the all three DL/I language interfaces, but the usage is internal and no calls are necessary to initiate the functions specified with the DL/I calls.

Related concepts

[“IMS solutions for Java development overview” on page 553](#)


You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

Application programming for assembler language


Application programs in assembly language use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

In assembler language programs, all DL/I call parameters that are passed as addresses can be passed in a register, which, if used, must be enclosed in parentheses.


Format



A



B



C



Notes:

¹ Assembler language programs must use either *parmcount* or VL.

Parameters

parmcount

Specifies the address of a 4-byte field in user-defined storage that contains the number of parameters in the parameter list that follows *parmcount*. Assembler language application programs must use either *parmcount* or VL.

function

Specifies the address of a 4-byte field in user-defined storage that contains the call function to be used. The call function must be left-justified and padded with blanks. For example, (GUBb) is a call function.

i/o pcb

Specifies the address of the I/O program communication block (PCB). The I/O PCB address is the first address passed on entry to the application program in the PCB list, given the following circumstances:

- A program executing in DLI or database management batch (DBB) regions where CMPAT=YES is coded on the PSB.
- Any program executing in batch message processing program (BMP), message processing program (MPP), or IMS Fast Path (IFP) regions regardless of the CMPAT= value.

alternate pcb

Specifies the address of the alternate PCB to be used for the call. The PCB address must be one of the PCB addresses passed on entry to the application program in the PCB list.

aib

Specifies the address of the application interface block (AIB) in user-defined storage.

i/o area

Specifies the address of the I/O area in user-defined storage used for the call. The I/O area must be large enough to contain the returned data.

i/o area length

Specifies the address of a 4-byte field in user-defined storage that contains the I/O area length (specified in binary).

area length

Specifies the address of a 4-byte field in user-defined storage that contains the length (specified in binary) of the area immediately following it in the parameter list. Up to seven area length/area pairs can be specified.

area

Specifies the address of the area in user-defined storage to be checkpointed. Up to seven area length/area pairs can be specified.

token

Specifies the address of a 4-byte field in user-defined storage that contains a user token.

options list

Specifies the address of the *options list* in user-defined storage that contains processing options used with the call.

feedback area

Specifies the address of the feedback area in user-defined storage that receives information about options list processing errors.

mod name

Specifies the address of an 8-byte area in user-defined storage that contains the user-defined MOD name used with the call. The *mod name* parameter is used only with MFS.

destination name

Specifies the address of an 8-byte field in user-defined storage that contains the name of the logical terminal or transaction code to which messages resulting from the call are sent.

VL

Signifies the end of the parameter list. Assembler language programs must use either *parmcount* or VL.

Example DL/I call formats

DL/I AIBTDLI interface:

```
CALL AIBTDLI,(function,aib,i/o area),VL
```

DL/I language-specific interface:

```
CALL ASMTDLI,(function,i/o pcb,i/o area),VL
```

Related concepts

[“AIBTDLI interface” on page 226](#)

Use AIBTDLI as the interface between your application program and IMS.

Related reference


[DL/I calls for transaction management \(Application Programming APIs\)](#)

[DL/I calls for IMS TM system services \(Application Programming APIs\)](#)


Application programming for C language

Application programs in C use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.


Format



A




B



C



D



area length

Specifies the name of a fixed-binary (31) variable in user-defined storage that contains the length of the area immediately following it in the parameter list. Up to seven area length/area pairs can be specified.

area

Specifies the name of the pointer variable that contains the address of the structure that defines the user-defined storage to be checkpointed. Up to seven area length/area pairs can be specified.

token

Specifies the name of a character (4) variable in user-defined storage that contains a user token.

options list

Specifies the name of the pointer variable that contains the address of the structure that defines the user-defined storage that contains processing options used with the call.

feedback area

Specifies the name of the pointer variable that contains the address of the structure that defines the user-defined storage that receives information about options list processing errors.

mod name

Specifies the name of a character (8) variable in user-defined storage that contains the user-defined MOD name used with the call. The *mod name* parameter is used only with MFS.

destination name

Specifies the name of a character (8) variable in user-defined storage that contains the name of the logical or terminal transaction code to which messages resulting from the call are sent.

I/O area

In C language, the I/O area can be of any type, including structure or array. The `ceetdli` declarations in `leawi.h` and the `ctdli` declarations in `ims.h` do not have any prototype information, so no type checking of the parameters is done. The I/O area can be auto, static, or allocated (with `malloc` or `calloc`). Give special consideration to C-strings because DL/I does not recognize the C convention of terminating strings with nulls (`'\0'`). Instead of using the `strcpy` and `strcmp` functions, you might want to use the `memcpy` and `memcmp` functions.

Example DL/I call formats**DL/I CEEDTLI interface:**

```
#include <leawi.h>
ceetdli(function,aib,i/o_area)
```

DL/I AIBTDLI interface:

```
int rc;
:
rc = aibtdli(parmcount,function,aib,i/o_area)
```

DL/I language-specific interface:

```
#include <ims.h>
int rc;
:
rc = ctdli(function,i/o_pcb,i/o_area)
```

Related concepts

“AIBTDLI interface” on page 226

Use AIBTDLI as the interface between your application program and IMS.

Related reference

[DL/I calls for transaction management \(Application Programming APIs\)](#)


[DL/I calls for IMS TM system services \(Application Programming APIs\)](#)

Application programming for COBOL

Application programs in COBOL use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.


Format

➤ CALL ➔




➤ . ➔

A



B



C



Parameters

parmcount

Specifies the identifier of a usage binary (4) byte data item in user-defined storage that contains the number of parameters in the parameter list that follows *parmcount*.

function

Specifies the identifier of a usage display (4) byte data item, left-justified, in user-defined storage, which contains the call function to be used. The call function must be padded with blanks. For example, (GUbb) is a call function.

i/o pcb

Specifies the address of the I/O program communication block (PCB). The I/O PCB address is the first address passed on entry to the application program in the PCB list, given the following circumstances:

- A program executing in DLI or database management batch (DBB) regions where CMPAT=YES is coded on the PSB.
- Any program executing in batch message processing program (BMP), message processing program (MPP), or IMS Fast Path (IFP) regions regardless of the CMPAT= value.

alternate pcb

Specifies the identifier of the I/O PCB or alternate PCB group item from the PCB list that is passed to the application program on entry. This identifier is used for the call.

aib

Specifies the identifier of the group item that defines the application interface block (AIB) in user-defined storage.

i/o area

Specifies the identifier of a group item, table, or usage display data item that defines the I/O area to be used for the call. The I/O area must be large enough to contain the returned data.

i/o area length

Specifies the identifier of a usage binary (4) byte data item in user-defined storage that contains the I/O area length.

area length

Specifies the identifier of a usage binary (4) byte data item in user-defined storage that contains the length of the area immediately following it in the parameter list. Up to seven area length/area pairs can be specified.

area

Specifies the identifier of the group item that defines the area to be checkpointed. Up to seven area length/area pairs can be specified.

token

Specifies the identifier of a usage display (4) byte data item that contains a user token.

options list

Specifies the identifier of the group item that defines the user-defined storage that contains processing options used with the call.

feedback area

Specifies the identifier of the group item that defines the user-defined storage that receives information about options list processing errors.

mod name

Specifies the identifier of a usage display (8) byte data item in user-defined storage that contains the user-defined MOD name used with the call.

destination name

Specifies the identifier of a usage display (8) byte data item that contains the name of the logical terminal or transaction code to which messages resulting from the call are sent.

Example DL/I call formats**DL/I CEETDLI interface:**

```
CALL 'CEETDLI' USING function, aib,i/o area.
```

DL/I AIBTDLI interface:

```
CALL 'AIBTDLI' USING function, aib,i/o area.
```

DL/I language-specific interface:

```
CALL 'CBLTDLI' USING function, i/o pcb, i/o area.
```

Related concepts

[“AIBTDLI interface” on page 226](#)

Use AIBTDLI as the interface between your application program and IMS.

Related reference

[DL/I calls for transaction management \(Application Programming APIs\)](#)

[DL/I calls for IMS TM system services \(Application Programming APIs\)](#)

Java application programming for IMS

IMS provides support for developing applications using the Java programming language.

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.


Related concepts

[“IMS solutions for Java development overview” on page 553](#)


You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

Application programming for Pascal


Application programs in Pascal use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

Format


A




B



C



D



Notes:

¹ For AIBTDLI, *parmcount* is required for applications.

Parameters

parmcount

specifies the address of a fixed-binary (31) variable in user-defined storage that contains the number of parameters in the parameter list that follows *parmcount*.

function

Specifies the name of a character (4) variable, left-justified, in user-defined storage, which contains the call function to be used. The call function must be padded with blanks. For example, (Gubb) is a call function.

i/o pcb

Specifies the address of the program communication block (I/O PCB). The I/O PCB address is the first address passed on entry to the application program in the PCB list, given the following circumstances:

- A program executing in DLI or database management batch (DBB) regions where CMPAT=YES is coded on the PSB.
- Any program executing in batch message processing program (BMP), message processing program (MPP), or IMS Fast Path (IFP) regions regardless of the CMPAT= value.

alternate pcb

Specifies the name of a pointer variable that contains the address of the I/O PCB defined in the call procedure statement.

aib

Specifies the name of a pointer variable that contains the address of the structure that defines the application interface block (AIB) in user-defined storage.

i/o area

Specifies the name of a pointer variable to a major structure, array, or character string that defines the I/O area in user-defined storage to be used for the call. The I/O area must be large enough to contain the returned data.

i/o area length

Specifies the name of a fixed-binary (31) variable in user-defined storage that contains the I/O area length.

area length

Specifies the name of a fixed binary (31) variable in user-defined storage that contains the length (specified in binary) of the area immediately following it in the parameter list. Up to seven area length/area pairs can be specified.

area

Specifies the name of a pointer variable that contains the address of the structure that defines the area in user-defined storage to be checkpointed. Up to seven area length/area pairs can be specified.

token

Specifies the name of a character (4) variable in user-defined storage that contains a user token.

options list

Specifies the name of a pointer variable that contains the address of the structure that defines the user-defined storage that contains processing options used with the call.

feedback area

Specifies the name of the pointer variable that contains the address of the structure that defines the user-defined storage that receives information about options list processing errors.

mod name

Specifies the name of a character (8) variable in user-defined storage that contains the user-defined MOD name used with the call.

destination name

Specifies the name of a character (8) variable in user-defined storage that contains the name of the logical terminal or transaction code to which messages resulting from the call are sent.

Example DL/I call formats**DL/I AIBTDLI interface:**

```
AIBTDLI(CONST function,
        VAR aib,
        VAR I/O area);
```

DL/I language-specific interface:

```
PASTDLI(CONST function,
area    VAR I/O PCB
        VAR I/O area);
```

Related concepts

[“AIBTDLI interface” on page 226](#)

Use AIBTDLI as the interface between your application program and IMS.

Related reference


[DL/I calls for transaction management \(Application Programming APIs\)](#)

[DL/I calls for IMS TM system services \(Application Programming APIs\)](#)

Application programming for PL/I


Application programs in PL/I use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

Format




—) — ; —


A



B



C



Parameters

parmcount

Specifies the name of a fixed-binary (31-byte) variable that contains the number of arguments that follow *parmcount*.

function

Specifies the name of a character (4-byte) variable, left justified, blank padded character string that contains the call function to be used. For example, (GUBB) is a call function.

i/o pcb

Specifies the address of the program communication block (I/O PCB). The I/O PCB address is the first address passed on entry to the application program in the PCB list, given the following circumstances:

- A program executing in DLI or DBB regions where CMPAT=YES is coded on the PSB.
- Any program executing in batch message processing program (BMP), message processing program (MPP), or IMS Fast Path (IFP) regions regardless of the CMPAT= value.

alternate pcb

Specifies the structure associated with the I/O PCB or alternate PCB to be used for the call. This structure is based on a PCB address that must be one of the PCB addresses passed on entry to the application program.

aib

Specifies the name of the structure that defines the application interface block (AIB).

i/o area

Specifies the name of the I/O area used for the call. The I/O area must be large enough to contain the returned data.

i/o area length

Specifies the name of a fixed binary (31) variable in user-defined storage that contains the I/O area length (specified in binary).

area length

Specifies the name of a fixed binary (31) variable that contains the length (specified in binary) of the area immediately following it in the parameter list. Up to seven area length/area pairs can be specified.

area

Specifies the name of the area to be checkpointed. Up to seven area length/area pairs can be specified.

token

Specifies the name of a character (4) variable that contains a user token.

options list

Specifies the name of a structure that contains processing options used with the call.

feedback area

Specifies the name of a structure that receives information about options list processing errors.

mod name

Specifies the name of a character (8) variable character string containing the user-defined MOD name used with the call.

destination name

Specifies the name of a character (8) variable character string containing the logical terminal or transaction code to which messages resulting from the call are sent.

Example DL/I call formats

DL/I CEETDLI interface:

```
%INCLUDE CEEIBMAW;  
CALL CEETDLI (function, i/o pcb, i/o area);
```

DL/I AIBTDLI interface:

```
CALL AIBTDLI (parmcount, function, aib, i/o area);
```

DL/I language-specific interface:

```
CALL PLITDLI (parmcount, function, i/o pcb, i/o area);
```

Relationship of calls to PCB types

The following table shows the relationship of DL/I calls to I/O and alternate program communication blocks (PCBs).

The PCB can be specified as a parameter in the call list, or in the AIB, depending on which xxxTDLI interface is used:

Table 65. Call relationship to PCBs and AIBs

Call	I/O PCBs	ALT PCBs
APSB ¹		
AUTH	X	
CHKP (basic)	X	
CHKP (symbolic)	X	
CHNG ²		X
CMD	X	
DPSB ¹		
GCMD	X	
GN	X	
GSCD	X	
GU	X	
INIT	X	
INQY	X	X
ISRT	X	X
LOG	X	
PURG	X	X
ROLB	X	
ROLS	X	
ROLL ¹		
SETO	X	X
SETS	X	

Table 65. Call relationship to PCBs and AIBs (continued)

Call	I/O PCBs	ALT PCBs
SETU	X	
SYNC	X	
XRST	X	

Notes:

1. This call is not associated with a PCB.
2. The alternate PCB used by this call must be modifiable.

Specifying the I/O PCB mask

After your program issues a call with the I/O program communications block (PCB), IMS returns information about the results of the call to the I/O PCB. To determine the results of the call, your program must check the information that IMS returns.

Issuing a system service call requires an I/O PCB. Because the I/O PCB resides outside your program, you must define a mask of the PCB in your program to check the results of IMS calls. The mask must contain the same fields, in the same order, as the I/O PCB. Your program can then refer to the fields in the PCB through the PCB mask.

An I/O PCB contains the fields listed in the following table. The table describes these fields, their lengths, and which environments are applicable for each field.

Table 66. I/O PCB mask

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
Logical terminal name ¹	8	X		X		
Reserved for IMS ²	2	X		X		
Status code ³	2	X	X	X	X	X
8-Byte Local date and time ⁴						
Date	4	X		X		
Time	4	X		X		
Input message sequence number ⁵	4	X		X		
Message output descriptor name ⁶	8	X		X		
Userid ⁷	8	X		X		
Group name ⁸	8	X		X		
12-Byte Time Stamp ⁹						
Date	4	X		X		
Time	6	X		X		
UTC Offset	2	X		X		

Table 66. I/O PCB mask (continued)

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
Userid Indicator ¹⁰	1	X		X		
Reserved for IMS ²	3					

Note:

1. Logical Terminal Name

This field contains the name of the terminal that sent the message. When your program retrieves an input message, IMS places the name of the logical terminal that sent the message in this field. When you want to send a message back to this terminal, you refer to the I/O PCB when you issue the ISRT call, and IMS takes the name of the logical terminal from the I/O PCB as the destination.

2. Reserved for IMS

These fields are reserved.

3. Status Code

IMS places the status code describing the result of the DL/I call in this field. IMS updates the status code after each DL/I call that the program issues. Your program should always test the status code after issuing a DL/I call.

The three status code categories are:

- Successful status codes or status codes with exceptional but valid conditions. This category does not contain errors. If the call was completely successful, this field contains blanks. Many of the codes in this category are for information only. For example, a QC status code means that no more messages exist in the message queue for the program. When your program receives this status code, it should terminate.
- Programming errors. The errors in this category are usually ones that you can correct. For example, an AD status code indicates an invalid function code.
- I/O or system errors.

For the second and third categories, your program should have an error routine that prints information about the last call that was issued program termination. Most installations have a standard error routine that all application programs at the installation use.

4. Local Date and Time

The current local date and time are in the prefix of all input messages except those originating from non-message-driven BMPs. The local date is a packed-decimal, right-aligned date, in the format yyddd. The local time is a packed-decimal time in the format hhmmss.t. The current local date and time indicate when IMS received the entire message and enqueued it as input for the program, rather than the time that the application program received the message. To obtain the application processing time, you must use the time facility of the programming language you are using.

For a conversation, for an input message originating from a program or for a message received using Multiple System Coupling (MSC), the time and date indicate when the original message was received from the terminal.

Note: Be careful when comparing the local date and time in the I/O PCB with the current time returned by the operating system. The I/O PCB date and time may not be consistent with the current time. It may even be greater than the current time for the following reasons:

- The time stamp in the I/O PCB is the local time that the message was received by IMS. If the local time was changed after the message arrived, it is possible for the current time to appear to be earlier than the I/O PCB time. This effect would be likely to occur in the hour immediately after the fall time change, when the clock is set back by one hour.

- The time stamp in the I/O PCB is derived from an internal IMS time stamp stored with the message. This internal time stamp is in Coordinated Universal Time (UTC), and contains the time zone offset that was in effect at the time the message was enqueued. This time zone offset is added to the UTC time to obtain the local time that is placed in the I/O PCB. However, the time zone offset that is stored is only fifteen minutes. If the real time zone offset was not an integer multiple of fifteen minutes, the local time passed back in the I/O PCB will differ from the actual time by plus or minus 7.5 minutes. This could cause the I/O PCB time to be later than the current time. See *IMS Version 15.3 Operations and Automation* for further explanation.

Concerns about the value in the local time stamp in the I/O PCB can be reduced by using the extended time stamp introduced in IMS V6. The system administrator can choose the format of the extended time stamp to be either local time or UTC. In some situations, it may be advantageous for the application to request the time in UTC from the operating system and compare it to the UTC form of the extended time stamp. This is an option available in installations where there is no ETR to keep the IMS UTC offset in sync with the z/OS UTC offset over changes in local time.

5. Input Message Sequence Number

The input message sequence number is in the prefix of all input messages except those originating from non-message-driven BMPs. This field contains the sequence number IMS assigned to the input message. The number is binary. IMS assigns sequence numbers by physical terminal, which are continuous since the time of the most recent IMS startup.

6. Message Output Descriptor Name

You only use this field when you use MFS. When you issue a GU call with a message output descriptor (MOD), IMS places its name in this area. If your program encounters an error, it can change the format of the screen and send an error message to the terminal by using this field. To do this, the program must change the MOD name by including the MOD name parameter on an ISRT or PURG call.

Although MFS does not support APPC, LU 6.2 programs 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 is to be formatted.

Related reading: For more information on the MOD name and the LTERM interface, see *IMS Version 15.3 Communications and Connections*.

7. Userid

The use of this field is connected with RACF signon security. If signon is not active in the system, this field contains blanks.

If signon is active in the system, the field contains one of the following:

- The user's identification from the source terminal.
- The LTERM name of the source terminal if signon is not active for that terminal.
- The authorization ID. For batch-oriented BMPs, the authorization ID is dependent on the value specified for the BMPUSID= keyword in the DFSDCxxx PROCLIB member:
 - If BMPUSID=USERID is specified, the value from the USER= keyword on the JOB statement is used.
 - If USER= is not specified on the JOB statement, the program's PSB name is used.
 - If BMPUSID=PSBNAME is specified, or if BMPUSID= is not specified at all, the program's PSB name is used. If PSBNAME is not defined to RACF, the Userid of the current Address Space will be used; this will be the Home Dependent Region one, or the Control Region one if LSO=Y or if PARDLI=1 has been specified for the BMP. Userid of the current Address Space will be used also if DFSBSEX0 has returned RC08.

Related Reading: For more information about authorizing resource use in a dependent region, see *IMS Version 15.3 System Administration*.

8. Group Name

The group name, which is used by DB2 to provide security for SQL calls, is created through IMS transactions.

Three instances that apply to the group name are:

- If you use RACF and signon on your IMS system, the RACROUTE SAF (extract) call returns an eight-character group name.
- If you use your own security package on your IMS system, the RACROUTE SAF call returns any eight-character name from the package and treats it as a group name. If the RACROUTE SAF call returns a return code of 4 or 8, a group name was not returned, and IMS blanks out the group name field.
- If you use LU 6.2, the transaction header can contain a group name.

Related reading: See *IMS Version 15.3 Communications and Connections* for more information on LU 6.2.

9. 12-Byte Time Stamp

This field contains the current date and time fields, but in the IMS internal packed-decimal format. The time stamp has the following parts:

Date

yyyydddf

This packed-decimal date contains the year (yyyy), day of the year (ddd), and a valid packed-decimal + sign such as (f).

Time

hhmmssthmiju

This packed-decimal time consists of hours, minutes, and seconds (hhmmss) and fractions of the second to the microsecond (thmiju). No packed-decimal sign is affixed to this part of the time stamp.

UTC Offset

aqq\$

The packed-decimal UTC offset is prefixed by 4 bits of attributes (a). If the 4th bit of (a) is 0, the time stamp is UTC; otherwise, the time stamp is local time. The control region parameter, TSR= (U/L), specified in the DFSPBxxx PROCLIB member, controls the representation of the time stamp with respect to local time versus UTC time.

The offset value (qq\$) is the number of quarter hours of offset to be added to UTC or local time to convert to local or UTC time respectively.

The offset sign (\$) follows the convention for a packed-decimal plus or minus sign.

Field 4 on the I/O PCB Mask always contains the local date and time. For a description of field 4, see the notes for the previous table.

Related reading: For a more detailed description of the internal packed-decimal time-format, see *IMS Version 15.3 System Utilities*.

10. Userid Indicator

The Userid Indicator is provided in the I/O PCB and in the response to the INQY call. The Userid Indicator contains one of the following:

- U - The user's identification from the source terminal during signon
- L - The LTERM name of the source terminal if signon is not active
- P - The PSBNAME of the source BMP or transaction
- O - Other name

The value contained in the Userid Indicator field indicates the contents of the userid field.

Related concepts

[“Results of a message: I/O PCB” on page 381](#)

After your program issues a call, IMS TM returns information about the results of the call in the I/O PCB. To find out about the results of the call, your application program must check the information that IMS TM returns to the I/O PCB.

Specifying the alternate PCB mask

An alternate program communication block (PCB) mask contains three fields.

The following table describes these fields, the field length, and in which environment the field applies.

Table 67. Alternate PCB mask

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
Logical terminal name ¹	8 bytes	X		X		
Reserved for IMS ²	2 bytes	X		X		
Status code ³	2 bytes	X		X		

Note:

1. Logical Terminal Name

This field contains the name of the logical terminal, LU 6.2 descriptor or the transaction code to which you want to send the message.

Related reading: For more information on LU 6.2, see *IMS Version 15.3 Communications and Connections*.

2. Reserved for IMS

This 2-byte field is reserved.

3. Status Code

This field contains the 2-byte status code that describes the results of the call that used this PCB most recently.

Related concepts

[“Sending messages to other terminals and programs” on page 391](#)

When an application program processes a message from a terminal, it usually sends the response to the terminal that sent the input message. But sometimes you might want to send output messages to a terminal other than the originating terminal, or to other terminals in addition to the originating terminal. You might also want to send messages to other application programs.

Specifying the AIB mask

The AIB is used by your program to communicate with IMS, when your application does not have a program communication block (PCB) address or the call function does not use a PCB.

The application program can use the returned PCB address, when available, to inspect the status code in the PCB and to obtain any other information needed by the application program. The AIB mask enables your program to interpret the control block defined. The AIB structure must be defined in working storage, on a fullword boundary, and initialized according to the order and byte length of the fields as shown in the following table. For a description of the contents of each field, see the list that follows the table.

Table 68. AIB fields

Descriptor	Byte Length	DB/DC	DBCTL	DCCTL	DB Batch	TM Batch
AIB identifier ¹	8	X	X	X	X	X
DFSAIB allocated length ²	4	X	X	X	X	X
Subfunction code ³	8	X	X	X	X	X
Resource name 1 ⁴	8	X	X	X	X	X
Resource name 2 ⁵	8	X	X	X	X	X
Reserved ⁶	8					
Maximum output area length ⁷	4	X	X	X	X	X
Output area length used ⁸	4	X	X	X	X	X
Resource field ⁹	4	X	X	X	X	X
Optional area length ¹⁰	4	X	X	X	X	X
Reserved ¹¹	4	X	X	X	X	X
Return code ¹²	4	X	X	X	X	X
Reason code ¹³	4	X	X	X	X	X
Error code extension ¹⁴	4	X	X	X		
Resource address 1 ¹⁵	4	X	X	X	X	X
Resource address 2 ¹⁶	4	X	X	X	X	X
Resource address 3 ¹⁷	4	X	X	X	X	X
User defined token ¹⁸	16	X	X	X	X	X
Return token ¹⁹	8	X	X		X	
Reserved ²⁰	16					

AIB Identifier (AIBID)

This 8-byte field contains the AIB identifier. You must initialize AIBID in your application program to the value DFSAIBbb before you issue DL/I calls. This field is required. When the call is completed, the information returned in this field is unchanged.

DFSAIB Allocated Length (AIBLEN)

This field contains the actual 4-byte length of the AIB as defined by your program. You must initialize AIBLEN in your application program before you issue DL/I calls. The minimum length required is 128 bytes. When the call is completed, the information returned in this field is unchanged. This field is required.

Subfunction Code (AIBSFUNC)

This 8-byte field contains the subfunction code for those calls that use a subfunction. You must initialize AIBSFUNC in your application program before you issue DL/I calls. When the call is completed, the information returned in this field is unchanged.

Resource Name (AIBRSNM1)

This 8-byte field contains the name of a resource. The resource varies depending on the call. You must initialize AIBRSNM1 in your application program before you issue DL/I calls. When the call is complete, the information returned in this field is unchanged. This field is required.

For PCB related calls where the AIB is used to pass the PCB name instead of passing the PCB address in the call list, this field contains the PCB name. The PCB name for the I/O PCB is IOPCBbb. The PCB name for other types of PCBs is defined in the PCBNAME= parameter in PSBGEN.

Resource Name 2 (AIBRSNM2)

This 8-byte field contains the name of a resource. The resource varies depending on the call. You must initialize AIBRSNM2 in your application program before you issue DL/I calls.

Reserved

This 8-byte field is reserved.

Maximum Output Area Length (AIBOALEN)

This 4-byte field contains the length of the output area in bytes that was specified in the call list. You must initialize AIBOALEN in your application program for all calls that return data to the output area. When the call is completed, the information returned in this area is unchanged.

Used Output Area Length (AIBOAUSE)

This 4-byte field contains the length of the data returned by IMS for all calls that return data to the output area. When the call is completed this field contains the length of the I/O area used for this call.

Resource Field (AIBRSFLD)

This 4-byte field contains a resource. The resource varies depending on the call. You must initialize AIBRSFLD in your application program before you issue DL/I calls.

Optional Area Length (AIBOPLEN)

This 4-byte field contains a resource. The resource varies depending on the call. You must initialize AIBOPLEN in your application program before you issue DL/I calls.

Reserved

This 4-byte field is reserved.

Return code (AIBRETRN)

When the call is completed, this 4-byte field contains the return code.

Reason Code (AIBREASN)

When the call is completed, this 4-byte field contains the reason code.

Error Code Extension (AIBERRXT)

This 4-byte field contains additional error information depending on the return code in AIBRETRN and the reason code in AIBREASN.

Resource Address 1 (AIBRSA1)

When the call is completed, this 4-byte field contains call-specific information. For PCB related calls where the AIB is used to pass the PCB name instead of passing the PCB address in the call list, this field returns the PCB address.

Resource Address 2 (AIBRSA2)

When the call is completed, this 4-byte field contains call-specific information.

Resource Address 3 (AIBRSA3)

When the call is completed, this 4-byte field contains call-specific information.

User Defined Token (AIBUTKN)

This 16-byte field contains a user defined token. The token varies depending on the call.

AIB return token (AIBRTKN)

AIB return token. This 8-byte field contains a token returned by a DL/I call. The usage is specific to the DL/I call for which the token was returned.

Reserved

This 16-byte field is reserved.

Related reference

[DL/I calls for transaction management \(Application Programming APIs\)](#)

[DL/I calls for database management \(Application Programming APIs\)](#)

Specifying the I/O areas

Use an I/O area to pass segments between the application program and IMS.

What the I/O area contains depends on the type of call you are issuing:

- When you retrieve a segment, IMS places the segment you requested in the I/O area.
- When you add a new segment, you first build the new segment in the I/O area.
- Before modifying a segment, your program must first retrieve it. When you retrieve the segment, IMS places the segment in an I/O area.

The format of the record segments you pass between your program and IMS can be fixed length or variable length. Only one difference is important to the application program: a message segment containing a 2-byte length field (or 4 bytes for the PLITDLI interface) at the beginning of the data area of the segment.

The I/O area for IMS calls must be large enough to hold the largest segment your program retrieves from or sends to IMS.

If your program issues any Get or ISRT calls that use the D command code, the I/O area must be large enough to hold the largest path of segments that the program retrieves or inserts.

AIBTDLI interface

Use AIBTDLI as the interface between your application program and IMS.

Restriction: No fields in the application interface block (AIB) can be used by the application program except as defined by IMS.

When you use the AIBTDLI interface, you specify the program communication block (PCB) requested for the call by placing the PCB name (as defined by PSBGEN) in the resource name field of the AIB. You do not specify the PCB address. Because the AIB contains the PCB name, your application program can refer to the PCB name rather than the PCB address. Your application program does not need to know the relative PCB position in the PCB list. At completion of the call, the AIB returns the PCB address that corresponds to the PCB name passed by the application program.

The names of DB PCBs and alternate PCBs are defined by the user during PSBGEN. All I/O PCBs are generated with the PCB name *bbb*. For a generated program specification block (GPSB), the I/O PCB is generated with the PCB name *IOPCBbbb*, and the modifiable alternate PCB is generated with the PCB name *TPPCB1bb*.

The ability to pass the PCB name means that you do not need to know the relative PCB number in the PCB list. In addition, the AIBTDLI interface enables your application program to make calls on PCBs that do not reside in the PCB list. The **LIST=** keyword, which is defined in the PCB macro during PSBGEN, controls whether the PCB is included in the PCB list.

The AIB resides in user-defined storage that is passed to IMS for DL/I calls that use the AIBTDLI interface. Upon call completion, IMS updates the AIB. Allocate at least 128 bytes of storage for the AIB.

Related concepts

[“PCB masks for GSAM databases” on page 283](#)

For the most part, you process GSAM databases in the same way that you process other IMS databases. You use calls that are very similar to DL/I calls to communicate your requests. GSAM describes the results of those calls in a GSAM DB PCB.

Related reference

[“Application programming for PL/I” on page 360](#)

Application programs in PL/I use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

[“Application programming for Pascal” on page 357](#)

Application programs in Pascal use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

[“Application programming for C language” on page 352](#)

Application programs in C use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

[“Application programming for assembler language” on page 349](#)

Application programs in assembly language use the following format, parameters, and DL/I calls to communicate with the IMS Transaction Manager.

[“Assembler language application programming” on page 195](#)

Application programs in assembly language use the following format, parameters, and DL/I calls to communicate with IMS databases.

Specifying language-specific entry points

IMS gives control to an application program through an entry point. Use the correct format for coding entry statements in assembler language, C language, COBOL, Java, Pascal, and PL/I.

Your entry point must refer to the program communication blocks (PCBs) in the order in which they are defined in the PSB.

IMS passes the PCB pointers to a PL/I program differently than it passes them to an assembler language, C language, COBOL, Java, or Pascal program. In addition, Pascal requires that IMS pass an integer before passing the PCB pointers. IMS uses the LANG keyword or the PSBGEN statement of PSBGEN to determine the type of program to which it is passing control. Therefore, you must be sure that the language specified during PSBGEN is consistent with the language of the program.

Application interfaces that use the AIB structure (AIBTDLI or CEETDLI) use the PCB name rather than the PCB structure and do not require the PCB list to be passed at entry to the application program.

When you code each DL/I call, you must provide the PCB you want to use for that call. For all IMS TM application programs, the list of PCBs the program can access is passed to the program at its entry point.

Assembler language

You can use any name for the entry statement to an assembler language DL/I program. When IMS passes control to the application program, register 1 contains the address of a variable-length fullword parameter list. Each word in the list contains the address of a PCB. Save the parameter list address before you overwrite the contents of register 1. IMS sets the high-order byte of the last fullword in the list to X'80' to indicate the end of the list. Use standard z/OS linkage conventions with forward and backward chaining.

C language

When IMS passes control to your program, it passes the addresses, in the form of pointers, for each of the PCBs your program uses. The usual argc and argv arguments are not available to a program invoked by IMS. The IMS parameter list is made accessible by using the __pcblist macro. You can directly reference the PCBs by __pcblist[0], __pcblist[1], or you can define macros to give these more meaningful names. I/O PCBs must be cast to get the proper type:

```
(IO_PCB_TYPE *) (__pcblist[0])
```

The entry statement for a C language program is the main statement.

```
#pragma runopts(env(IMS),plist(IMS))
#include <ims.h>

main()
{
  :
  :
}
```

The `env` option specifies the operating environment in which your C language program is to run. For example, if your C language program is invoked under IMS and uses IMS facilities, specify `env(IMS)`. The `plist` option specifies the format of the invocation parameters received by your C language program when it is invoked. When your program is invoked by a system support services program such as IMS, the format of the parameters passed to your main program must be converted into the C language format: `argv`, `argc`, and `envp`. To do this conversion, you must specify the format of the parameter list received by your C language program. The `ims.h` include file contains declarations for PCB masks.

You can finish program execution in three ways:

- End the main procedure without an explicit return statement.
- Execute a return statement from main.
- Execute an `exit` or an `abort` call from anywhere, or alternately issue a `longjmp` back to main, and then do a normal return.

One C language program can pass control to another by using the `system` function. The normal rules for passing parameters apply. For example, when using the `system` function, the `argc` and `argv` arguments can be used to pass information. The initial `__pcblist` is made available to the invoked program.

COBOL

The procedure statement must refer to the I/O PCB first, then to any alternate PCB it uses, and finally to the DB PCBs it uses. The alternate PCBs and DB PCBs must be listed in the order in which they are defined in the PSB.

```
Procedure division using the PCB-NAME-1 [...,PCB-NAME-N]
```

On previous versions of IMS, the `using` keyword might be coded on the entry statement to reference PCBs. However, IMS continues to accept such coding on the entry statement.

Recommendation: Use the procedure statement rather than the entry statement to reference the PCBs.

Java

The entry into a Java program is the `main (String args[])` method. When IMS passes control to your main program, it does not pass the PCB list, and the `main (String args[])` method is always called with a `String` array with 0 elements. Java programs are different in this regard as they use the AIBTDLI interface which addresses the PCB by a PCB name or PCB label coded in the PSB rather than a PCB pointer, to communicate with IMS. For more information, see: [“AIBTDLI interface” on page 226](#)

```
The entry into a Java program is the main() method.
public class HelloWorldJmp {
    public static void main(String[] args) {
        ...
        System.out.println("Hello. I am running in a JMP");
    }
}
```

Pascal

The entry point must be declared as a `REENTRANT` procedure. When IMS passes control to a Pascal procedure, the first address in the parameter list is reserved for Pascal's use and the other addresses are the PCBs the program uses. The PCB types must be defined before this entry statement. The IMS interface routine `PASTDLI` must be declared with the `GENERIC` directive.

```
procedure ANYNAME(var SAVE: INTEGER;
                  var pcb1-name: pcb1-name-type[;
                  ...
                  var pcbn-name: pcbn-name-type]); REENTRANT;
procedure ANYNAME;
(* Any local declarations *)
procedure PASTDLI; GENERIC;
begin
```



```
(* Code for ANYNAME *)  
end;
```

PL/I

The entry statement can be any valid PL/I name and must appear as the first executable statement in the program. When IMS passes control to your program, it passes the addresses of each of the PCBs your program uses in the form of pointers. When you code the entry statement, make sure you code the parameters of this statement as pointers to the PCBs, and not the PCB names.

```
anyname: PROCEDURE (pcb1_ptr [, ..., pcbn_ptr]) OPTIONS (MAIN);  
:  
RETURN;
```

CCETDLI and AIBTDLI interface considerations

The CCETDLI considerations are:

- For PL/I programs, the CEETDLI entry point is defined in the CEEIBMMAW include file. Alternatively, you can declare it yourself. But it must be declared as an assembler language entry (DCL CEETDLI OPTIONS(ASM);).
- For C language applications, you must specify env(IMS) and plist(IMS); these specifications enable the application to accept the PCB list of arguments. The CEETDLI function is defined in <leawi.h>; the CTDLI function is defined in <ims.h>.

The AIBTDLI considerations are:

- When using the AIBTDLI interface for C/MVS, COBOL, or PL/I language applications, the language run-time options for suppressing abend interception (that is, NOSPIE and NOSTAE) must be specified. However, for Language Environment-conforming applications, the NOSPIE and NOSTAE restriction is removed.
- The AIBTDLI entry point for PL/I programs must be declared as an assembler language entry (DCL AIBTDLI OPTIONS(ASM);).
- For C language applications, you must specify env(IMS) and plist(IMS); these specifications enable the application to accept the PCB list of arguments.

Program communication block (PCB) lists

Use the correct format of program communication block (PCB) lists and generated program specification block (GPSB) PCB lists in your application program.

PCB list format

This is the format of a PCB:

```
[IOPCB]  
[Alternate PCB ... Alternate PCB]  
[DB PCB ... DB PCB]  
[GSAM PCB ... GSAM PCB]
```

Each PSB must contain at least one PCB. An I/O PCB or alternate PCB is required for transaction management calls, and an I/O PCB is required for most system service calls. DB PCBs for DL/I databases are used only with the IMS Database Manager, but can be present even though your program is running under DCCTL or TM Batch. (A DB PCB can be a full-function PCB, a DEDB PCB, or an MSDB PCB.) GSAM PCBs can be used with DCCTL or TM batch.

Format of a GPSB PCB list

A generated program specification block (GPSB) has the following format:

[IOPCB]
[Alternate PCB]

A GPSB contains only an I/O PCB and one modifiable alternate PCB. It can be used by all transaction management application programs, and permits access to the PCBs specified without the need for PSBGEN.

The PCBs in a GPSB have predefined PCB names. The name of the I/O PCB is IOPCBbb. The name of the alternate PCB is TPPCB1bb.

PCB summary

I/O PCBs and alternate PCBs can be used in various types of application programs.

TM Batch Programs

Alternate PCBs are always included in the list of PCBs supplied to the program by IMS TM. The I/O PCB is always present in the PCB list regardless of the CMPAT options specified in PSBGEN.

BMPs, MPPs, and IFPs

The I/O PCB is always present in the PCB list and is always the first address in the list, regardless of the CMPAT options specified in the PSB. The PCB list always contains the address of the I/O PCB followed by the addresses of any alternate PCBs, followed by the addresses of the DB PCBs.

Language environments

IBM Language Environment provides the strategic execution environment for running your application programs written in one or more high level languages.

It provides not only language-specific run-time support, but also cross-language run-time services for your applications, such as support for initialization, termination, message handling, condition handling, storage management, and National Language Support. Many of Language Environment's services are accessible explicitly through a set of Language Environment interfaces that are common across programming languages; these services are accessible from any Language Environment-conforming program.

Language Environment-conforming programs can be compiled with the following compilers:

- IBM C++/MVS
- IBM COBOL
- IBM PL/I

By default, the Language Environment infrastructure uses the 31-bit addressing mode. By specifying JVM=64, the Language Environment is changed to use the 64-bit addressing mode. By specifying JVM=3164, the dependent region will initialize a 31-bit Language Environment and a secondary 64-bit Language Environment to support 31-bit COBOL and 64-bit Java interoperability.

Language Environment supports C, C++, and assembly language interoperability in a 64-bit addressing mode, but does not support COBOL and PL/I interoperability in a 64-bit addressing mode. Do not switch to JVM=64 if your Java application invokes either COBOL or PL/I. If the regions are switched to use JVM=64 inadvertently, and incompatible interoperable applications are running, the application might receive system or user abends.

The CEETDLI interface to IMS

The language-independent CEETDLI interface to IMS is provided by Language Environment. It is the only IMS interface that supports the advanced error handling capabilities provided by Language Environment. The CEETDLI interface supports the same functionality as the other IMS application interfaces, and it has the following characteristics:

- The parmcount variable is optional.
- Length fields are 2 bytes long.

- Direct pointers are used.

Related reading: For more information about Language Environment, see *z/OS Language Environment Programming Guide*.

LANG= option on PSBGEN for PL/I compatibility

For IMS PL/I applications running in a compatibility mode that uses the PLICALLA entry point, you must specify LANG=PLI on the PSBGEN. Your other option is to change the entry point and add SYSTEM(IMS) to the EXEC PARM of the compile step so that you can specify LANG=blank or LANG=PLI on the PSBGEN. The following table summarizes when you can use LANG=blank and LANG=PLI.

Table 69. Using LANG= option in a Language Environment for PL/I compatibility

Compile exec statement is PARM=(...,SYSTEM(IMS)...	Entry point name is PLICALLA	Valid LANG= value
Yes	Yes	LANG=PLI
Yes	No	LANG=blank or LANG=PLI
No	No	Note: Not valid for IMS PL/I applications
No	Yes	LANG=PLI

PLICALLA is only valid for PL/I compatibility with Language Environment. If a PL/I application using PLICALLA entry at bind time is bound using Language Environment with the PLICALLA entry, the bind will work; however, you must specify LANG=PLI in the PSB. If the application is re-compiled using PL/I for z/OS & VM Version 1 Release 1 or later, and then bound using Language Environment Version 1 Release 2 or later, the bind will fail. You must remove the PLICALLA entry statement from the bind.

Special DL/I situations for IMS TM programming

Special considerations during application programming for IMS Transaction Manager include mixed-language programming, using the extended addressing capabilities of z/OS, COBOL compiler options for preloaded programs, and considerations for the DCCTL environment.

Mixed-language programming

When an application program uses the Language Environment language-independent interface, CEETDLI, IMS does not need to know the language of the calling program.

When the application program calls IMS in a language-dependent interface, IMS determines the language of the calling program according to the entry name specified in the CALL statement:

- CALL CBLTDLI indicates the program is in COBOL.
- CALL PLITDLI indicates the program is in PL/I.
- CALL PASTDLI indicates the program is in Pascal.
- ctdli(...) indicates the program is in C language.
- CALL ASMTDLI indicates the program is in assembler language.

If a PL/I program calls an assembler language subroutine and the assembler language subroutine makes DL/I calls by using CALL ASMTDLI, the assembler language subroutine should use the assembler language calling convention, not the PL/I convention.

In this situation, where the I/O area uses the LLZZ format, the LL is a halfword, not the fullword that is used for PLITDLI.

Using Language Environment routine retention

If you run programs in an IMS TM dependent region that requires Language Environment (such as an IMS message processing region), you can improve performance if you use Language Environment library routine retention along with the existing PREINIT feature of IMS TM.

Related reading: For more information about Language Environment, see *z/OS Language Environment Programming Guide*.

Using the extended addressing capabilities of z/OS

The two modes in z/OS with extended addressing capabilities are: the addressing mode (AMODE) and the residency mode (RMODE).

IMS places no constraints on the RMODE and AMODE of an application program. The program can reside in the extended virtual storage area. The parameters referenced in the call can also be in the extended virtual storage area.

Related reading: For more information about Language Environment, see *z/OS MVS Programming: Assembler Services Guide*.

COBOL compiler options for preloaded programs

If you compile your COBOL program with the COBOL for z/OS & VM compiler and preload it, you must use the COBOL compiler option RENT. Alternatively, if you compile your COBOL program with the VS COBOL II compiler and preload it, you must use the COBOL compiler options RES and RENT.

DCCTL

In a DCCTL environment, the application can only reference the address of an I/O PCB, alternate PCB, or GSAM PCB. An application program can use a PSB that contains PCBs referencing databases; however, these PCBs cannot be used during processing. Entry statements for COBOL, PL/I, C, and Pascal must refer to all PCBs included in the PSB, including PCBs which you might not be able to process, as PCBs must be included in the order in which they are listed in the PSB. This includes all PCBs prior to the last referenced PCB and can include DB PCBs. If you used a GSAM PCB, all PCBs ahead of it must be referenced.

Chapter 25. Message processing with IMS TM

IMS Transaction Manager application programs can be written in assembler language, C language, COBOL, Pascal, and PL/I to process messages.

How your program processes messages

To retrieve and send messages, an IMS TM application program issues calls to IMS TM. When your program issues a call to retrieve a message, IMS TM places the input message in the I/O area you name in the call. Before you issue a call to send a message, you must build the output message in an I/O area in your program.

Message types

An operator at a terminal can send four kinds of messages to IMS TM.

The destination of an IMS TM message identifies which kind of message is being sent:

- **Another terminal.** A logical terminal name in the first 8 bytes means that this is a message switch destined for another terminal. For a user at a logical terminal to send a message to another logical terminal, the user enters the name of the receiving logical terminal followed by the message. The IMS TM control region routes the message to the specified logical terminal. This kind of message does not result in the scheduling of any activity in a message processing program (MPP).
- **An application program.** A transaction code in the first 8 bytes means that the message is destined for an application program. IMS TM uses a transaction code to identify MPPs and transaction-oriented batch message processing programs (BMPs). To use a particular application program to process requests, the user enters the transaction code for that application program.
- **IMS TM.** A "/" (slash) in the first byte means that the message is a command destined for IMS TM.
- **Message switch service.** A system service DFSAPPC request is destined for the message switch service.

An application program can send three kinds of messages:

- **Commands.** A "/" in the first byte of the message text means that the message is a command for IMS TM. Programmers design applications to issue commands when they want a program to perform tasks that an operator at a terminal usually performs. This is called automated operator interface (AOI) and is described in *IMS Version 15.3 Communications and Connections* and *IMS Version 15.3 Operations and Automation*.

Use the CMD call to issue commands. Do not use the ISRT call for issuing commands, because a message created with ISRT can contain a slash in the first byte without being a command.

- Messages to logical terminals by specifying a logical terminal name.
- Program-to-program switches using a transaction code.

The messages that your program receives and sends are made up of segments. Use a GU call to retrieve the first segment of a new message, and use GN calls to retrieve the remaining segments of the message. The following figure shows three messages. Message A contains one segment, message B contains two segments, and message C contains three segments.




Figure 73. Message segments

To retrieve message A, you only have to issue a GU call. To retrieve messages B and C, issue one GU call to retrieve the first segment, then a GN call for each remaining segment. This assumes that you know how many segments each message contains. If you do not know this, issue GN calls until IMS TM returns a QD status code, indicating that all of the segments for that message have been retrieved.

If you inadvertently issues a GU call after retrieving the first segment of the multi-segment messages, IMS TM returns a QC status code. This status indicates that no more messages are present, without your program retrieving the additional segments associated with the message. Data would have been lost without any indication that it happened.

Input message format and contents

The input message that an application program receives from a terminal or another program is in most case containing these fields: the length field, the ZZ field, the transaction code field, and the text field.

However, if the input message is the result of a Change (CHNG) call that sets the destination of a modifiable alternate PCB to the logical terminal, LU 6.2 descriptor, or transaction code that you specify, then the transaction code field could be absent.

The tables that follow show the message input layouts. The input message field names are in the first row of each table. The number below each field name is the length in bytes that has been defined for that field. The following table shows the format of an input message for the AIBTDLI, ASMTDLI, CBLTDLI, CEETDLI, CTDLI, and PASTDLI interfaces. The message is slightly different for the PLITDLI interface.

Table 70. Input message format

Field Name	Field Length
LL	2
ZZ	2
TRANCODE	8
Text	Variable

Table 71. Input message format for the PLITDLI interface

Field Name	Field Length
LLLL	4
ZZ	2
TRANCODE	8
Text	Variable

The contents of the input message fields are:

LL or LLLL

The length field contains the length of the input message segment in binary, including LL (or LLLL) and ZZ. IMS TM supplies this number in the length field when you retrieve the input message.

For the AIBTDLI, ASMTDLI, CBLTDLI, CEETDLI, CTDLI, and PASTDLI interfaces, define the LL field as 2 bytes long.

For the PLITDLI interface, define the LLLL field as 4 bytes long. The value in the LLLL field is the input message length minus 2 bytes. For example, if the text is 12 bytes, then the fullword LLLL contains a value of 24 bytes. This value is the total of LLLL (4 bytes) + ZZ (2 bytes) + TRANCODE (8 bytes) + text (12 bytes) - 2 bytes.

ZZ

The ZZ field is a 2-byte field that is reserved for IMS TM. Your program does not modify this field.

TRANCODE

The TRANCODE is the transaction code for the incoming message.

Text

This field contains the message text sent from the terminal to the application program. The first segment of a message can also contain the transaction code associated with the program in the beginning of the text portion of the message. Input messages do not have to include the transaction code, but you can provide it for consistency.

The text field's contents in the input message and the formatting of the contents when your program receives the message depends on the editing routine your program uses.

Output message format and contents

The format of the output message that you build to send back to a terminal or to another program is similar to the format of the input message, but the fields contain different information.

Output messages contain four fields: the length field, the Z1 field, the Z2 field, and the text field. The following tables show the message output layouts. The output message field names are in the first row of each table. The number below each field name is the length in bytes that has been defined for that field. The following table shows the format of an output message for AIBTDLI, ASMTDLI, CBLTDLI, CEETDLI, CTDLI, and PASTDLI interfaces. The format for PLITDLI is slightly different.

Table 72. Output message format

Field Name	Field Length
LL	2
Z1	1
Z2	1
Text	Variable

Table 73. Output message format for PLITDLI

Field Name	Field Length
LLLL	4
Z1	1
Z2	1
Text	Variable

The contents of the output message fields are:

LL or LLLL

The field length contains the length of the message in binary, including the LL (or LLLL), Z1, and Z2 fields. For output message segments, supply this length when you are ready to send the message segment.

For the AIBTDLI, ASMTDLI, CBLTDLI, CEETDLI, CTDLI, and PASTDLI interfaces, the LL field must be 2 bytes long. For the PLITDLI interface, the LLLL field must be 4 bytes long and contain the length of the message segment, minus 2 bytes.

Z1

The Z1 field is a 1-byte field that must contain binary zeros. It is reserved for IMS TM.

Z2

The Z2 field is a 1-byte field that can contain special device-dependent instructions (such as instructions to ring the alarm bell, instructions to disconnect a switched line, or paging instructions) or device-dependent information (such as information about structured field data or bypassing MFS).

If you do not use any of these instructions, the Z2 field must contain binary zeros. For MFS, this field contains the number of the option that is being used for this message.

Text

The text portion of the message segment contains the data that you want to send to the logical terminal or to an application program. (Text messages are typically EBCDIC characters.) The length of the text depends on the data that you want to send.

When a message is processed

A program's response to a message will depend on the type of message the program receives. A transaction code associates a request for information from a terminal with the application program that can process and respond to that request. IMS TM schedules an MPP when there are messages to be processed that contain the transaction code associated with that MPP.

Example: Suppose you have an MPP that processes the transaction code "INVINQ" for inventory inquiry. The MPP receives a request from a user at a terminal for information on the inventory of parts. When the user enters the transaction code for that application program, IMS TM schedules the application program that can process the request.

When you enter INVINQ and one or more part numbers, the MPP sends your program the quantity of each part on hand and the quantity on order.

When you enter INVINQ at the terminal, IMS TM puts the message on the message queue for the MPP that processes INVINQ. Then, after IMS TM has scheduled the MPP, the MPP issues GU and GN calls to retrieve the message. To retrieve the messages from LTERM1, the application program issues a GU for the first segment of a message, then issues GN calls until IMS TM returns a QD status code. This means that the program has retrieved all of the segments of that message. The program then processes the request, and sends the output message to the queue for your logical terminal. (The logical terminal name is in the I/O PCB.) When the MPP sends the output message, IMS TM sends it to the queue for that logical terminal, and the message goes to the physical terminal. The following figure shows the flow of a message between the terminal and the MPP.




Figure 74. Transaction message flow

The following example shows the calls you use, the status codes, and what the input and output for the inventory inquiry would look like. To show you how to use GU and GN to retrieve messages, and how you insert multiple-segment messages, this example shows messages containing three segments. If input and output messages in this example were single segment messages, the program would issue only a GU to retrieve the entire message, and only one ISRT to send the message.

The message formats shown are examples; not all messages are in this format. When the program receives the input message in the I/O area, the first field of each segment contains the length of that segment. This is the LL field in the figure. For clarity, the figure shows this length in decimal; in the input message, however, it is in binary. The second field (ZZ) is reserved for IMS TM; it is 2 bytes long. The text of the message follows the reserved 2 bytes. The first message segment contains the transaction code in the 8 bytes following the ZZ field. These are the first 8 bytes of the text portion of the message.

The format of the output messages is the same. You do not need to include the name of the logical terminal, because it is in the first 8 bytes of the I/O PCB.

PART, QTY, and ON ORDER in the example are headings. These are values that you can define as constants that you want to appear on the terminal screen. To include headings in MFS output messages, define them as literals.




Figure 75. Inventory inquiry MPP example

Results of a message: I/O PCB

After your program issues a call, IMS TM returns information about the results of the call in the I/O PCB. To find out about the results of the call, your application program must check the information that IMS TM returns to the I/O PCB.

When your application program retrieves a message, IMS TM returns this information about the message to the I/O PCB:

- The name of the terminal that sent the message.
- A 2-character status code describing the results of the call. If the program receives a status code of QC after issuing a call to retrieve a message, no more messages are available for the program to process.
- The current date, time, and sequence number for the message.
- The user ID of the person at the terminal or the transaction code for the program that sent the message.

Because the I/O PCB resides in storage outside of your program, you define a mask of the PCB in your program based at this address to check the results of IMS TM calls. The mask contains the same fields in the same order as the I/O PCB.

Related reference

“Specifying the I/O PCB mask” on page 363

After your program issues a call with the I/O program communications block (PCB), IMS returns information about the results of the call to the I/O PCB. To determine the results of the call, your program must check the information that IMS returns.

How IMS TM edits messages

When an application program passes messages to and from a terminal, IMS TM edits the messages before the program receives the message from the terminal and before the terminal receives the message from the application program.

IMS TM gives you many choices about how you want your messages to appear both on the terminal screen and in the program's I/O area. You need to know which editing routines have been specified for your program and how they affect your programming.

The three editing routines available to non-LU 6.2 terminals in IMS TM are:

Basic Edit

Performs basic edit functions if you do not use MFS and if the message does not originate at an LU 6.1 device. You must provide control characters for some formatting functions.

Intersystem Communication (ISC) Edit

Provides the default edit for messages that originate from an LU 6.1 device. You can enter binary data in addition to text.

Message Format Service (MFS)

Formats messages through control blocks. You define the way the messages look with the control blocks.

For LU 6.2 devices, use the LU 6.2 Edit exit routine to edit input and output messages.

Related reading: For more information on LU 6.2, see *IMS Version 15.3 Communications and Connections*. For more information on LU 6.2 Edit exit routine, see *IMS Version 15.3 Exit Routines*.

Printing output messages

To print output messages, you must provide the horizontal and vertical control characters that are necessary to format your output messages.

To print your output at a printer terminal, include these control characters where necessary within the text of the message:

X'05'

Skip to the tab stop, but stay on the same line.

X'15'

Start a new line at the left margin.

X'25'

Skip to a new line, but stay at the same place horizontally.

If you want to skip multiple lines, you can start a new line (X'15'), then skip as many lines as necessary (X'25').

Using Basic Edit

If you do not use MFS or an LU 6.1 device, IMS TM does some editing automatically. The editing IMS TM does to the first message segment is different from the editing IMS TM does for subsequent message segments.

See *IMS Version 15.3 Communications and Connections* for a complete description of Basic Edit.

Editing input messages

When IMS TM receives the first segment of an input message for your application program, IMS TM:

- Removes leading and trailing control characters.
- Removes leading blanks.
- Removes backspaces (from a printer terminal).
- Translates to uppercase, if this is specified with the EDIT=UC specification on the system definition TRANSACT macro.

If the message segment contains a password, IMS TM edits the segment by:

- Removing the password and inserting a blank in place of the password.
- Removing the password if the first character of the text is a blank. IMS TM does not insert the blank.
- Left-justifying the text of the segment.

For subsequent input message segments, IMS TM does not remove leading blanks from the text of the message. The other formatting features are the same.

Editing output messages

For output messages, Basic Edit:

- Changes nongraphic characters in the output message before the data goes to the output 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.

Using Intersystem Communication Edit

Intersystem Communication (ISC) Edit is the default edit for messages from LU 6.1 devices. It is not valid for any other device types. One advantage of using ISC edit is that IMS TM does not edit the text of a message, allowing you to enter binary data.

Editing input messages

The editing IMS TM does to input messages depends on whether the Function Management (FM) header contains the SNA-defined primary resource name (PRN) parameter. In either case, IMS TM removes the FM header before the input message is received by the application program.

If the FM header does not contain the PRN parameter:

- IMS TM removes leading control characters and blanks when it receives the first segment of an input message for your application program.
- If the message segment contains a password, IMS TM removes the password and inserts a blank where the password was.
- IMS TM does not edit the text of the message (the data following the password).

If the FM header contains the PRN parameter:

- The PRN is treated as the transaction code and is received by your application program as the first field in the message segment.
- The message segment is not edited by IMS TM.

Editing output messages

ISC edit does not edit output messages.

Using Message Format Service

Format the messages that you send to MPP using the Message Format Service (MFS). You define the format in control blocks.

The MFS control blocks indicate to IMS TM how you want your input and output messages arranged:

- For input messages, MFS control blocks define how the message that the terminal sends to your MPP is arranged in the I/O area.
- For output messages, MFS control blocks define how the message that your MPP sends to the terminal is arranged on the screen or at the printer. You can also define words or other data that appear on the screen (headings, for example) but do not appear in the program's I/O area. This data, called a literal, can be a field in the output message from the application program or a field in the input message from the terminal.

Terminals and MFS

Whether your program uses MFS depends on the types of terminals and secondary logical units (SLUs) your network uses. You can bypass MFS formatting of an output message for a 3270 device or for SLU Type 2 devices. When MFS is bypassed, you construct the entire 3270 data stream from within your program.

Restriction: MFS cannot be used with LU 6.2 devices (APPC).

Related reading: For more information on LU 6.2 and APPC, see *IMS Version 15.3 Communications and Connections*.

Using MFS involves high-level design decisions that are separate from the tasks of application design and application programming; many installations that use MFS have a specialist who designs MFS screens and message formats for all applications that use MFS.

MFS makes it possible for an MPP to communicate with different types of terminals without having to change the way it reads and builds messages. When the MPP receives a message from a terminal, the message's format in the MPP I/O area depends on the MFS options specified and not on what kind of terminal sent it. MFS shields the MPP from the physical device that is sending the message in the same way that a DB PCB shields the program from what the data in the database actually looks like and how it is stored.

MFS input message formats

You define a message to MFS in fields just as you would define fields within a database segment.

When you define the fields that make up a message segment, you give MFS information such as:

- The field length
- The fill character used when the length of the input data is less than the length defined for the field
- Whether the data in the field is left-justified or right-justified
- If the field is truncated, whether it is truncated on the left or right

The order and length of these fields within the message segment depends on the MFS option that your program is using. You specify the MFS option in the MID. The decision of which option to use for an application program is based on:

- How complex the input data is
- How much the input data varies
- The language the application program is written in
- The complexity of the application program
- Performance factors

The Z2 field in MFS messages contains the MFS formatting option being used to format the messages to and from your program. If something is wrong in the way that IMS TM returns the messages to your I/O area, and you suspect that the problem might be with the MFS option used, you can check this field to see if IMS TM is using the correct option. A X'00' in this field means that MFS did not format the message at all.

One way to understand how each of the MFS options formats your input and output messages is to look at examples of each option.

Example: Suppose that you have defined the four message segments shown in the following table. Each of the segments contains a 2-byte length field and a 2-byte ZZ field. The first segment contains the transaction code that the person at the terminal entered to invoke the application program. The number of bytes defined for each field appears below the name of the field in the figure.

When you use the PLITDLI interface, you must define the length field as a binary fullword, LLLL. When you use the AIBTDLI, ASMTDLI, CBLTDLI, CEETDLI, CTDLI, or PASTDLI interfaces, you must define the length field as a halfword, LL. The value provided by the PL/I application program must represent the actual segment length minus 2 bytes. For example, if the output text is 10 bytes, then the value of the fullword LLLL is 14 and is the sum of the length of LLLL (4 bytes - 2 bytes) + Z1 (1 byte) + Z2 (1 byte) + TEXT (10 bytes).

Table 74. Four-segment message

Segment Number	Field Name	Field Length	Field Value
1	LL	2	0027
	ZZ	2	XXXX
	TRANCODE	8	YYYY
	Text	5	PATIENT#
	Text	10	NAME
2	LL	2	0054
	ZZ	2	XXXX
	Text	50	ADDRESAF
3	LL	2	0016
	ZZ	2	XXXX
	Text	6	CHARGES
	Text	6	PAYMENTS
4	LL	2	0024
	ZZ	2	XXXX
	Text	10	TREATMENT
	Text	10	DOCTOR

For these examples, assume that:

- The transaction code is defined in the MID as a literal.
- All of the fields are left-justified.
- The fill character is defined as a blank. When the length of the data in a field is less than the length that has been defined for that field, MFS pads the field with fill characters. Fill characters can be:
 - Blanks
 - An EBCDIC character
 - An EBCDIC graphic character
 - A null, specified as X'3F'

When you specify that the fill character is to be a null, MFS compresses the field to the length of the data if that length is less than the field length.

The fields for segment 4 of the message in the previous table are arranged on the terminal screen in the format shown in the following figure.

Example: Assume the person enters the name of a patient, and the charges and payments associated with that patient.

```

PATIENT#:          NAME:  MC ROSS
ADDRESAF:
CHARGES: 106.50    PAYMENTS:  90.00
TREATMENT:
DOCTOR:

```

Figure 76. Terminal screen for MFS example

MFS provides three options for message formatting.

MFS option 1

Use this option when the program receives and transmits most of the fields in the message segments. The way that option 1 formats messages depends on whether you have defined a null as the fill character for any of the fields in the segment.

If none of the fields in the message were defined as having a fill character of null:

- The program receives all the segments in the message.
- Each segment is the length that was specified for it in the MID.
- Each segment contains all its fields.
- Each field contains data, data and fill characters, or all fill characters.

The following table shows the Option 1 Format of segments received by the application program.

Table 75. MFS option 1 message format

Segment Number	Field Name	Field Length	Field Value
1	LL	2	0027
	Z1	1	XX
	Z2	1	01
	TRANCODE	8	YYYY
	Text	5	blanks
	Text	10	MCROSSbbbb
2	LL	2	0054
	Z1	1	XX
	Z2	1	01
	Text	50	blanks
3	LL	2	0016
	Z1	1	XX
	Z2	1	01
	Text	6	010650
	Text	6	009000

Table 75. MFS option 1 message format (continued)

Segment Number	Field Name	Field Length	Field Value
4	LL	2	0024
	Z1	1	XX
	Z2	1	01
	Text	10	blanks
	Text	10	blanks

The message format for option 1 output messages is the same as the input message format. The program builds output messages in an I/O area in the format shown for segment 4 in the previous figure. The program can truncate or omit fields in one of two ways:

- Inserting a short segment
- Placing a null character in the field

If one or more of the fields are defined as having a null fill character, the message is different. In this case, the message has these characteristics:

- If a field has been defined as having a fill character of null and the terminal offers not data, the field is eliminated from the message segment.
- If all of the fields in a segment have a null fill character and none of the fields contains any literals, the segment is eliminated from the message.
- If only some of the fields in a segment have a null fill character, any field containing nulls is eliminated from the segment. The relative positions of the fields remaining within the segments are changed.
- When the length of the data that is received from the originating terminal is less than the length that is been defined for the field, the field is truncated to the length of the data.

MFS option 2

Use this option when the program processes multisegment messages where most of the fields are transmitted but some of the segments are omitted. Option 2 formats messages in the same way that option 1 does, unless the segment contains no input data from the terminal after IMS TM has removed the literals. If this is true, and if no additional segments in the message contain input data from the terminal, IMS TM ends the message. The last segment that the program receives is the last segment that contains input data from the terminal.

Sometimes a segment that does not have any input data from the terminal is followed by segments that do contain input data from the terminal. When this happens, MFS gives the program the length field and the Z fields for the segment, followed by a 1-byte field containing X'3F'. This indicates to the program that this is a null segment.

If the message segments shown in [Table 74 on page 385](#) are formatted by option 2, they appear in the format shown in the table below.

Table 76. MFS option 2 message format

Segment Number	Field Name	Field Length	Field Value
1	LL	2	0027
	Z1	1	XX
	Z2	1	02
	TRANCODE	8	YYYY
	Text	5	blanks
	Text	10	MCROSSbbbb
2	LL	2	0005
	Z1	1	XX
	Z2	1	02
	Text	1	X'3F'
3	LL	2	0016
	Z1	1	XX
	Z2	1	02
	Text	6	010650
	Text	6	009000

Segment 2 in the previous table contains only a X'3F' because that segment is null, but Segment 3 contains data. This message does not contain a segment 4 because it is null.

MFS option 3

Use this option when the program receives and transmits only a few of the fields within a segment. When you use option 3, the program receives only those fields that have been received from the terminal. The program receives only segments that contain fields received from the originating terminal. Segments and fields can be of variable length if you have defined option 3 as having a null fill character.

A segment in an option 3 message is identified by its relative segment number—in other words, what position in the message it occupies. The fields within a segment are identified by their offset count within the segment.

Example: The NAME field in segment 1 is (MCROSSbbbb). The value 17 is the sum of the lengths of the fields preceding the NAME field and includes an 8-byte transaction code and a 5-byte field of blanks. It does not include the 2-byte relative segment number field (field A in the following table), the 2-byte length field (field B), or the 2-byte relative offset field (field C).

Option 3 messages do not contain literals defined in the MID. This means that the transaction code is removed from the message, except during a conversation. If the transaction that the program is processing is a conversational transaction, the transaction code is not removed from the message. The transaction code still appears in the scratchpad area (SPA).

Each segment the program receives contains the relative number of this segment in the message (field A in the following table). In addition, each data field within the segment is preceded by two fields:

- A 2-byte length field (B). Including the length field itself, the 2-byte relative field offset, and the data in the field.
- A 2-byte relative field offset (C), giving the field's position in the segment as defined in the MID.

These two fields are followed by the data field. MFS includes these fields for each field that is returned to the application program.

If the message segments shown in Table 74 on page 385 are formatted by option 3, they appear in the format shown in the following table. The notes for the tables explain the letters A, B, C, and D, which are in the first row of segment 1 and segment 3.

Table 77. MFS option 3 message format

Segment Number	Field Name	Field Length	Field Value
1	LL	2	0020
	Z1	1	XX
	Z2	1	03
	A	2	0001
	B	2	0014
	C	2	0017
	D	10	MCROSSbbbb
2	LL	2	0000
	Z1	1	XX
	Z2	1	03
	A	2	0003
	B	2	0010
	C	2	0004
	D	6	010650
	B	2	0010
	C	2	0010
	D	6	009000

Notes to the previous table:

- The fields marked A contain the relative segment number. This number gives the segment's position within the message.
- The fields marked B contain the field length. This length is the sum of the lengths of B field (2 bytes) + C field (2 bytes) + D field (the length of the data).
- The fields marked C contain the relative field offset. This gives each field's position within the segment.
- The fields marked D contain the data from the terminal. In this example, the fill character was defined as blank, so the data field is always its defined length. IMS TM does not truncate it. If you define the fill character as null, the lengths of the data fields can differ from the lengths defined for them in the segment. With a null fill character, if the length of the data from the terminal is less than the length defined for the field, IMS TM truncates the field to the length of the data. Using a null fill with option 3 reduces the space required for the message even further.

MFS output message formats

The output message format is used to define what segments and fields MFS will receive from the application program.

If using option 1 or option 2, the output message format is the same as it is for input messages. Present all fields and segments to MFS. You can present null segments. All fields in output messages are fixed length and fixed position. Output messages do not contain option numbers.

Option 3 output messages are similar to input messages, except that they do not contain option numbers. The program submits the fields as required in their segments with the position information.

Using LU 6.2 User Edit exit routine (optional)

This exit routine edits input and output messages from LU 6.2 devices when the implicit application program interface support is used.

If it is not provided, then messages are presented without modification. IMS does not invoke the exit for CPI-C driven transactions because IMS does not participate in the data flows when the application program uses the CPI directly.

The LU 6.2 User Edit exit routine is called once for each message segment or inbound control flow. You can call the exit routine for data messages and use it to:

- Examine the contents of a message segment.
- Change the contents of a message segment.
- Expand or compact the contents of a message segment.
- Discard a message segment and process subsequent segments, if any.
- Use the Deallocate_Aband command to end the conversation.

For more information on LU 6.2 User Edit exit routine, see *IMS Version 15.3 Communications and Connections* and *IMS Version 15.3 Operations and Automation*.

Message processing considerations for DB2

For the most part, the message processing function of a dependent region that accesses DB2 databases is similar to that of a dependent region that accesses only DL/I databases.

The method each program uses to retrieve and send messages and back out database changes is the same. The differences are:

- DL/I statements are coded differently from SQL (structured query language) statements.
- When an IMS TM application program receives control from IMS TM, IMS has already acquired the resources the program is able to access. IMS TM schedules the program, although some of the databases are not available. DB2 does not allocate resources for the program until the program issues its first SQL statement. If DB2 cannot allocate the resources your program needs, your program can optionally receive an initialization error when it issues its first SQL call.
- When an application issues a successful checkpoint call or a successful message GU call, DB2 closes any cursors that the program is using. This means that your program should issue its OPEN CURSOR statement after a checkpoint call or a message GU.

IMS TM and DB2 work together to keep data integrity in these ways:

- When your program reaches a commit point, IMS TM makes any changes that the program has made to DL/I databases permanent, releases output messages for their destinations, and notifies DB2 that the program has reached a commit point. DB2 then makes permanent any changes that the program has made to DB2 databases.
- When your program terminates abnormally or issues one of the IMS TM rollback calls (ROLB, ROLS without a token, or ROLL), IMS TM cancels any output messages your program has produced, backs out changes your program has made to DL/I databases since the last commit point, and notifies DB2. DB2 backs out the changes that the program has made to DB2 databases since the last commit point.

Through the Automated Operator Interface (AOI), IMS TM application programs can issue DB2 commands and IMS TM commands. To issue DB2 commands, the program issues the IMS TM /SSR command followed by the DB2 command. The output of the /SSR command is routed to the master terminal operator (MTO).

Sending messages to other terminals and programs

When an application program processes a message from a terminal, it usually sends the response to the terminal that sent the input message. But sometimes you might want to send output messages to a terminal other than the originating terminal, or to other terminals in addition to the originating terminal. You might also want to send messages to other application programs.

When you use an alternate PCB:

- If you want to send output messages to one alternate destination, define the alternate PCB for that destination.
- If you want to send output messages to more than one alternate destination, and you want to be able to change the destination of the alternate PCB, define the alternate PCB as modifiable during program specification block (PSB) generation. Then, before you issue the ISRT call, you issue a CHNG call to set the destination of the alternate modifiable PCB for the destination program or terminal.

The *express alternate PCB* is a special kind of alternate PCB that is defined during PSB generation, by specifying EXPRESS=YES.

When you use an express alternate PCB, messages you send using that PCB are sent to their final destinations immediately. Messages sent with other PCBs are sent to temporary destinations until the program reaches a commit point. Messages sent with express PCBs are sent if the program subsequently terminates abnormally, or issues one of the rollback calls: ROLL, ROLB, or ROLS. Using an express alternate PCB in this kind of situation is a way to ensure that the program can notify the person at the terminal, even if abnormal termination occurs. For all PCBs, when a program abnormally terminates or issues a ROLL, ROLB, or ROLS call, messages inserted but not made available for transmission are cancelled, while messages made available for transmission are never cancelled.

For a nonexpress PCB, the message is not made available for transmission to its destination until the program reaches a commit point. The commit point occurs when the program terminates, issues a CHKP call, or requests the next input message and the transaction has been defined with MODE=SNGL.

For an express PCB, when IMS TM knows that it has the complete message, it makes the message available for transmission to the destination. In addition to occurring at a commit point, this also occurs when the application program issues a PURG call using that PCB or requests the next input message.

A PSBGEN can also specify an alternate PCB as an alternate response PCB defined during PSB generation.

- If you want to send a message to an LU 6.2 device, you can specify the LU 6.2 descriptor name that is associated with that device. IMS internally performs the uppercase translation of the destination name (CNT or SMB).

Related reference

[“Specifying the alternate PCB mask” on page 367](#)

An alternate program communication block (PCB) mask contains three fields.

Sending messages to other terminals

To reply to a different terminal, also use the ISRT call, but use an alternate program communication block (PCB) instead of the TP PCB.

Just as the TP PCB represents the terminal that sent the message, an alternate PCB represents the terminal to which you want to send the message.

Single alternate terminal

If you are going to send messages to only one alternate terminal, you can define the alternate PCB for that terminal during PSB generation. When you define an alternate PCB for a particular destination, you cannot change that destination during program execution. Each time you issue an ISRT call that references that PCB, the message goes to the logical terminal whose name was specified for the alternate PCB. To send a message to that terminal, place one message segment at a time in the I/O area, and issue an ISRT call referring to the alternate PCB, instead of the TP PCB.

Several alternate terminals

To send messages to several terminals, you can define the alternate PCB as modifiable during PSB generation. Therefore, the alternate PCB represents more than one alternate terminal. You can change the destination while your program is running.

Before you can set or change the destination of an alternate PCB, you must indicate to IMS TM that the message you have been building so far with that PCB is finished. To do this, issue a PURG call.

PURG allows you to send multiple output messages while processing one input message. When you do not use PURG, IMS TM groups message segments into a message and sends them when the program issues a GU for a new message, terminates, or reaches a commit point. A PURG call tells IMS TM that the message built against this TP PCB or alternate PCB (by issuing one ISRT call per message segment) is complete. IMS TM collects the message segments that you have inserted into one PCB as one message and sends it to the destination represented by the alternate PCB you have referenced.

A PURG call that does not contain the address of an I/O area indicates to IMS TM that this message is complete. If you include an I/O area in the call, PURG acts as an ISRT call as well. IMS TM treats the data in the I/O area as the first segment of a new message. When you include an I/O area on a PURG call, you can also include a MOD name to change the format of the screen for this message. Although specifying the MOD name is optional, when you use it, you can specify it only once per message or in only the first ISRT or PURG that begins the message.

To set the destination of a modifiable alternate PCB during program execution, you use a CHNG call. When you issue the CHNG call you supply the name of the logical terminal to which you want to send the message. The alternate PCB you use then remains set with that destination until you do one of the following:

- Issue another CHNG call to reset the destination.
- Issue another GU to the message queue to start processing a new message. In this case, the name still appears in the alternate PCB, even though it is no longer valid.
- Terminate your program. When you do this, IMS TM resets the destination to blanks.

The first 8 bytes of the alternate PCB contain the name of the logical terminal to which you want to send the message.

When you issue a CHNG call, give IMS TM the address of the alternate PCB you are using and the destination name you want set for that alternate PCB.

When you use the PURG call, you give IMS TM only the address of the alternate PCB. IMS TM sends the message you have built using that PCB.

To indicate an error situation, you can send a message by issuing an ISRT call followed by a PURG call against an express PCB. These calls send the message to its final destination immediately.

Example: The program could go through these steps:

1. The program issues a GU call (and GN calls, if necessary) to retrieve an input message.
2. While processing the message, the program encounters an abnormal situation.
3. The program issues a PURG call to indicate to IMS TM the start of a new message.
4. The program issues a CHNG call to set the destination of an express PCB to the name of the originating logical terminal. The program can get this name from the first 8 bytes of the I/O PCB.

5. The program issues ISRT calls as necessary to send message segments. The ISRT calls reference the express PCB.
6. The program issues a PURG call referencing the express PCB. IMS TM then sends the message to its final destination.
7. The program can then terminate abnormally, or it can issue a ROLL, ROLB, or ROLS call to back out its database updates and cancel the output messages it has created since the last commit point.

If your output messages contained three segments, and you used the PURG call to indicate the end of a message (and not to send the next message segment), you could use this call sequence:

```
CHNG ALTPCB1, LTERMA
ISRT ALTPCB1, SEG1
ISRT ALTPCB1, SEG2
ISRT ALTPCB1, SEG3
PURG ALTPCB1
CHNG ALTPCB1, LTERMB
ISRT ALTPCB1, SEG4
ISRT ALTPCB1, SEG5
ISRT ALTPCB1, SEG6
```

Sending messages to other IMS application programs

A program-to-program switch occurs when an IMS application running in an IMS dependent region sends a message to another IMS application running in an IMS dependent region.

To send a request to another IMS transaction and receive the response in the same unit of work, use the DL/I ICAL request. For more information, see [Synchronous program switch requests \(Communications and Connections\)](#).

You can issue a program-to-program switch to send and receive messages with any of the following types of IMS applications:

- message processing program (MPP)
- batch message processing (BMP) program
- Java message processing (JMP) program
- Java batch processing (JBP) program

To send a message to another online program, use an alternate program communication block (PCB) in a similar way as when sending messages to alternate terminals. If you send messages to only one application program, then you can define the alternate PCB with the transaction code for that application program during PSB generation. If you send messages to more than one application program, you can define the alternate PCB as modifiable.

If you use an alternate modifiable PCB, IMS TM makes a security check when you issue the CHNG call to set the destination of the alternate modifiable PCB. The terminal that enters the transaction code that causes the message switch must be authorized to enter the transaction code that the CHNG call places in the alternate modifiable PCB. IMS TM does not check for security when you issue the ISRT call.

When an IMS TM application program issues a CHNG call, the Resource Access Control Facility (RACF) is invoked and a check is made to determine whether the originating terminal is authorized for the transaction code that was issued. If, instead of using the CHNG call, the program issues an ISRT call against a preset alternate PCB, no security check is made, regardless of the environment.

When you issue a program-to-program message switch, you have the same considerations as when you communicate with a logical terminal. Keep in mind the following points:

- Create an I/O area large enough to hold the largest segment that you are sending.
- Use an alternate PCB, not the TP PCB, to send the message.
- Issue a CHNG call before the ISRT call to place the transaction code of the program in the first field of the alternate PCB. If the alternate PCB was set to this transaction code in the PSBGEN, issue the ISRT call.
- IMS TM must know the transaction code. Define it at system definition.

- A nonconversational program can do a program-to-program message switch to another nonconversational program, but not to a conversational program.
- A conversational program can do a program-to-program message switch to either another conversational program or a nonconversational program.

Open Transaction Manager Access (OTMA) program-to-program switching has the following restrictions:

- In a shared queues environment that has both synchronous APPC/OTMA support (AOS=Y on the DFSDCxxx PROCLIB member) and 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.
- 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 message processing region (MPR).

A message switch to another conversational program transfers the scratchpad area (SPA) and the responsibility to respond to the originating terminal to the new application program. A message switch to a nonconversational program does not change the responsibilities of the conversational program. The conversational program must still return the SPA to IMS TM (if the SPA has been modified) and must respond to the originating terminal. The following tables show the format for an output message to an application program.

Table 78. Message Format for program-to-program message switch for AIBTDLI, ASMTDLI, CBLTDLI, CEETDLI, CTDLI, and PASTDLI interfaces

Field Name	Field Length
LL	2
Z1	1
Z2	1
Text	Variable

Table 79. Message format for program-to-program message switch for the PLITDLI interface

Field Name	Field Length
LLLL	4
Z1	1
Z2	1
Text	Variable

The format is the same as for output messages to terminals. Z1 and Z2 are fields that must contain binary zeros. These fields are reserved for IMS. The text field contains the message segment that you want to send to the application program.

If the program that is processing the message expects the transaction code, include the transaction code of the recipient program as part of the message text of the first segment of the message, because IMS TM does not automatically include the transaction code in the first segment of a switched message. Including the transaction code in the message text of the first segment keeps the first segments of all messages in the same format, regardless of whether they are sent from terminals or other programs.

Related concepts

[“Passing the conversation to another conversational program” on page 404](#)

A conversational program can pass the conversation to another conversational program in by performing a deferred switch or a immediate switch.

Related tasks

“Program switching in JMP and JBP applications” on page 694

IMS allows you to switch programs in JMP and JBP applications. You can perform immediate program switches in JMP and JBP applications, and you can also make a deferred program switch in a conversational JMP application.

How the VTAM I/O facility affects your VTAM terminal

VTAM terminals can fail to respond to requests sent by IMS. The master terminal operator or an automated operator interface application program can optionally activate a "timeout" facility. This allows a message stating a specific amount of time has passed to be sent to the master terminal operator.

IMS TM can be set up to do one of the following:

- Do nothing, which means that your terminal remains inactive. This is the default.
- Send a message to the master terminal operator stating that the specified period of time has passed. The operator can then determine what action, if any, should be taken.
- Send a message to the master terminal operator stating that the specified period of time has passed. IMS TM then issues the VTAM VARY NET, INACT command followed by a VTAM VARY NET, ACT command. If the terminal is defined to IMS TM as non-shared and operable, and if IMS TM is not shutting down, IMS TM issues an OPNDST for the terminal.

Restriction: This option does not apply to ISC terminals. If your installation chooses this option and an ISC terminal times out, a message is sent to the master terminal stating that the specified period of time has passed. The operator can determine what action, if any, should be taken.

Communicating with other IMS TM systems using Multiple Systems Coupling

In addition to communicating with programs and terminals in your IMS TM system, your program can communicate with terminals and programs in other IMS TM systems through Multiple Systems Coupling (MSC).

MSC makes this possible by establishing links between two or more separate IMS TM systems. The terminals and transaction codes within each IMS TM system are defined as belonging to that system. Terminals and transaction codes within your system are called "local," and terminals and transaction codes defined in other IMS TM systems connected by MSC links are called "remote."

Related reading: For an overview of MSC, see *IMS Version 15.3 Communications and Connections*.

Implications of MSC for program coding

For the most part, communicating with a remote terminal or program does not affect how you code your program. MSC handles the message routing between systems.

For example, if you receive an input message from a remote terminal, and you want to reply to that terminal, you issue an ISRT call against the I/O PCB—just as you would reply to a terminal in your system.

In the following two situations, MSC might affect your programming:

- When your program needs to know whether an input message is from a remote terminal or a local terminal. For example, if two terminals in separate IMS TM systems had the same logical terminal name, your program's processing might be affected by knowing which system sent the message.
- When you want to send a message to an alternate destination in another IMS TM system.

Restriction: If a transaction allocated by an LU 6.2 device is destined to a remote system through MSC links, IMS rejects the transaction with the message TP_NOT_Avail_No_Retry.

Directed routing makes it possible for your program to find out whether an input message is from your system or from a remote system, and to set the destination of an output message for an alternate destination in another IMS TM system. With directed routing, you can send a message to an alternate destination in another IMS TM system, even if that destination is not defined in your system as remote.

Restriction: MSC directed routing does not support a program-to-program switch between conversational transactions.

Related Reading: For more information about LU 6.2 and about MSC directed routing, see *IMS Version 15.3 Communications and Connections*.

Receiving messages from other IMS TM systems

When an application program retrieves an input message, the program can determine whether the input message is from a terminal or program in its IMS TM system, or from a terminal or program in another IMS TM system. There might be situations in which the application program's processing is changed if the input message is from a remote terminal, rather than from a local terminal.

For example, suppose that your IMS TM system is system A, and that it is linked to another IMS TM system called system B. MSC links are one-way links. The link from system A to system B is called LINK1, and the link from system B to system A is called LINK2. The application program named MPP1 runs in system A. The logical terminal name of the master terminals in both systems is MASTER. The following figure shows systems A and B.




Figure 77. MSC example

If the MASTER terminal in system B sends a message indicating that the system is shutting down to MPP1 in system A, MPP1 needs to know that the message is from MASTER in system B and not MASTER in system A.

If you have specified ROUTING=YES on the TRANSACT macro during IMS TM system definition, IMS TM does two things to indicate to the program that the message is from a terminal in another IMS TM system.

First, instead of placing the logical terminal name in the first field of the I/O PCB, IMS TM places the name of the MSC logical link in this field. In the example, this is LINK1. This is the logical link name that was specified on the MSNAME macro at system definition. However, if the message is subsequently sent back to the originating system, the originating LTERM name is reinstated in the first field of the I/O PCB.

Second, IMS TM turns on a bit in the field of the I/O PCB that is reserved for IMS. This is the second bit in the first byte of the 2-byte field. The following figure shows the location of this bit within the reserved field.




Figure 78. Directed routing bit in I/O PCB

MPP1 tests this bit to determine if the message is from MASTER in system A. If it is, MPP1 should terminate immediately. However, if the message is from MASTER in system B, MPP1 could perform some local processing and send transactions for system B to a message queue so that those transactions could be processed later on, when system B is up.

Sending messages to alternate destinations in other IMS TM systems

To send an output message to an alternate terminal in another IMS TM system, your system must have an MSC link with the system to which you want to send the message.

To do this, issue a CHNG call against an alternate PCB and supply the name of the MSC link (in the example this is LINK1) that connects the two IMS TM systems.

For example, if you were sending a message to TERMINAL 1 in system B after you received a message from some other terminal, you would first issue this CHNG call:

```
CHNG altpcb, LINK1
```

Then issue an ISRT call (or calls) to send the message just as you would send a message to a local terminal. The following tables show the format of the Direct Routing Output Message.

Table 80. Directed routing output message format for AIBTDLI, ASMTDLI, CBLTDLI, CEETDLI, CTDLI, and PASTDLI interfaces

Field Name	Field Length
LL	2
ZZ	2
DESTNAME	1 - 8
b	1
Text	Variable

Table 81. Directed routing output message format for the PLITDLI interface

Field Name	Field Length
LLLL	4
ZZ	2
DESTNAME	1 - 8
b	1
Text	Variable

The field formats in a directed routing output message are listed below:

- The LL and ZZ fields are 2 bytes each (For the PLITDLI interface, use the 4-byte field LLLL). LL (or LLLL) contains the total length of the message. This is the sum of all of the fields in the message, including the LL field (and in PL/I, LLLL contains the total length minus 2). ZZ is reserved for IMS.
- The destination name, DESTNAME, is the name of the logical terminal to which you are sending the message. This field is from 1 to 8 bytes long and it must be followed by a blank.

If the destination in the other system is a terminal, IMS TM removes the DESTNAME from the message. If the destination in the other system is a program, IMS TM does not remove the DESTNAME.

- The TEXT field contains the text of the message. Its length depends on the message you are sending.

If your message contains a security violation, MSC detects it in the receiving system (in this case, system B), and reports it to the person at the originating terminal (system A).

IMS conversational processing

You can write conversational programs to perform conversational processing with IMS Transaction Manager.

The difference between a conversational and nonconversational program is:

Conversational program

A message processing program (MPP) that processes transactions made up of several steps. It does not process the entire transaction at the same time. A conversational program divides processing into a connected series of terminal-to-program-to-terminal interactions. You use conversational processing when one transaction contains several parts.

Nonconversational program

A message processing program that receives a message from a terminal, processes the request, and sends a message back to the terminal. A conversational program receives a message from a terminal, and replies to the terminal, but saves the data from the transaction in a scratchpad area (SPA). Then, when the person at the terminal enters more data, the program has the data it saved from the last message in the SPA, so it can continue processing the request without the person at the terminal having to enter the data again.

A conversational example

The following example shows how to use conversational processing to find out if a customer can qualify for a car loan.

This inquiry contains two parts. First, you give the name and address of the person requesting the loan and the number of years for which the person wants the loan. After you provide this information, IMS TM asks you for the information on the car: model, year, and cost. You enter this information, IMS TM invokes the program that processes this information, and the program tells you whether the loan can be granted.

If you use MFS, the process involves these steps:

1. Enter the format command (/FORMAT) and the MOD name. This tells IMS to format the screen in the way defined by this message output descriptor (MOD).

If the MOD name is CL, the command is:

```
/FORMAT CL
```

IMS TM then takes that MOD from the MFS library and formats your screen in the way defined by the MOD. When the MOD for the car loan application formats your screen, it looks like this:

```
CARLOAN
NAME:
ADDRESS:
YEARS:
```

The word "CARLOAN" is the transaction code for this application. Each transaction code is associated with an application program, so when IMS TM receives the transaction code "CARLOAN", IMS TM knows what application program to schedule for this request.

2. Enter the customer's name and address, and the length of the loan. When you enter this information, your screen looks like this:

```
CARLOAN
NAME:    JOHN EDWARDS
ADDRESS: 463 PINWOOD
YEARS:   5
```

3. IMS TM reads the transaction code, CARLOAN, and invokes the program that handles that transaction code. MFS formats the information from the screen for the MPP's I/O area by using the DIF and the MID.

When the MPP issues its first call, which is usually a GU for the SPA, IMS TM clears the SPA to binary zeros and passes it to the application program.

4. Next, the MPP processes the input data from the terminal and does two things. It moves the data that it will need to save to the SPA, and it builds the output message for the terminal in the I/O area. The information that the MPP saves in the SPA is the information the MPP will need when the second part of the request comes in from the terminal. You do not save information in the SPA that you can get from the database. In this example, you save the name of the customer applying for the loan, because if the customer is granted the loan, the program uses the customer name to locate the information to be updated in the database.

The program then issues an ISRT call to return the SPA to IMS, and another ISRT call to send the output message to the terminal.

The response that the MPP sends to the terminal gives IMS TM the name of the MOD to format the screen for the next cycle of the conversation. In that cycle, you need to supply the model, year, and cost of the car that John Edwards wants to buy. Your screen looks like this:

```
MODEL :
YEAR :
COST :
```

5. IMS TM again uses the device input format (DIF) and message input descriptor (MID) associated with the transaction code, and sends the information back to the MPP. The MPP has not been running all this time. When IMS TM receives the terminal input with the transaction code CARLOAN, IMS TM invokes the MPP that processes that transaction again for this cycle of the conversation.
6. IMS TM returns the updated SPA to the MPP when the MPP issues a GU, then returns the message to the MPP when the MPP issues a GN. The MPP does the required processing (in this case, determining whether the loan can be granted and updating the database if necessary), and is then ready to end the conversation. To do this, the MPP blanks out the transaction code in the SPA, inserts it back to IMS, then sends a message to the terminal saying whether the loan can be granted.

Conversational structure

Structuring your conversational program depends on the interactions between your program and the person at the terminal.

Before structuring your program, you need to know:

- What should the program do in an error situation?

When a program in a conversation terminates abnormally, IMS TM backs out only the last cycle of the conversation. A cycle in a conversation is one terminal/program interaction. Because the conversation can terminate abnormally during any cycle, you should be aware of some things you can do to simplify recovery of the conversation:

- The ROLB or ROLS call can be used in conversational programs to back out database updates that the program has made since the last commit point. ROLL can also be used in conversational programs, but terminates the conversation.

- If possible, updating the database should be part of the last cycle of the conversation so that you do not have different levels of database updates resulting from the conversation.
- If your program encounters an error situation and it has to terminate, it can use an express alternate (program communication block) PCB to send a message to the originating terminal, and, if desired, to the master terminal operator.

To do this, the program issues a CHNG call against the express alternate PCB and supplies the name of the logical terminal from the TP PCB, then an ISRT call that references that PCB and the I/O area that contains the message. The program can then issue another CHNG call to set the destination of the express alternate PCB for the master terminal, and another ISRT call that references that PCB, and the I/O area that contains the output message.

- Does your application program process each cycle of the conversation?

A conversation can be processed by one or several application programs. If your program processes each stage of the conversation (in other words, your program processes each input message from the terminal), the program has to know what stage of the conversation it is processing when it receives each input message.

When the person at the terminal enters the transaction code that starts the conversation, IMS TM clears the SPA to binary zeros and passes the SPA to the program when the program issues a GU call. On subsequent passes, however, the program has to be able to tell which stage of the conversation it is on so that it can branch to the section of the program that handles that processing.

One technique that the program can use to determine which cycle of the conversation it is processing is to keep a counter in the SPA. The program increments this counter at each stage of the conversation. Then, each time the program begins a new cycle of the conversation (by issuing a GU call to retrieve the SPA), the program can check the counter in the SPA to determine which cycle it is processing, then branch to the appropriate section.

- How can your program pass control of the conversation to another conversation program?

Sometimes it is more efficient to use several application programs to process a conversation. This does not affect the person at the terminal. It depends on the processing that is required.

In the car loan example, one MPP could process the first part of the conversation (processing the name, address, and number of years), and another MPP could process the second part of the conversation (processing the data about the car and responding with the status of the loan).

A conversational program can perform two types of program switching:

Deferred program switch

Responds to the originating terminal but causes the next input from the terminal to go to another conversational program.

Immediate program switch

Passes the conversation directly to another conversational program. The program passes the SPA (and, optionally, a message) to another conversational program without responding to the terminal. In this case, it is the next program's responsibility to respond to the originating terminal.

A conversational program must:

1. Retrieve the SPA and the message using GU and GN calls.

If your MPP is starting this conversation, test the variable area of the SPA for zeros to determine if this is the beginning of the conversation. If the SPA does not contain zeros, it means that you started the conversation earlier and that you are now at a later stage in the conversation. If this is true, you would branch to the part of your program that processes this stage of the conversation to continue the conversation.

If another MPP has passed control to your MPP to continue the conversation, the SPA contains the data you need to process the message, so you do not have to test it for zeros. Start processing the message immediately.

2. Process the message, including handling any necessary database access.

3. Send the output message to the terminal by using an ISRT call against the I/O PCB. This step can follow step 4.
4. Store the data (that your program, or the program that you pass control to, needs to continue processing) in the SPA using an ISRT call to the I/O PCB. (This step can precede step 3.) IMS TM determines which segment is the SPA by examining the ZZZZ field of the segment shown in the tables below.

To end the conversation, move blanks to the area of the SPA that contains the transaction code, and then insert the SPA back to IMS TM by issuing an ISRT call and referencing the I/O PCB.

If your MPP passes the conversation to another conversational program, the steps after the program processes the message are somewhat different.

Also, your program should be designed to handle the situation that occurs when the first GU call to the I/O PCB does not return a message to the application program. This can happen if the person at the terminal cancels the conversation by entering the /EXIT command before the program issues a GU call. (This happens if the message from this terminal was the only message in the message queue for the program.)

The contents of SPA

The SPA that IMS TM gives your program when you issue a GU contains the four parts shown in the following tables.

Table 82. SPA format for AIBTDLI, ASMTDLI, CBLTDLI, CEETDLI, CTDLI, and PASTDLI interfaces

Field Name	Field Length
LL	2
ZZZZ	4
TRANCODE	8
User Work Area	Variable

Table 83. SPA format for the PLITDLI interface

Field Name	Field Length
LLLL	4
ZZZZ	4
TRANCODE	8
User Work Area	Variable

The SPA format fields are:

LL or LLLL

A length field that gives the total length of the SPA. This length includes 2 bytes for the LL field. (For the PLITDLI interface, use a 4-byte field. Its contents include 4 bytes for LLLL, minus 2.)

ZZZZ

A 4-byte field reserved for IMS TM that your program must not modify.

TRANCODE

The 8-byte transaction code for this conversation.

User Work Area

A work area that you use to save the information that you need to continue the conversation. The length of this area depends on the length of the data you want to save. This length is defined at system definition.

When your program retrieves the SPA with a GU to start the conversation, IMS TM removes the transaction code from the message. In your first message segment, you receive only the data from the message that the person at the terminal entered.

The following list indicates the ways that an application program processes the SPA. The program must:

- Not modify the first 6 bytes of the SPA (LL and ZZZZ). IMS TM uses these fields to identify the SPA.

If the program modifies the SPA, the program must return the SPA to IMS TM (or, for a program switch, to the other program).

- Not return the SPA to IMS TM more than once during one cycle of the conversation.
- Not insert the SPA to an alternate PCB that represents a nonconversational transaction code or a logical terminal. The program can use an alternate response PCB if it represents that same physical terminal as the originating logical terminal.

Restriction: If you are using MFS, the IMS TM does not always remove the transaction code.

The appearance of messages in a conversation

Because the first segment contains the SPA, conversational input messages are made up of at least two segments. The input message starts in the second message segment.

The input message segment in a conversation contains only the data from the terminal. During the first step in the conversation, IMS TM removes the transaction code from the input message and places it in the SPA. When the program issues the first GU, IMS TM returns the SPA. To retrieve the first message segment, the program must issue a GN.

The format for the output messages that you send to the terminal is no different than the format for output messages in nonconversational programs.

Saving information in the SPA

After you have processed the message and are ready to reply to the terminal, you can save the necessary data in the SPA. The part of the SPA in which you save data is the work area portion. Use the ISRT call to save data to the work area. This is a special use of the ISRT call, because you are not sending the SPA to a terminal, but rather saving it for future use.

If your program processes each stage of the conversation, you just issue an ISRT call to the I/O PCB and give the name of the I/O area that contains the SPA. For example:

```
ISRT    I/O PCB, I/O AREA
```

This returns the updated SPA to IMS TM so that IMS TM can pass it to your program at the next cycle of the conversation.

If you do not modify the SPA, you do not need to return it to IMS. However, the SPA will be passed by IMS TM to your program at the next cycle of the conversation.

Related concepts

[“Conversational processing using ROLB, ROLL, and ROLS” on page 403](#)

Issuing a ROLB or ROLS in a conversational program causes IMS TM to back out the messages that the application program has sent.

[“Passing the conversation to another conversational program” on page 404](#)

A conversational program can pass the conversation to another conversational program in by performing a deferred switch or a immediate switch.

Replying to the terminal

For a conversation to continue, the originating terminal must receive a response to each of its input messages. The person at the terminal cannot enter any more data to be processed (except IMS TM commands) until the response has been received at the terminal.

To continue the conversation, the program must respond to the originating terminal by issuing the required ISRT calls to send the output message to the terminal. To send a message to the originating terminal, the ISRT calls must reference either the TP PCB or an alternate response PCB. Use an alternate response PCB in a conversation when the terminal you are responding to has two components—for example, a printer and a punch—and you want to send the output message to a component that is separate from the component that sent the input message. If the program references an alternate response PCB, the PCB must be defined for the same physical terminal as the logical terminal that sent the input message.

The program can send only one output message to the terminal for each input message. Output messages can contain multiple segments, but the program cannot use the PURG call to send multiple output messages. If a conversational program issues a PURG call, IMS TM returns an AZ status code to the application program and does not process the call.

Conversational processing using ROLB, ROLL, and ROLS

Issuing a ROLB or ROLS in a conversational program causes IMS TM to back out the messages that the application program has sent.

If the application program issues a ROLB or ROLS and then reaches a commit point without sending the required response to the originating terminal, IMS TM terminates the conversation and sends the message DFS2171I NO RESPONSE CONVERSATION TERMINATED to the originating terminal.

If you issue ROLL during a conversation, IMS TM backs out the updates and cancels output messages, but it also terminates the conversation.

Conversational processing for modified message-driven IMS applications

The following processing considerations apply to modified message-driven IMS applications issuing the IMS ROLB call that can receive protected input messages from OTMA or APPC/MVS and issue outbound protected work to other z/OS Resource Recovery Services (RRS) resource managers:

- If a modified message-driven IMS application program with protected input issues a ROLB call, the ROLB call is isolated to the IMS application without affecting the entire protected unit of work. After the ROLB call is issued, the protected input message remains in process for the IMS application until a commit point is reached.
- If a modified message-driven IMS application program issues an outbound protected conversation, the outbound protected conversation is not included in the ROLB processing (that is, the outbound protected conversation is not backed out as part of the ROLB call). The modified message-driven IMS application program is responsible for explicitly cleaning up any outbound protected work to be backed out.

Related concepts

[“Conversational structure” on page 399](#)

Structuring your conversational program depends on the interactions between your program and the person at the terminal.

Passing the conversation to another conversational program

A conversational program can pass the conversation to another conversational program in by performing a deferred switch or a immediate switch.

A conversational program can pass the conversation to another conversational program in two ways:

- A deferred switch.

The program can respond to the terminal but cause the next input from the terminal to go to another conversational program by:

- Issuing an ISRT call against the I/O PCB to respond to the terminal
- Placing the transaction code for the new conversational program in the SPA
- Issuing an ISRT call referencing the I/O PCB and the SPA to return the SPA to IMS TM

IMS TM then routes the next input message from the terminal to the program associated with the transaction code that was specified in the SPA. Other conversational programs can continue to make program switches by changing the transaction code in the SPA.

- An immediate switch.

The program can pass the conversation directly to another conversational program by issuing an ISRT call against the alternate PCB that has its destination set to the other conversational program.

The first ISRT call must send the SPA to the other program, but the program passing control can issue subsequent ISRT calls to send a message to the new program. If the program does this, in addition to routing the SPA to the other conversational program, IMS TM updates the SPA as if the program had returned the SPA to IMS. If the program does an immediate switch, the program cannot also return the SPA to IMS TM or respond to the original terminal.

Restrictions on passing the conversation

These are restrictions that apply to passing the conversation to another conversational program:

- When an immediate program switch occurs and the MPP receives an XE status code, the program attempts to insert the SPA to an alternate express PCB. Remove the EXPRESS=YES option from the PCB or define and use another PCB that is not express. This restriction prevents the second transaction from continuing the conversation if the first transaction abends after inserting the SPA.

The person at the terminal can issue the /SET CONV XX command, where XX is the program that is to be scheduled in order to process the next step of the conversation.

- APPC or OTMA protected transactions do not allow immediate program or deferred program switches. If either of these switches occur, the MPP receives an X6 status code.

Defining the SPA size

Define the SPA size with the TRANSACT macro. An option to capture truncated data is also defined with the TRANSACT macro. The format is:

```
TRANSACT SPA=(size,STRUNC|RTRUNC)
```

The default is to support truncated data (STRUNC). When a conversation is initially started, and on each program switch, the truncated data option is checked and set or reset as specified. When the truncated data option is set, it remains set for the life of the conversation, or until a program switch occurs to a transaction that specifies that the option be reset.

For example, assume you have three transactions defined as follows:

```
TRANA SPA=100
```


TRANB SPA=050

TRANC SPA=150

For TRANC to receive the truncated data (which is the second 50 bytes from TRANA that TRANB does not receive) from TRANA, one of the following sets of specifications can be used:

- TRANA - STRUNC or none, TRANB - STRUNC or none, TRANC - STRUNC or none
- TRANA - RTRUNC, TRANB - STRUNC, TRANC - STRUNC or none

Conversational processing and MSC

If your installation has two or more IMS TM systems, and they are linked to each other through MSC, a program in one system can process a conversation that originated in another system.

- If a conversational program in system A issues an ISRT call that references a response alternate PCB in system B, system B does the necessary verification. This is because the destination is implicit in the input system. The verification that system B does includes determining whether the logical terminal that is represented by the response alternate PCB is assigned to the same physical terminal as the logical terminal that sent the input message. If it is not, system B (the originating system) terminates the conversation abnormally without issuing a status code to the application program.
- Suppose program A processes a conversation that originates from a terminal in system B. Program A passes the conversation to another conversational program by changing the transaction code in the SPA. If the transaction code that program A supplies is invalid, system B (the originating system) terminates the conversation abnormally without returning a status code to the application program.

Ending the conversation

To end the conversation, a program blanks out the transaction code in the SPA and returns it to IMS TM by issuing an ISRT call and referencing the I/O PCB and the SPA. This terminates the conversation as soon as the terminal has received the response.

The program can also end the conversation by placing a nonconversational transaction code in the transaction field of the SPA and returning the SPA to IMS. This causes the conversation to remain active until the person at the terminal has entered the next message. The transaction code will be inserted from the SPA into the first segment of the input message. IMS TM then routes this message from the terminal to the MPP or BMP that processes the transaction code that was specified in the SPA.

In addition to being ended by the program, a conversation can be ended by the person at the originating terminal, the master terminal operator, and IMS.

- The person at the originating terminal can end the conversation by issuing one of several commands:

/EXIT

The person at the terminal can enter the /EXIT command by itself, or the /EXIT command followed by the conversational identification number assigned by the IMS TM system.

/HOLD

The /HOLD command stops the conversation temporarily to allow the person at the terminal to enter other transactions while IMS TM holds the conversation. When IMS TM responds to the /HOLD command, it supplies an identifier that the person at the terminal can later use to reactivate the conversation. The /RELEASE command followed by this identifier reactivates the conversation.

- **/START LINE.** The master terminal operator can end the conversation by entering a /START LINE command (without specifying a PTERM) or /START NODE command for the terminal in the conversation or a /START USER command for a signed-off dynamic user in conversation.
- IMS TM ends a conversation if, after the program successfully issues a GU call or an ISRT call to return the SPA, the program does not send a response to the terminal. In this situation, IMS TM sends the message DFS2171I NO RESPONSE, CONVERSATION TERMINATED to the terminal. IMS TM then terminates the conversation and performs commit point processing for the application program.

Related concepts

[“Sending messages to other IMS application programs” on page 393](#)

A program-to-program switch occurs when an IMS application running in an IMS dependent region sends a message to another IMS application running in an IMS dependent region.

[“Conversational structure” on page 399](#)

Structuring your conversational program depends on the interactions between your program and the person at the terminal.

Related tasks

[“Deferred program switching for conversational JMP applications” on page 696](#)

You can make a deferred program switch in a conversational JMP application. A deferred program switch changes the transaction code in the scratchpad area (SPA) before the SPA is returned to IMS. When an application makes a deferred program switch, the application replies to the terminal and passes the conversation to another conversational application.

[“Immediate program switching for JMP and JBP applications” on page 694](#)

The IMS Java dependent region resource adapter supports immediate program switching in JMP and JBP applications. An immediate program switch passes the conversation directly to another conversational program that is specified by an alternate PCB.

Message switching in APPC conversations

With the system service DFSAPPC, you can transfer messages between separate LU 6.2 devices and between an LU 6.2 device and another terminal supported by IMS TM. Message delivery with DFSAPPC is asynchronous, so messages are held on the IMS TM message queue until they can be delivered.


To send a message with DFSAPPC, specify the logical terminal name of an IMS TM terminal or the Transaction Program (TP) name of an LU 6.2 device.

DFSAPPC format

The message format for DFSAPPC is as follows:

```
DFSAPPC (options)user_data
```

DFSAPPC can be coded as follows:



A blank (b) is required between DFSAPPC and the specified options.

Blanks are valid within the specified options except within keywords or values. Either commas or blanks can be used as delimiters between options, but because the use of commas is valid, the TP name must be followed by at least one blank.

If an LU 6.2 conversation has not been established from other sources (for example, during a CPI-C driven application program), DFSAPPC is used to establish the conversation with a partner LU 6.2 device. If no options are specified with DFSAPPC, IMS TM default options are used.

Option keywords

LTERM=

Specifies the LTERM name of an IMS TM logical terminal. An LTERM name can contain up to eight alphanumeric or national (@, \$, #) characters. If you specify LTERM, you cannot specify the other option keywords.

LU=

Specifies the LU name of the partner in an LU 6.2 conversation. The LU name can contain up to eight alphanumeric or national characters, but the first character must be a letter or a national character. If both LU and SIDE options are specified, LU overrides the LU name contained in the side information entry but does not change that LU name.

If the LU name is a network-qualified name, it can be up to 17 characters long and consist of the network ID of the originating system, followed by a '.', and the LU name (for example, netwrkid.luname). The LU name and the network ID can be up to eight characters long.

MODE=

Specifies the MODE name of the partner in an LU 6.2 conversation. The MODE name can contain up to eight alphanumeric or national characters, but the first character must be a letter or a national character. If both MODE and SIDE option keywords are specified, MODE overrides the MODE name contained in the side information entry but does not change that MODE name.

TPN=

Specifies the transaction program (TP) name of the partner in an LU 6.2 conversation. The TP name can contain up to 64 characters from the 00640 character set. Because the character set allows commas, at least one blank must follow the TP name. If both TPN and SIDE option keywords are specified, TPN overrides the TP name contained in the side information entry but does not change that name.

Related Reading: The *CPI Communications Specification* describes the 00640 character set, which contains all alphanumeric and national characters and 20 special characters.

SIDE=

Specifies the name of the side information entry for the partner in an LU 6.2 conversation. The side information entry name can contain up to eight characters from the 01134 character set. If the SIDE option keyword is specified, it can be overridden with LU, MODE, and TPN option keywords.

Related Reading: The *CPI Communications Specification* describes the 01134 character set, which contains the uppercase alphabet and the digits, 0-9.

SYNC=N | C

Specifies the synchronization level of the LU 6.2 conversation. N selects none as the synchronization level, and C selects confirm as the synchronization level.

TYPE=B | M

Specifies the conversation type for the LU 6.2 conversation. B selects a basic conversation type, and M selects a mapped conversation type.

Processing conversations with APPC

APPC/IMS supports standard, modified, and CPI Communications driven application programs.

The three types of application programs supported by APPC/IMS

- 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.

In the modified or CPI Communications driven application programs, if an APPC conversation is allocated with SYNCLVL=SYNCP, z/OS manages the sync-point process for the APPC conversation participants: the

application program and IMS. Transaction rollback and rescheduling is possible, because IMS issues the SRRCMIT or SRRBACK calls on behalf of the modified IMS APPC application program. If the CPI-C driven program is linked with the IMS stub code (DFSCPIR0) as required in previous releases, IMS also issues the SRRCMIT or SRRBACK calls. If the program is not linked with the stub code, then IMS is driven by the z/OS sync point manager when the application issues these calls. With z/OS as the sync point manager, failures can also be backed out.

You can schedule your standard and modified application programs locally and remotely using MSC or APPC/MVS. The logic flow for local scheduling differs from the logic flow for remote scheduling.

Scheduling programs remotely through MSC is not supported if an APPC/MVS conversation with SYNCLVL=SYNCPT is specified.

Ending the APPC conversation

You can end a conversation using LU 6.2 devices by issuing the CPI-C verb, DEALLOCATE, or by inserting a blank transaction code into the SPA for IMS conversational transactions.

Restriction: You cannot use the /EXIT command for LU 6.2 conversations.

Several error conditions can exist at the end of an LU 6.2 conversation:

- If your application program sends data to the LU 6.2 device just before deallocating conversation, IMS TM issues a SENDERROR and SENDDATA of the DFS1966 error message. This indicates that the transaction ended, but that the last message could not be delivered. For SENDERROR to be activated, specify a synchronization level of CONFIRM.
- If IMS TM encounters an error sending output from an IMS TM conversational transaction to the LU 6.2 device, the output is discarded, and the conversation is terminated for both IMS TM and LU 6.2.
- If an IMS TM conversational application program abends during an LU 6.2 conversation, a DFS555 error message is sent to the originating LU 6.2 device, and the conversation is terminated for both IMS TM and LU 6.2.

Coding a conversational program

Before coding a conversational program, you need to obtain the following information.

- The transaction code to use for a program to which you pass control
- The data that you should save in the SPA
- The maximum length of that data

A SPA contains four fields:

- The 2-byte length field.
- The 4-byte field that is reserved for IMS TM.
- The 8-byte transaction code.
- The work area where you store the conversation data. The length of this field is defined at system definition.

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 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, regardless of whether LU 6.2 is used. 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.

A non-message-driven BMP is considered a standard IMS application program when it does not use the explicit API.

When an advanced program-to-program communication (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 system. 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 system. From this point on, the transaction goes through normal MSC processing. After the remote IMS system executes the transaction, the output is returned to the local IMS system and is then delivered to the originating LU 6.2 application program.

Modified IMS application programs

Modified IMS application programs use a DL/I GU call to get the incoming transaction. These modified IMS application programs also use DL/I ISRT calls to generate output messages to the same or different terminals, regardless of whether LU 6.2 is used.

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

Modified IMS programs are scheduled by IMS TM, and the DL/I calls are processed by the DL/I language interface. The conversation, however, is maintained by APPC/MVS, and any failures that involve APPC/MVS are not backed out by IMS TM. The general format of a modified IMS application program is shown in the following code example.

- GU IOPCB
 ALLOCATE
 SEND
 RECEIVE
 DEALLOCATE
- ISRT IOPCB

Figure 79. General format of a modified DL/I application program

Restriction: The APPC conversation cannot span sync points. If the conversation is not deallocated before a sync point is reached, IMS causes the conversation to be terminated by issuing a clean TP call (ATBCMTP). A new APPC conversation can be allocated after each sync point.

When an APPC program enters an IMS transaction that executes on a remote IMS system, an LU 6.2 conversation is established between the APPC program and the local IMS system. The local IMS system is considered the partner LU of the LU 6.2 conversation. The transaction is then queued on the local IMS system's remote transaction queue. From this point on, the transaction goes through normal MSC processing. After the remote IMS system executes the transaction, the output is returned to the local IMS and is then delivered to the originating LU 6.2 program.

Related Reading: For more information on failure recovery and modified DL/I application program design, see *IMS Version 15.3 Application Programming APIs*.

CPI-C 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. Their 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 through the APPLCTN macro for system definition and through PSBGEN or ACBGEN when APPLCTN PSB= is specified. When APPLCTN GPSB= is specified, a PSBGEN or ACBGEN is not required.

CPI-C driven application programs must begin with the CPI-C verbs, ACCEPT and RECEIVE, to initiate the LU 6.2 conversation. You can then issue the APSB call to allocate a PSB for use by the application program. After the APSB call is issued, you can issue additional DL/I calls using the PCBs that were allocated. You then issue the SRRCMIT verb to commit changes or the SRRBACK verb to back out changes. To use SRRCMIT and SRRBACK, your application program must be linked with DFSCPIR0.

Restriction: The I/O PCB cannot be used for message processing calls by CPI-C driven application programs. See the description of each call for specific CPI restrictions.

To deallocate the PSB in use, issue the DPSB call. You can then issue another APSB call, or use the CPI-C verb, DEALLOCATE, to end the conversation.

CPI-C driven application programs are considered discardable (unless they are allocated with a SYNCLVL=SYNCPT) by IMS TM and are therefore not recovered automatically at system failure. If they are allocated with a SYNCLVL=SYNCPT, a two-phase commit process is used to recover from any failures. The general format of a CPI-C driven application program is shown in the following code example.

- ACCEPT
- RECEIVE
 - APSB
 - GU DBPCB
 - REPL DBPCB
 - SRRCMIT
 - DPSB
- DEALLOCATE

Figure 80. General format of a CPI-C driven application program

Related concepts

[CPI-C driven application programs \(Communications and Connections\)](#)

Processing conversations with OTMA

You can run IMS conversational transactions through OTMA.

Refer to *IMS Version 15.3 Communications and Connections*.

Backing out to a prior commit point: ROLL, ROLB, and ROLS calls

When a program determines that some of its processing is invalid, you can use these calls to remove the effects of its incorrect processing: Roll Back calls ROLL, ROLS using a database PCB, ROLS with no I/O area or token, and ROLB.

When you issue one of these calls, IMS does the following:

- Backs out the database updates that the program has made since the program's most recent commit point.
- Cancels the non-express output messages that the program has created since the program's most recent commit point.

The main difference among these calls is that ROLB returns control to the application program after backing out updates and canceling output messages, ROLS does not return control to the application program, and ROLL terminates the program with a user abend code of 0778. ROLB can return to the program the first message segment since the most recent commit point, but ROLL and ROLS cannot.

The ROLL and ROLB calls, and the ROLS call without a token specified, are valid when the PSB contains PCBs for Generalized Sequential Access Method (GSAM) data sets. However, segments inserted in the GSAM data sets since the last commit point are not backed out by these calls. An extended checkpoint-restart can be used to reposition the GSAM data sets when restarting.

You can use a ROLS call either to back out to the prior commit point or to back out to an intermediate backout point established by a prior SETS call. This section refers only to the form of ROLS that backs out to the prior commit point.

Related concepts

[“Backing out to an intermediate backout point: SETS, SETU, and ROLS” on page 267](#)

You can use a ROLS call either to back out to an intermediate backout point that was established by a prior SETS or SETU call, or to back out to the prior commit point.

Comparison of ROLB, ROLL, and ROLS

The following table provides a comparison of the ROLB, ROLL, and ROLS calls.

<i>Table 84. Comparison of ROLB, ROLL, and ROLS</i>			
Actions taken	ROLB	ROLL	ROLS
Back out database updates since the last commit point.	X	X	X
Cancel output messages created since the last commit point.	X ¹	X ¹	X ¹
Delete the message in process from the queue. Previous messages (if any) processed since the last commit point are returned to the queue to be reprocessed.		X	
Return the first segment of the first input message since the most recent commit point.	X ²		
3303 abnormal termination and returns the processed input messages to the message queue.			X ³
778 abnormal termination, no dump.		X	
Noabend; program continues processing.	X		

Notes:

1. ROLB, ROLL, or ROLS cancel output messages sent with an express PCB unless the program issued a PURG.

For example, if the program issues the following call sequence, MSG1 would be sent to its destination because the PURG tells IMS that MSG1 is complete and the I/O area now contains the first segment of the next message (which in this example is MSG2). MSG2, however, would be canceled:

```
ISRT    EXPRESS PCB, MSG1
PURG    EXPRESS PCB, MSG2
ROLB    I/O PCB
```

Because IMS has the complete message (MSG1) and because an express PCB is being used, the message can be sent before a commit point.

2. Returned only if you supply the address of an I/O area as one of the call parameters.
3. The transaction is suspended and requeued for subsequent processing.

ROLL

A ROLL call backs out the database updates and cancels any non-express output messages the program has created since the last commit point. It also deletes the current input message. Any other input messages processed since the last commit point are returned to the queue to be reprocessed. IMS then terminates the program with a user abend code 0778. This type of abnormal termination terminates the program without a storage dump.

When you issue a ROLL call, the only parameter you supply is the call function, ROLL.

You can use the ROLL call in a batch program. If your system log is on direct access storage, and if dynamic backout has been specified through the use of the BKO execution parameter, database changes since the last commit point will be backed out. Otherwise they will not be backed out. One reason for issuing ROLL in a batch program is for compatibility.

After backout is complete, the original transaction is discarded if it is discardable, and it is not re-executed. IMS issues the APPC/MVS verb ATBCMTP TYPE(ABEND) specifying the TPI to notify remote transaction programs. Issuing the APPC/MVS verb causes all active conversations (including any spawned by the application program) to be DEALLOCATED TYP(ABEND_SVC).

ROLB

The advantage of using ROLB is that IMS returns control to the program after executing ROLB, so the program can continue processing.

The parameters for ROL are:

- The call function ROLB
- The name of the I/O PCB or AIB

The total effect of the ROLB call depends on the type of IMS application that issued it.

- For current IMS application programs:

After IMS backout is complete, the original transaction is represented to the IMS application program. Any resources that cannot be rolled back by IMS are ignored. For example, output sent to an express alternate PCB and a PURG call is issued before the ROLB.

- For modified IMS application programs:

The same consideration for the current IMS application programs applies. It is the responsibility of the application program to notify any spawned conversations that a ROLB was issued.

- For CPI-C driven IMS application programs:

Only IMS resources are affected. All database changes are backed out. Any messages inserted to nonexpress alternate PCBs are discarded. Also, any messages inserted to express PCBs that have not had a PURGE call are discarded. It is the responsibility of the application program to notify the originating remote program and any spawned conversations that a ROLB call was issued.

In MPPs and transaction-oriented BMPs

If the program supplies the address of an I/O area as one of the ROLB parameters, the ROLB call acts as a message retrieval call and returns the first segment of the first input message since the most recent commit point. This is true only if the program has issued a GU call to the message queue since the last commit point; if it has not, it was not processing a message when it issued the ROLB call.

If the program issues a GN to the message queue after issuing the ROLB, IMS returns the next segment of the message that was being processed when ROLB was issued. If there are no more segments for that message, IMS returns a QD status code.

If the program issues a GU to the message queue after the ROLB call, IMS returns the first segment of the next message to the application program. If there are no more messages on the message queue for the program to process, IMS returns a QC status code to the program.

If you include the I/O area parameter, but you have not issued a successful GU call to the message queue since the last commit point, IMS returns a QE status code to your program.

If you do not include the address of an I/O area in the ROLB call, IMS does the same things for you. If the program has issued a successful GU in the commit travel, and then issues a GN, IMS returns a QD status code. If the program issues a GU after the ROLB, IMS returns the first segment of the next message, or a QC status code if there are no more messages for the program.

If you have not issued a successful GU since the last commit point, and you do not include an I/O area parameter on the ROLB call, IMS backs out the database updates and cancels the output messages created since the last commit point.

In batch programs

If your system log is on direct access storage, and if dynamic backout has been specified through the use of the BKO execution parameter, you can use the ROLB call in a batch program. The ROLB call does not process messages as it does for message processing programs (MPPs); it backs out the database updates since the last commit point and returns control to your program. You cannot specify the address of an I/O area as one of the parameters on the call; if you do, an AD status code is returned to your program. You must, however, have an I/O PCB for your program. Specify CMPAT=YES on the CMPAT keyword in the PSBGEN statement for your program's PSB.

Related Reading: For more information on using the CMPAT keyword, see *IMS Version 15.3 System Utilities*. For information on coding the ROLB call, see the topic "ROLB Call" in *IMS Version 15.3 Application Programming APIs*.

ROLS

You can use the ROLS call to back out to the prior commit point and return the processed input messages to IMS for later reprocessing.

In your program, you can either:

- Issue the ROLS call using the I/O PCB but without an I/O area or token in the call. The parameters for this form of the ROLS call are:
 - The call function ROLS
 - The name of the I/O PCB or AIB
- Issue the ROLS call using a database PCB that has received one of the data-unavailable status codes. This has the same result as if unavailable data were encountered, and the INIT call was not issued. ROLS must be the next call for that PCB. Intervening calls using other PCBs are permitted.

On a ROLS with a token, message queue repositioning can occur for all non-express messages including all messages processed by IMS. This processing using APPC/MVS calls and includes the initial message segments. The original input transaction can be represented to the IMS application program. Input and output positioning is determined by the SETS call. This positioning applies to current and modified IMS application programs but does not apply to CPI-C driven IMS programs. The IMS application program must notify all remote transaction programs of the ROLS.

On a ROLS without a token, IMS issues the APPC/MVS verb, ATBCMTP TYPE(ABEND), specifying the TPI. Issuing this verb causes all conversations associated with the application program to be DEALLOCATED TYPE(ABEND_SVC). If the original transaction was entered from an LU 6.2 device and IMS received the message from APPC/MVS, a discardable transaction is discarded rather than being placed on the suspend queue like a non-discardable transaction.

Related Reading: For more information on LU 6.2, see *IMS Version 15.3 Communications and Connections*.

The parameters for this form of the ROLS call are:

- The call function, ROLS
- The name of the DB PCB that received the BA or BB status code

In both of the ways to use ROLS calls, the ROLS call causes a 3303 abnormal termination and does not return control to the application program. IMS keeps the input message for future processing.

Backing out to an intermediate backout point: SETS/SETU and ROLS

You can use a ROLS call either to back out to an intermediate backout point established by a prior SETS or SETU call or to back out to the prior commit point.

This section refers only to the form of ROLS that backs out to the intermediate backout point. For information about the other form of ROLS, see 'Backing out to a prior commit point: ROLL, ROLB, and ROLS calls'.

The ROLS call that backs out to an intermediate point backs out only DL/I changes. This version of the ROLS call does not affect CICS changes using CICS file control or CICS transient data.

The SETS and ROLS calls set intermediate backout points within the call processing of the application program and then backout database changes to any of these points. Up to nine intermediate backout points can be set. The SETS call specifies a token for each point. IMS then associates this token with the current processing point. A subsequent ROLS call, using the same token, backs out all database changes and discards all non-express messages that were performed following the SETS call with the same token. The figure below shows how the SETS and ROLS calls work together.

In addition, to assist the application program in reestablishing other variables following a ROLS call, user data can be included in the I/O area of the SETS call. This data is then returned when the ROLS call with the same token is issued.




Figure 81. SETS and ROLS calls working together

SETS/SETU

The SETS call sets up to nine intermediate backout points or cancels all existing backout points. By using the SETS call, you can back out pieces of work. If the necessary data to complete one piece of work is unavailable, you can complete a different piece of work and then return to the former piece.

To set an intermediate backout point, issue the call using the I/O PCB and include an I/O area and a token. The I/O area has the format LLZZ user-data, where LL is the length of the data in the I/O area including the length of the LLZZ portion. The ZZ field must contain binary zeros. The data in the I/O area is returned to the application program on the related ROLS call. If you do not want to save some data to be returned on the ROLS call, you must set the LL that defines the length of the I/O area to 4.

For PLITDLI, you must define the LL field as a fullword rather than a halfword as it is for the other languages. The content of the LL field for PLITDLI is consistent with the I/O area for other calls using the LLZZ format; that is, the content is the total length of the area including the length of the 4-byte LL field minus 2.

A 4-byte token associated with the current processing point is also required. This token can be a new token for this program execution or match a token issued by a preceding SETS call. If the token is new, no preceding SETS calls are canceled. If the token matches the token of a preceding SETS call, the current SETS call assumes that position. In this case, all SETS calls that were issued subsequent to the SETS call with the matching token are canceled.

The parameters for this form of the SETS call are:

- The call function SETS
- The name of the I/O PCB or AIB
- The name of the I/O area containing the user data
- The name of an area containing the token

For the SETS call format, see the topic 'SETS/SETU Call' in *IMS Version 15.3 Application Programming APIs*.

To cancel all previous backout points, the call is issued using the I/O PCB but does not include an I/O area or a token. When no I/O area is included in the call, all intermediate backout points set by prior SETS calls are canceled.

The parameters for this form of the SETS call are:

- The call function SETS
- The name of the I/O PCB or AIB

Because it is not possible to back out committed data, commit point processing causes all outstanding SETS to be canceled.

If PCBs for DEDB, MSDB, and GSAM organizations are in the PSB, or if the program accesses an attached subsystem, a partial backout is not possible. In that case, the SETS call is rejected with an SC status code. If the SETU call is used instead, it is not rejected because of unsupported PCBs, but returns an SC status code as a warning that the PSB contains unsupported PCBs and the function is not applicable to these unsupported PCBs.

Related Reading: For the status codes that are returned after the SETS call and the explanation of those status codes and the response required, see *IMS Version 15.3 Application Programming APIs*.

ROLS

The ROLS call backs out database changes to a processing point set by a previous SETS or SETU call, or to the prior commit point and returns the processed input messages to the message queue.

To back out database changes and message activity that have occurred since a prior SETS call, you issue the ROLS call using the I/O PCB and specifying an I/O area and token in the call. If the token does not match a token set by a preceding SETS call, an error status is returned. If the token does match the token of a preceding SETS call, the database updates made since this corresponding SETS call are backed

out, and all non-express messages inserted since the corresponding SETS are discarded. The ROLS call returns blanks if the call is processed, and returns a status code if an error or warning occurs. If you are using SETU with ROLS and have an external subsystem, the ROLS call will not be rejected, but an RC status code will be returned as a warning. All SETS points that were issued as part of the processing that was backed out are then canceled, and the existing database position for all supported PCBs is reset. For the ROLS call format, see the topic "ROLB Call" in *IMS Version 15.3 Application Programming APIs*.

The parameters for this form of the ROLS call are:

- The call function ROLS
- The name of the I/O PCB or AIB
- The name of the I/O area to receive the user data
- The name of an area containing the 4-byte token

Related reading: For the status codes that are returned after the ROLS call and the explanations of those status codes and the response required, see *IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes*.

Writing message-driven programs

A message-driven program is similar to an MPP: it retrieves messages and processes them, and it can read and update MSDBs, DEDBs, and full-function databases.

Message-driven programs can send messages to these destinations:

- The logical terminal that sent the input message, by issuing an ISRT call referencing the I/O PCB
- A different component of the physical terminal that sent the input message, by issuing an ISRT call referencing an alternate response PCB
- A different physical terminal from the one that sent the input message, by issuing an ISRT call referencing an alternate PCB

The message processing functions available to a message-driven program have some restrictions. These restrictions apply only to messages received or sent by the I/O PCB. The input message for a message-driven program must be a single segment message. Therefore, GU is the only call you can use to obtain the input message. The response message sent by the I/O PCB also must be a single segment message.

The transactions are in the response mode. This means that you must respond before the next message can be sent. You cannot use SPAs because a message-driven program cannot be a conversational program.

Not all of the system service calls are available. These system service calls are valid in a message-driven region:

CHKP (basic)
DEQ
INIT
LOG
SETS
ROLB
ROLS

However, other conditions might restrict their function in this environment. The options or calls issued using alternate terminal PCBs have no constraints.

Coding DC calls and data areas

The way you code DC calls and data areas depends on the application programming language you use.

Before coding your program

In addition to the information you need about the database processing that your program does, you need to know about message processing. Before you start to code, be sure you are not missing any of this information. Also, be aware of the standards at your installation that affect your program.

Information you need about your program's design:

- The names of the logical terminals that your program will communicate with
- The transaction codes, if any, for the application program's MPP skeleton to which your program will send messages
- The DC call structure for your program
- The destination for each output message that you send
- The names of any alternate destinations to which your program sends messages

Information you need about input messages:

- The size and layout of the input messages your program will receive (if possible)
- The format in which your program will receive the input messages
- The editing routine your program uses
- The range of valid data in input messages
- The type of data that input messages will contain
- The maximum and minimum length of input message segments
- The number of segments in a message

Information you need about output messages:

- The format in which IMS expects to receive output from your application program MPP skeleton
- The destination for the output messages
- The maximum and minimum length of output message segments

MPP code examples

Your MPP application can be written in assembler language, COBOL, C, Pascal, and PL/I.

In the following code examples, the programs do not have all the processing logic that a typical MPP has. The purpose of providing these programs is to show you the basic MPP structure in assembler language, COBOL, C language, Pascal, and PL/I. All the programs follow these steps:

1. The program retrieves an input message segment from a terminal by issuing a GU call to the I/O PCB. This retrieves the first segment of the message. Unless this message contains only one segment, your program issues GN calls to the I/O PCB to retrieve the remaining segments of the message. IMS places the input message segment in the I/O area that you specify in the call. In each of skeleton MPP examples, this is the MSG-SEG-IO-AREA.
2. The program retrieves a segment from the database by issuing a GU call to the DB PCB. This call specifies an SSA, SSA-NAME, to qualify the request. IMS places the database segment in the I/O area specified in the call. In this case, the I/O area is called DB-SEG-IO-AREA.
3. The program sends an output message to an alternate destination by issuing an ISRT call to the alternate PCB. Before issuing the ISRT call, the program must build the output message segment in an I/O area, and then the program specifies the I/O area in the ISRT call. The I/O area for this call is ALT-MSG-SEG-OUT.

The sample program is simplified for demonstration purposes; for example, the call to initiate sync point is not shown in the sample program. Include other IMS calls in a complete application program.

Coding your MPP program in assembler language

The coding conventions of an assembler language MPP are the same as those for a DL/I assembler program.

An assembler language MPP receives a PCB parameter list address in register 1 when it executes its entry statement. The first address in this list is a pointer to the TP PCB; the addresses of any alternate PCBs that the program uses come after the I/O PCB address, and the addresses of the database PCBs that the program uses follow. Bit 0 of the last address parameter is set to 1.

Coding your MPP program in C language

The program shown below is a skeleton MPP written in C language.

The numbers to the right of the program refer to the notes that follow the program. All storage areas that are referenced in the parameter list of your C language application program call to IMS can reside in the extended virtual storage area.

Skeleton MPP written in C

```

1  #pragma runopts(env(IMS),plist(IMS))
2  #include <ims.h>
3  #include <stdio.h>
4  /*
5  /*                                     ENTRY POINT
6  /*                                     */
7  main() {
8      static const char func_GU[4] = "GU ";
9      static const char func_ISRT[4] = "ISRT";
10     #define io_pcb ((IO_PCB_TYPE *)(_pcblist[0]))
11     #define alt_pcb (_pcblist[1])
12     #define db_pcb (_pcblist[2])
13     int rc;
14     #define io_pcb ((IO_PCB_TYPE *)(_pcblist[0]))
15     #define alt_pcb (_pcblist[1])
16     #define db_pcb (_pcblist[2])
17     rc = ctdli(func_GU, io_pcb, msg_seg_io_area);
18     rc = ctdli(func_GU, db_pcb, db_seg_io_area, ssa_name);
19     rc = ctdli(func_ISRT, alt_pcb, alt_msg_seg_out);
20 }
21 C language interface

```

Note:

1. The `env(IMS)` establishes the correct operating environment and the `plist(IMS)` establishes the correct parameter list, when invoked under IMS. The `ims.h` header file contains declarations for PCB layouts, `__pcblist`, and the `ctdli` routine. The PCB layouts define masks for the DB PCBs that the program uses as structures. These definitions make it possible for the program to check fields in the DB PCBs.

The `stdio.h` header file contains declarations for `sprintf`, which is useful for building SSAs.

2. After IMS has loaded the application program's PSB, IMS passes control to the application program through this entry point.
3. These are convenient definitions for the function codes and could be in one of your include files.

4. These could be structures, with no loss of efficiency.
5. The return code (status value) from DL/I calls can be returned and used separately.
6. The C language run-time sets up the __pcblist values. The order in which you refer to the PCBs must be the same order in which they have been defined in the PSB: first the TP PCB, then any alternate PCBs that your program uses, and finally the database PCBs that your program uses.
7. The program issues a GU call to the I/O PCB to retrieve the first message segment. You can leave out the rc =, and check the status in some other way.
8. The program issues a GU call to the DB PCB to retrieve a database segment. The function codes for these two calls are identical; the way that IMS identifies them is by the PCB to which each call refers.
9. The program then sends an output message to an alternate destination by issuing an ISRT call to an alternate PCB.
10. When there are no more messages for the program to process, the program returns control to IMS by returning from main or by calling exit().
11. IMS provides a language interface module (DFSLI000) that gives a common interface to IMS. This module must be made available to the application program at bind time.

Coding your MPP program in COBOL

The program shown below is a skeleton MPP in COBOL that shows the main elements of an MPP.

The numbers to the right of each part of the program refer to the notes that follow the program. If you plan to preload your IBM COBOL for z/OS & VM program, you must use the compiler option RENT. Alternatively, if you plan to preload your VS COBOL II program, you must use the compiler options RES and RENT. Enterprise COBOL programs version 4.2 or later must be compiled with the RENT compiler option to generate reentrant code. For more information about the RENT option, see [Compiling and linking COBOL programs for running under IMS](#).

If you want to use the IBM COBOL for z/OS & VM compiler to compile a program that is to execute in AMODE(31) on z/OS, you must use the compiler option RENT. Alternatively, if you want to use the VS COBOL II compiler to compile a program that is to execute in AMODE(31) on z/OS, you must use the compiler options RES and RENT. All storage areas that are referenced in the parameter lists of your calls to IMS can optionally reside in the extended virtual storage area.

IBM COBOL for z/OS & VM and VS COBOL II programs can coexist in the same application.

Important:

Skeleton MPP written in COBOL

NOTES:	
ENVIRONMENT DIVISION.	
.	
.	
.	
DATA DIVISION.	
WORKING-STORAGE SECTION.	1
77 GU-CALL PICTURE XXXX VALUE 'GU '.	
77 ISRT-CALL PICTURE XXXX VALUE 'ISRT'.	
77 CT PICTURE S9(5) COMPUTATIONAL VALUE +4.	
01 SSA-NAME.	
.	
01 MSG-SEG-IO-AREA.	2
.	
01 DB-SEG-IO-AREA.	
.	
01 ALT-MSG-SEG-OUT.	
.	
LINKAGE SECTION.	
01 IO-PCB.	3
.	
01 ALT-PCB.	
.	
01 DB-PCB.	
.	
PROCEDURE DIVISION USING IO-PCB, ALT-PCB, DB-PCB	4

CALL 'CBLTDLI' USING GU-CALL, IO-PCB, MSG-SEG-IO-AREA.	5
CALL 'CBLTDLI' USING GU-CALL, DB-PCB, DB-SEG-IO-AREA, SSA-NAME.	6
CALL 'CBLTDLI' USING ISRT-CALL, ALT-PCB, ALT-MSG-SEG-OUT.	7
GOBACK.	8
COBOL LANGUAGE INTERFACE	9

Note:

1. To define each of the call functions that your program uses, use a 77 or 01 level working-storage statement. Assign the value to the call function in a picture clause defined as four alphanumeric characters.
2. Use a 01 level working-storage statement for each I/O area that you will use for message segments.
3. In the linkage section of the program, use a 01 level entry for each PCB that your program uses. You can list the PCBs in the order that you list them in the entry statement, but this is not a requirement.
4. On the procedure statement, list the PCBs that your program uses in the order they are defined in the program's PSB: first the TP PCB, then any alternate PCBs, and finally the database PCBs that your program uses.
5. The program issues a GU call to the I/O PCB to retrieve the first segment of an input message.
6. The program issues a GU call to the DB PCB to retrieve the segment that would be described in the SSA-NAME area.
7. The program sends an output message segment to an alternate destination by using an alternate PCB.
8. When no more messages are available for your MPP to process, you return control to IMS by issuing the GOBACK statement.

If you compile all of your COBOL programs in the task with VS COBOL II, you can use the GOBACK statement with its normal COBOL-defined semantics.



Attention: The STOP RUN and EXIT PROGRAM statements are not supported. Using these statements might cause unpredictable results or abends.

9. If the COBOL compiler option NODYNAM is specified, you must link edit the language interface module, DFSLI000, with your compiled COBOL application program. If the COBOL compiler option DYNAM is specified, do not link edit DFSLI000 with your compiled COBOL program.

Coding your MPP program in Pascal

The program shown below is a skeleton MPP written in Pascal.

The numbers to the right of the program refer to the notes that follow the program. All storage areas that are referenced in the parameter list of your Pascal application program's call to IMS can reside in the extended virtual storage area.

Skeleton MPP written in Pascal

segment PASCIMS;	1
type	
CHAR4 = packed array [1..4] of CHAR;	2
CHARn = packed array [1..n] of CHAR;	
IOPCBTYPE = record	3
(* Field declarations *)	
end;	
ALTPCBTYPE = record	
(* Field declarations *)	
end;	
DBPCBTYPE = record	
(* Field declarations *)	
end;	

NOTES:


```

procedure PASCIMS (var SAVE: INTEGER;                                4
                   var IOPCB: IOPCBTYPE;
                   var ALTPCB: ALTPCBTYPE;
                   var DBPCB: DBPCBTYPE); REENTRANT;

procedure PASCIMS;                                                5
type
  SSATYPE    = record
    (* Field declarations *)
  end;

  MSG_SEG_IO_AREA_TYPE = record
    (* Field declarations *)
  end;

  DB_SEG_IO_AREA_TYPE = record
    (* Field declarations *)
  end;

  ALT_MSG_SEG_OUT_TYPE = record
    (* Field declarations *)
  end;
var
  MSG_SEG_IO_AREA : MSG_SEG_IO_AREA_TYPE;                          6
  DB_SEG_IO_AREA  : DB_SEG_IO_AREA_TYPE;
  ALT_MSG_SEG_OUT : ALT_MSG_SEG_OUT_TYPE;
const
  GU  = 'GU  ';                                                    7
  ISRT = 'ISRT';
  SSANAME = SSATYPE(...);
procedure PASTDLI; GENERIC;                                         8
begin
  PASTDLI(const GU,                                                9
          var IOPCB,
          var MSG_SEG_IO_AREA);
  PASTDLI(const GU,                                              10
          var DBPCB,
          var DB_SEG_IO_AREA,
          const SSANAME);
  PASTDLI(const ISRT,                                           11
          var ALTPCB,
          var ALT_MSG_SEG_OUT);
end;                                                                12
Pascal language interface                                         13

```

Note:

1. Define the name of the Pascal compile unit.
2. Define the data types needed for the PCBs used in your program.
3. Define the PCB data types used in your program.
4. Declare the procedure heading for the REENTRANT procedure called by IMS. The first word in the parameter list should be an INTEGER, which is reserved for VS Pascal's use, and the rest of the parameters will be the addresses of the PCBs received from IMS.
5. Define the data types needed for the SSAs and I/O areas.
6. Declare the variables used for the SSAs and I/O areas.
7. Define the constants (function codes, SSAs, and so forth) used in the PASTDLI DL/I calls.
8. Declare the IMS interface routine with the GENERIC Directive. GENERIC identifies external routines that allow multiple parameter list formats. A GENERIC routine's parameters are "declared" only when the routine is called.
9. The program issues a GU call to the I/O PCB to retrieve the first segment of an input message. The declaration of the parameters in your program might differ from this example.
10. The program can issue a GU call to a DB PCB to retrieve a database segment. The function codes for these two calls are identical; the way that IMS distinguishes between them is by the PCB to which each call refers. The declaration of the parameters in your program might differ from this example.
11. The program sends an output message segment to an alternate destination by issuing an ISRT call to an alternate PCB. The declaration of the parameters in your program might differ from this example.
12. When there are no more messages for your MPP to process, you return control to IMS by exiting the PASCIMS procedure. You can also code a RETURN statement to leave at another point.

13. You must bind your program to the IMS language interface module, DFSLI000, after you have compiled your program.

Coding your MPP program in PL/I

The following program is a skeleton MPP written in PL/I.

The numbers to the right of the program refer to the notes for the program. All storage areas that are referenced in the parameter list of your PL/I application program call to IMS can optionally reside in the extended virtual storage area.

If you plan to execute PL/I programs in 31-bit addressing mode, see Enterprise PL/I for z/OS Programming Guide.

Skeleton MPP written in PL/I

	NOTES
/*	*/
/*	*/
/*	*/
UPDMASTER: PROCEDURE (IO_PTR, ALT_PTR, DB_PTR)	1
OPTIONS (MAIN);	
DCL FUNC_GU CHAR(4) INIT('GU');	2
DCL FUNC_ISRT CHAR(4) INIT('ISRT');	
.	
DCL SSA_NAME...;	
.	
DCL MSG_SEG_IO_AREA CHAR(n);	3
DCL DB_SEG_IO_AREA CHAR(n);	
DCL ALT_MSG_SEG_OUT CHAR(n);	
.	
DCL 1 IO_PCB BASED (IO_PTR),...;	4
DCL 1 ALT_PCB BASED (ALT_PTR),...;	
DCL 1 DB_PCB BASED (DB_PTR),...;	
.	
DCL THREE FIXED BINARY(31) INIT(3);	5
DCL FOUR FIXED BINARY(31) INIT(4);	
DCL PLITDLI ENTRY EXTERNAL;	
.	
CALL PLITDLI (THREE, FUNC_GU, IO_PTR, MSG_SEG_IO_AREA);	6
.	
CALL PLITDLI (FOUR, FUNC_GU, DB_PTR, DB_SEG_IO_AREA,	7
SSA_NAME);	
.	
CALL PLITDLI (THREE, FUNC_ISRT, ALT_PTR, ALT_MSG_SEG_OUT);	8
.	
END UPDMASTER;	9
PL/I LANGUAGE INTERFACE	10

Note:

1. This is the standard entry point to a PL/I Optimizing Compiler MPP. This statement includes a pointer for each PCB that the MPP uses. You must refer to the PCBs in the same order as they are listed in the PSB: first the TP PCB, then any alternate PCBs that your program uses, and finally the database PCBs that your program uses.
2. The program defines each call function that it uses in its data area. In PL/I, you define the function codes as character strings and assign the appropriate values to them.
3. Define PCB Masks as major structures based on the addresses passed in the PROCEDURE statement. Although not shown in the example, you will code the appropriate additional fields in the structure, depending on the type of PCB to which the mask is associated.
4. To define your PCBs, use major structure declarations.
5. PL/I calls have a parameter that is not required in COBOL programs or assembler language programs. This is the parmcount, and it is always the first parameter. You define the values that your program will need for the parmcount in each of its calls. The parmcount gives the number of parameters that follow parmcount itself.
6. The program issues a GU call to the I/O PCB to retrieve the first message segment.

7. The program can issue a GU call to a DB PCB to retrieve a database segment. The function codes for these two calls are identical; the way that IMS distinguishes between them is by the PCB to which each call refers.
8. The program then sends an output message to an alternate destination by issuing an ISRT call to an alternate PCB.
9. When there are no more messages for the program to process, the program returns control to IMS by issuing the END statement or the RETURN statement.
10. You must bind your program to the IMS language interface module, DFSLI000, after you have compiled your program.

Message processing considerations for DB2

For the most part, the message processing function of a dependent region that accesses DB2 databases is similar to that of a dependent region that accesses only DL/I databases. The method each program uses to retrieve and send messages and back out database changes is the same.

The differences are:

- DL/I statements are coded differently from SQL (structured query language) statements.
- When an IMS TM application program receives control from IMS TM, IMS has already acquired the resources the program is able to access. IMS TM schedules the program, although some of the databases are not available. DB2 does not allocate resources for the program until the program issues its first SQL statement. If DB2 cannot allocate the resources your program needs, your program can optionally receive an initialization error when it issues its first SQL call.
- When an application issues a successful checkpoint call or a successful message GU call, DB2 closes any cursors that the program is using. This means that your program should issue its OPEN CURSOR statement after a checkpoint call or a message GU.

IMS TM and DB2 work together to keep data integrity in these ways:

- When your program reaches a commit point, IMS TM makes any changes that the program has made to DL/I databases permanent, releases output messages for their destinations, and notifies DB2 that the program has reached a commit point. DB2 then makes permanent any changes that the program has made to DB2 databases.
- When your program terminates abnormally or issues one of the IMS TM rollback calls (ROLB, ROLS without a token, or ROLL), IMS TM cancels any output messages your program has produced, backs out changes your program has made to DL/I databases since the last commit point, and notifies DB2. DB2 backs out the changes that the program has made to DB2 databases since the last commit point.

Through the Automated Operator Interface (AOI), IMS TM application programs can issue DB2 commands and IMS TM commands. To issue DB2 commands, the program issues the IMS TM /SSR command followed by the DB2 command. The output of the /SSR command is routed to the master terminal operator (MTO).

Chapter 26. IMS Spool API

The IMS Spool API support provides feedback to the application program when IMS detects errors in the print data set options of the CHNG and SETO calls.

For convenience, your application program can display these errors by sending a message to an IMS printer or by performing another action that lets you examine the parameter lists and feedback area without looking at a dump listing. This information applies only to the calls as they are used with Spool API support.

Managing the IMS Spool API overall design

The IMS Spool API (application programming interface) is an expansion of the IMS application program interface that allows applications to interface directly to JES and create print data sets on the job entry subsystem (JES) spool. These print data sets can then be made available to print managers and spool servers to serve the needs of the application.

IMS Spool API design

The IMS Spool API design provides the application program with the ability to create print data sets on the JES spool using the standard DL/I call interface.

The functions provided are:

- Definition of the data set output characteristics
- Allocation of the data set
- Insertion of lines of print into the data set
- Closing and deallocation of the data set
- Backout of uncommitted data within the limits of the JES interface
- Assistance in controlling an in-doubt print data set

The IMS Spool API support uses existing DL/I calls to provide data set allocation information and to place data into the print data set. These calls are:

- The CHNG call. This call is expanded so that print data set characteristics can be specified for the print data set that will be allocated. The process uses the alternate PCB as the interface block associated with the print data set.
- The ISRT call. This call is expanded to perform dynamic allocation of the print data set on the first insert, and to write data to the data set. The data set is considered in-doubt until the unit of work (UOW) terminates. If possible, the sync point process deletes all in-doubt data sets for abending units of work and closes and deallocates data sets for normally terminating units of work.
- The SETO call. This is a call, SETO (Set Options), introduced by this support. Use this call to create dynamic output text units to be used with the subsequent CHNG call. If the same output descriptor is used for many print data sets, the overhead can be reduced by using the SETO call to prebuild the text units necessary for the dynamic output process.

Sending data to the JES spool data sets

Application programs can send data to the JES spool data sets using the same method that is used to send output to an alternate terminal. Use the DL/I call to change the output destination to a JES spool data set.

Use the DL/I ISRT or PURG call to insert a message.

The options list parameter on the CHNG and SETO calls contains the data set printer processing options. These options direct the output to the appropriate IMS Spool API data set. These options are validated for the DL/I call by the MVSScheduler JCL Facility (SJF). If the options are invalid, error codes are returned to the application. To receive the error information, the application program specifies a feedback area in the

CHNG or SETO DL/I call parameter list. If the feedback area is present, information about the options list error is returned directly to the application.

IMS Spool API performance considerations

The IMS Spool API interface uses z/OS services within an IMS application while minimizing the performance impact of the z/OS services on the other IMS transactions and services.

For this reason, the IMS Spool API support places the print data directly on the JES spool at insert time instead of using the IMS message queue for intermediate storage. The processing of IMS Spool API requests is performed under the TCB of the dependent region to ensure maximum usage of N-way processors. This design reduces the error recovery and JES job orientation problems.

JES initiator considerations

Because the dependent regions are normally long-running jobs, some of the initiator or job specifications might must be changed if the dependent region is using the IMS Spool API.

You might need to limit the amount of JES spool space used by the dependent region to contain the dynamic allocation and deallocation messages. For example, you can use the JOB statement MSGLEVEL to eliminate the dynamic allocation messages from the job log for the dependent region. You might be able to eliminate these messages for dependent regions executing as z/OS started tasks.

Another initiator consideration is the use of the JES job journal for the dependent region. If the job step has a journal associated with it, the information for z/OS checkpoint restart is recorded in the journal. Because IMS dependent regions cannot use z/OS checkpoint restart, specify JOURNAL=NO for the JES2 initiator procedure and the JES3 class associated with the dependent regions execution class. You can also specify the JOURNAL= on the JES3 //*MAIN statement for dependent regions executing as jobs.

Application managed text units

The application can manage the dynamic descriptor text units instead of IMS. If the application manages the text units, overhead for parsing and text unit build can be reduced.

Use the SETO call to have IMS build dynamic descriptor text units. After they are built, these text units can be used with subsequent CHNG calls to define the print characteristics for a data set.

To reduce overhead by managing the text units, the text units should be used with several change calls. An example of this is a wait-for-input (WFI) transaction. The same data set attributes can be used for all print data sets. For the first message processed, the application uses the SETO call to build the text units for dynamic descriptors and a subsequent CHNG call with the TXTU= parameter referencing the prebuilt text units. For all subsequent messages, only a CHNG call using the prebuilt text units is necessary.

Note: No testing has been done to determine the amount of overhead that might be saved using prebuilt text units.

BSAM I/O area

The I/O area for spool messages can be very large. It is not uncommon for the area to be 32 KB in length. To reduce the overhead incurred with moving large buffers, IMS attempts to write to the spool data set from the application's I/O area.

BSAM does not support I/O areas in 31-bit storage for SYSOUT files. If IMS finds that the application's I/O area is in 31-bit storage:

- A work area is obtained from 24-bit storage.
- The application's I/O area is moved to the work area.
- The spool data set is written from the work area.

If the application's I/O area can easily be placed in 24-bit storage, the need to move the I/O area can be avoided and possible performance improvements achieved.

Note: No testing has been done to determine the amount of performance improvement possible.

Since a record can be written by BSAM directly from the application's I/O area, the area must be in the format expected by BSAM. The format must contain:

- Variable length records
- A Block Descriptor Word (BDW)
- A Record Descriptor Word (RDW)

IMS Spool API application coding considerations

Your application can send data to a JES Spool or Print server using a print data set. You can set options for message integrity and recovering data when failures occur.

Print data formats

The IMS Spool API attempts to provide a transparent interface for the application to insert data to the JES spool. The data can be in line, page, IPDS, AFPDS, or any format that can be handled by a JES Spool or Print server that processes the print data set. The IMS Spool API does not translate or otherwise modify the data inserted to the JES spool.

Message integrity options

The IMS Spool API provides support for message integrity.

This is necessary because IMS cannot properly control the disposition of a print data set when:

- IMS abnormal termination does not execute because of a hardware or software problem.
- A dynamic deallocation error exists for a print data set.
- Logic errors are in the IMS code.

In these conditions, IMS might not be able to stop the JES subsystem from printing partial print data sets. Also, the JES subsystems do not support a two-phase sync point.

Print disposition

The most common applications using Advanced Function Printing (AFP) are TSO users and batch jobs. If any of these applications are creating print data sets when a failure occurs, the partial print data sets will probably print and be handled in a manual fashion. Many IMS applications creating print data sets can manage partial print data sets in the same manner. For those applications that need more control over the automatic printing by JES of partial print data sets, the IMS Spool API provides the following integrity options. However, these options alone might not guarantee the proper disposition of partial print data sets. These options are the b variable following the IAFP keyword used with the CHNG call.

b=0

Indicates no data set protection

This is probably the most common option. When this option is selected, IMS does not do any special handling during allocation or deallocation of the print data set. If this option is selected, and any condition occurs that prevents IMS from properly disposing the print data set, the partial data set probably prints and must be controlled manually.

b=1

Indicates SYSOUT HOLD protection

This option ensures that a partial print data set is not released for printing without a JES operator taking direct action. When the data set is allocated, the allocation request indicates to JES that this print data set be placed in SYSOUT HOLD status. The SYSOUT HOLD status is maintained for this data set if IMS cannot deallocate the data set for any reason. Because the print data set is in HOLD status, a JES operator must identify the partial data set and issue the JES commands to delete or print this data set.

If the print data set cannot be deleted or printed:

- Message DFS0012I is issued when a print data set cannot be deallocated.
- Message DFS0014I is issued during IMS emergency restart when an in-doubt print data set is found. The message provides information to help the JES operator find the proper print data set and effect the proper print disposition.

Some of the information includes:

- JOBNAME
- DSNAME
- DDNAME
- A recommendation on what IMS believes to be the proper disposition for the data set (for example, printing or deleting).

By using the Spool Display and Search Facility (SDSF), you can display the held data sets, identify the in-doubt print data set by DDNAME and DSNAME, and issue the proper JES command to either delete or release the print data set.

b=2

Indicates a non-selectable destination

This option prevents the automatic printing of partial print data sets. The IMS Spool API function requests a remote destination of IMSTEMP for the data set when the data set is allocated. The JES system must have a remote destination of IMSTEMP defined so that JES does not attempt to print any data sets that are sent to the destination.

If b=2, the name of the remote destination for the print data set must be specified in the destination name field of the call parameter list when the CHNG call is issued. When IMS deallocates the data set at sync point, and the data set prints, IMS requests that the data set be transferred to the requested final remote destination.

If the remote destination is not defined to the JES system, a dynamic allocation failure occurs. Because this remote destination is defined as non-selectable, and if IMS is unable to deallocate the print data set and control its proper disposition, the print data set remains associated with remote destination IMSTEMP when deallocated by z/OS.

When an deallocation error occurs, message DFS0012I is issued to provide details of the deallocation error and help identify the print data set that requires operator action. When partial print data sets are left on this special remote destination, the JES operator can display all the print data sets associated with this JES destination to locate the data set that requires action. The b=2 option simplifies the operator's task of locating partial print data sets.

Message options

The third option on the IAPF keyword controls informational messages issued by the IMS Spool API support. These messages inform the JES operator of in-doubt data sets that need action.

c=0

Indicates that no DFS0012I or DFS0014I messages are issued for the print data set. You can specify c=0 only if b=0 is specified.

c=m

Indicates that DFS0012I and DFS0014I messages are issued if necessary. You can specify c=m or if b=1 or if b=2, it is the default.

Option c does not affect issuing message DFS0013E.

IMS emergency restart

When IMS emergency restart is performed, DFS0014I messages might be issued if IMS finds that the proper disposition of a print data set is in-doubt, as a result of the restart. This message is only issued if the message option for the print data set was requested or c=m on the IAPF variable. When a DFS0014I message is received, a JES operator might need to find and properly dispose of the print data set. The DFS0014I message provides a recommended disposition (that is, deletion or printing).

Destination name (LTERM) usage

The standard CHNG call parameter list contains a destination name field. For traditional message calls, this field contains the LTERM or transaction code that becomes the destination of messages sent using this alternate PCB. When ISRT calls are issued against the PCB, the data is sent to the LTERM or transaction.

However, the destination name field has no meaning to the IMS Spool API function unless b=2 is specified following the IAFP keyword.

When b=2 is specified:

- The name must be a valid remote destination supported by the JES system that receives the print data sets.
- If the name is not a valid remote destination, an error occurs during dynamic deallocation.

If any option other than 2 is selected, the name is not used by IMS.

The LTERM name appears in error messages and log records. Use a name that identifies the routine creating the print data set. This information can aid in debugging application program errors.

Understanding parsing errors

When you are diagnosing multiple parsing error return codes, the first code returned is usually the most informative.

Keywords

The CHNG and SETO calls have two types of keywords. The type of keyword determines what type of keyword validation IMS should perform. The keyword types are:

- Keywords valid for the calls (for example, IAFP, PRT0, TXTU, and OUTN)
- Keywords valid as operands of the PRT0 keyword (for example CLASS and FORMS).

Incorrectly specified length fields can cause errors when IMS checks for valid keywords. When IMS is checking the validity of keywords on the CHNG and SETO calls, one set of keywords is valid. When IMS is checking the validity of keywords on the PRT0 keyword, another set of keywords is valid. For this reason, incorrectly specified length fields can cause a scan to terminate prematurely, and keywords that appear to be valid are actually invalid because of where they occur in the call list. IMS might report that a valid keyword is invalid if it detects a keyword with an incorrect length field or a keyword that occurs in the wrong place in the call list.

Status codes

The status code returned for the call can also suggest the location of the error. Although exceptions exist, generally, an AR status code is returned when the keyword is invalid for the call. An AS status code is returned when the keyword is invalid as a PRT0 option.

Error codes

This topic contains information on Spool API error codes that your application program can receive. The topic "Diagnosis examples" contains examples of errors and the resulting error codes provided to the application program.

Error Code

Reason

(0002)

Unrecognized option keyword.

Possible reasons for this error are:

- The keyword is misspelled.

- The keyword is spelled correctly but is followed by an invalid delimiter.
- The length specified field representing the PRT0 is shorter than the actual length of the options.
- A keyword is not valid for the indicated call.

(0004)

Either too few or too many characters were specified in the option variable. An option variable following a keyword in the options list for the call is not within the length limits for the option.

(0006)

The length field (LL) in the option variable is too large to be contained in the options list. The options list length field (LL) indicates that the options list ends before the end of the specified option variable.

(0008)

The option variable contains an invalid character or does not begin with an alphabetic character.

(000A)

A required option keyword was not specified.

Possible reasons for this error are:

- One or more additional keywords are required because one or more keywords were specified in the options list.
- The specified length of the options list is more than zero but the list does not contain any options.

(000C)

The specified combination of option keywords is invalid. Possible causes for this error are:

- The keyword is not allowed because of other keywords specified in the options list.
- The option keyword is specified more than once.

(000E)

IMS found an error in one or more operands while it was parsing the print data set descriptors. IMS usually uses z/OS services (SJF) to validate the print descriptors (PRT0= option variable). When IMS calls SJF, it requests the same validation as for the TSO OUTDES command. Therefore, IMS is insensitive to changes in output descriptors. Valid descriptors for your system are a function of the MVS release level. For a list of valid descriptors and proper syntax, use the TSO HELP OUTDES command.

IMS must first establish that the format of the PRT0 options is in a format that allows the use of SJF services. If it is not, IMS returns the status code AS, the error code (000E), and a descriptive error message. If the error is detected during the SJF process, the error message from SJF will include information of the form (R.C.=xxxx,REAS.=yyyyyyyy), and an error message indicating the error.

Related reading: For more information on SJF return and reason codes, see *z/OS MVS Programming: Authorized Assembler Services Guide*.

The range of some variables is controlled by the initialization parameters. Values for the maximum number of copies, allowable remote destination, classes, and form names are examples of variables influenced by the initialization parameters.

Diagnosis examples

The following examples illustrate mistakes that can generate the various spool API error codes, and diagnosis of the problems.

Some length fields are omitted when they are not necessary to illustrate the example. The feedback and options lists that are shown on multiple lines are contiguous.

Error code (0002)

Two examples of the error code 0002 are shown in this section.

For the first example the options list contains both the keywords PRTO and TXTU. The keyword, TXTU, is invalid for the SETO call.

```
CALL = SETO
  OPTIONS LIST = PRTO=04DEST(018),CLASS(A),TXTU=SET1
  FEEDBACK = TXTU(0002)
  STATUS CODE = AR
```

For the second example, the length field of the PRTO options is too short to contain all of the options. This means that IMS finds the COPIES and FORMS keywords outside the PRTO options list area and indicates that they are invalid on the CHNG call.

```
CALL = CHNG
  OPTIONS LIST = IAFP=N0M,PRTO=0FDEST(018),LINECT(200),CLASS(A),
                COPIES(80),FORMS(ANS)
  FEEDBACK = COPIES(0002),FORMS(0002)
  STATUS CODE = AR
```

Error code (0004)

For this example, the operand for the OUTN keyword is 9 bytes long and exceeds the maximum value for the OUTPUT JCL statement.

```
CALL = CHNG
  OPTIONS LIST = IAFP=N0M,OUTN=OUTPUTDD1
  FEEDBACK = OUTN(0004)
  STATUS CODE = AR
```

Error code (0006)

The length of the options list for this call is too short to contain all of the operands of the PRTO keyword.

This example shows an options list that is X'48' bytes long and is the correct length. The length field of the PRTO keyword incorrectly indicates a length of X'5A'. The length of the PRTO options exceeds the length of the entire options list so IMS ignores the PRTO keyword and scans the rest of the options list for valid keywords. The feedback area contains the PRTO(0006) code (indicating a length error) and the (0002) code (indicating that the PRTO keywords are in error). This is because the keywords beyond the first PRTO keyword, up to the length specified in the options list length field, have been scanned in search of valid keywords for the call. The status code of AR indicates that the keywords are considered invalid for the call and not the PRTO keyword.

```
CALL = CHNG
  OPTIONS LIST = 0400          05
                0800IAFP=N0M,PRTO=0ADEST(018),LINECT(200),CLASS(A),
                COPIES(3),FORMS(ANS)
  FEEDBACK = PRTO(0006),LINECT(0002),CLASS(0002),COPIES(0002),
                FORMS(0002)
  STATUS CODE = AR
```

Error code (0008)

In this example, the message option of the IAFP keyword is incorrectly specified as "Z".

```
CALL = CHNG
  OPTIONS LIST = IAFP=N0Z,PRTO=0BDEST(018)
  FEEDBACK = IAFP(0008) INVALID VARIABLE
  STATUS CODE = AR
```

Error code (000A)

In this example, the valid keyword TXTU is specified, but the call also requires that the IAFP keyword be specified if the TXTU keyword is used.

```
CALL = CHNG
  OPTIONS LIST = TXTU=SET1
  FEEDBACK = TXTU(000A)
  STATUS CODE = AR
```

Error code (000C)

The AR status code is returned with the (000C) error code. This implies that the problem is with the call options and not with the PRTO options.

The call options list contains the PRTO and TXTU keywords. These options cannot be used in the same options call list.

```
CALL = CHNG
  OPTIONS LIST = IAFP=A00,PRTO=0BCOPIES(3),TXTU=SET1
  FEEDBACK = TXTU(000C)
  STATUS CODE = AR
```

Error code (000E)

In this example, the COPIES parameter has the incorrect value "RG" specified as one of its operands. The error message indicates that the values for these operands must be numeric.

```
CALL = CHNG
  OPTIONS LIST = IAFP=A00,PRTO=0BCOPIES((3),(8,RG,18,80))
  FEEDBACK = PRTO(000E) (R.C.=0004,REAS.=00000204) COPIES/RG VALUE
                                MUST BE NUMERIC CHARACTERS
  STATUS CODE = AS
```

This example includes an invalid PRTO operand. The resulting reason code of X'000000D0' indicates that the XYZ operand is invalid.

```
CALL = CHNG
  OPTIONS LIST = IAFP=A00,PRTO=0AXYZ(018)
  FEEDBACK = PRTO(000E) (R.C.=0004,REAS.=000000D0) XYZ
  STATUS CODE = AS
```

Understanding allocation errors

The IMS Spool API interface defers dynamic allocation of the print data set until data is actually inserted into the data set. Incorrect data set print options on the CHNG or SETO call can cause errors during dynamic allocation. The print data set options can be parsed during the processing of the CHNG and SETO calls but some things, for example the *destination name* parameter, can be validated only during dynamic allocation.

If one of the print options is incorrect and dynamic allocation fails when the IMS performs the first insert for the data set, IMS returns a AX status code to the ISRT call. IMA also issues message DFS0013E and writes a diagnostic log record (67D0) that you can use to evaluate the problem. The format of the error message indicates the type of service that was invoked and the return and reason codes that were responsible for the error. The error message can indicate these services:

DYN

MVS dynamic allocation (SVC99)

OPN

MVS data set open

OUT

MVS dynamic output descriptors build (SVC109)

UNA

MVS dynamic unallocation (SVC99)

WRT

MVS BSAM write

If the DFS0013E message indicates an error return code from any of these services, you should consult the corresponding MVS documentation for more information on the error code. If the service is for dynamic allocation, dynamic unallocation, or dynamic output descriptor build, see *z/OS MVS Programming: Authorized Assembler Services Guide* for the appropriate return and reason codes.

One common mistake is the use of an invalid destination or selection of integrity option 2 (non-selectable destination) when the destination of IMSTEMP has not been defined to JES. If you specify an invalid destination in the **destination name** parameter, the call will result in a dynamic unallocation error when IMS unallocates the print data set.

Understanding dynamic output for print data sets

IMS can use the z/OS services for Dynamic Output (SVC109) for print data sets. IMS uses this service to specify the attributes provided by the application for the print data sets being created. The service can be used on the CHNG call with the PRT0, TXTU, and OUTN options.

Related reading: For more information, see *z/OS MVS Programming: Assembler Services Guide*.

CHNG call with PRT0 option

When you use the CHNG call and PRT0 option, IMS activates SJF to verify the print options to call z/OS services for Dynamic Output. This creates the output descriptors that are used when the print data set is allocated. This is the simplest way for the application to provide print data set characteristics. However, it also uses the most overhead because parsing must occur for each CHNG call. If your application is WFI or creates multiple data sets with the same print options, use another option to reduce the parsing impact. You must specify the IAFP option keyword with this option.

CHNG call with TXTU option

If your application can manage the text units necessary for Dynamic Output, then you can avoid parsing for many of the print data sets. You can do this in one of two ways:

- The application can build the text unit in the necessary format within the application area and pass these text units to IMS with the CHNG call and TXTU option.
- The application can provide the print options to IMS with a SET0 call and provide a work area for the construction of the text units. After z/OS has finished parsing and text construction, the work area passed will contain the text units necessary for Dynamic Output after a successful SET0 call. The application must not relocate this work area because the work area contains address sensitive information.

Regardless of the method the application uses to manage the text units, applications that can reuse the text units can often achieve better performance by using the TXTU option on the CHNG call.

You must specify the IAFP option keyword with this option.

CHNG call with OUTN option

The dependent region JCL can contain OUTPUT JCL statements. If your application can use this method, you can use the CHNG call and OUTN option to reference OUTPUT JCL statements. When you use the OUTN option, IMS will reference the OUTPUT JCL statements at dynamic allocation. JES will obtain the print data set characteristics from the OUTPUT JCL statement. You must specify the IAFP option keyword with this option.

Sample programs using the Spool API

The Spool API provides functions that allow an application program to write data to the IMS Spool using the same techniques for sending data to native IMS printers.

The Spool API provides functions such as error checking for invalid OUTDES parameters. Error checking makes application programs more complex. To simplify these application programs, develop a common routine to manage error information, then make the diagnostic information from the Spool API available for problem determination.

The sample programs in this section shows how DL/I calls can be coded to send data to the IMS Spool. Only the parts of the application program necessary to understand the DL/I call formats are included. The examples are in assembler language.

Application PCB structure

The application PCBs are as follows:

I/O PCB
ALTPCB1
ALTPCB2
ALTPCB3
ALTPCB4

GU call to I/O PCB

IMS application programs begin with initialization and a call to the I/O PCB to obtain the input message. The following code example shows how to issue a GU call to the I/O PCB.

After completing the GU call to the I/O PCB, the application program prepares output data for the IMS Spool.

Issuing a GU call to the I/O PCB

```
*****
*      ISSUE GU ON IOPCB                      *
*****
      L      9,IOPCB                      I/O PCB ADDRESS
      LA     9,0(9)
      MVC    FUNC,=CL4'GU'                GU FUNCTION
      CALL   ASMTDLI,(FUNC,(9),IOA1),VL
      BAL    10,STATUS                     CHECK STATUS
*      ADDITIONAL PROGRAM LOGIC HERE
FUNC   DC    CL4' '
IOA1   DC    AL2(IOA1LEN),AL2(0)
TRAN   DS    CL8                          TRANSACTION CODE AREA
DATA   DS    CL5                          DATA STARTS HERE
      DC    20F'0'
IOA1LEN EQU  *-IOA1
```

CHNG call to alternate PCB

In the same way that other programs specify the destination of the output using the CHNG call, this program specifies the IMS Spool as the output destination. For a native IMS printer, the DEST NAME parameter identifies the output LTERM name. When a CHNG call is issued that contains the IAFP= keyword, the DEST NAME parameter is used only if integrity option '2' is specified. If option '2' is not specified, the DEST NAME parameter can be used by the application program to identify something else, such as the routine producing the change call. The destination for the print data set is established using a combination of initialization parameters or OUTDES parameters.

The following code example shows how to issue a CHNG call to the alternate modifiable PCB.

After the CHNG call is issued, the application program creates the print data set by issuing ISRT calls.

Issuing a CHNG call to the alternate modifiable PCB

```
*****
*          ISSUE CHNG ON ALTPCB4                      *
*****
      L      9,ALTPCB4          ALT MODIFIABLE PCB
      LA      9,0(9)           CLEAR HIGH BYTE/BIT
      MVC     FUNC,=CL4'CHNG'   CHNG FUNCTION
      CALL    ASMTDLI,(FUNC,(9),DEST2,OPT1,FBA1),VL
      BAL     10,STATUS         CHECK STATUS OF CALL
*      ADDITIONAL PROGRAM LOGIC HERE
FUNC    DC     CL4' '
DEST2   DC     CL8'IAFP1'      LTERM NAME
*
      DC      C'OPT1'          OPTIONS LIST AREA
OPT1    DC     AL2(OPT1LEN),AL2(0)
      DC     C'IAFP='
OCC     DC     C'M'           DEFAULT TO MACHINE CHAR
OOPT    DC     C'1'           DEFAULT TO HOLD
OMSG    DC     C'M'           DEFAULT TO ISSUE MSG
      DC     C' '
      DC     C'PRT0='
PRT01   EQU    *
      DC     AL2(PRT01LEN)
      DC     C'COPIES(2),CLASS(T),DEST(RMT003)'
PRT01LEN EQU    *-PRT01
      DC     C' '
OPT1LEN EQU    *-OPT1
*
FBA1    DC     AL2(FBA1LEN),AL2(0)
      DC     CL40' '
FBA1LEN EQU    *-FBA1
```

ISRT call to alternate PCB

Once the IMS Spool is specified as the destination of the PCB, ISRT calls can be issued against the alternate PCB.

The following code example shows how to issue the ISRT call to the alternate modifiable PCB.

The print data streams can be stored in databases or generated by the application, depending on the requirements of the application program and the type of data set being created.

Issuing an ISRT call to the alternate modifiable PCB

```
*****
*          ISSUE ISRT TO ALTPCB4                      *
*****
      L      9,ALTPCB4          ALT MODIFIABLE PCB
      LA      9,0(9)           CLEAR HIGH BYTE/BIT
      MVC     FUNC,=CL4'ISRT'   ISRT FUNCTION
      CALL    ASMTDLI,(FUNC,(9),IOA2),VL
      BAL     10,STATUS         CHECK STATUS OF CALL
*      ADDITIONAL PROGRAM LOGIC HERE
FUNC    DC     CL4' '
IOA2    DC     AL2(IOA2LEN),AL2(0)
IOA21   DC     AL2(MSG2LEN),AL2(0)
      DC     C' '             CONTROL CHARACTER
      DC     C'MESSAGE TO SEND TO IMS SPOOL'
MSG2LEN EQU    *-IOA21
IOA2LEN EQU    *-IOA2
```

Program termination

After the calls are issued, the program sends a message back to originating terminal, issues a GU call to the I/O PCB, or terminates normally.

Chapter 27. IMS Message Format Service

The IMS Message Format Service (MFS) is a facility of the IMS Transaction Manager environment that formats messages to and from terminal devices, so that IMS application programs do not deal with device-specific characteristics in input or output messages.

In addition, MFS formats messages to and from user-written programs in remote controllers and subsystems, so that application programs do not deal with transmission-specific characteristics of the remote controller.

MFS uses control blocks you specify to indicate to IMS how input and output messages are arranged.

- For input messages, MFS control blocks define how the message 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 sent by the application program to the device is arranged on the screen or at the printer. Data that appears on the screen but not in the program's I/O area, such as a literal, can also be defined.

In IMS Transaction Manager systems, data passing between the application program and terminals or remote programs can be edited by MFS or basic edit. Whether an application program uses MFS depends on the type of terminals or secondary logical units (SLUs) your network uses.

Restriction: MFS does not support message formatting for LU 6.2 devices.

Advantages of using MFS

By using MFS, you can simplify the developing and maintaining of terminal-oriented applications, and improve online performance by using control blocks for online processing.

Simplify development and maintenance

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 using their different capabilities. Thus, MFS can minimize the number of required 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 receives a message from a terminal, how the message appears in the program's I/O area is independent of what kind of terminal sent it; it depends on the MFS options specified for the program. If the next message the application receives is from a different type of terminal, you do not need to do anything to the application. MFS shields the application from the physical device that is sending the message in the same way that a DB program communication block (PCB) shields a program from what the data in the database actually looks like and how it is stored.

Other common functions performed by MFS include left or right justification of data, padding, exits for validity checking, time and date stamping, page and message numbering, and data sequencing and segmenting. When MFS assumes 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 formatting the application program data for presentation to the output device or remote program.




Figure 82. Message formatting using MFS

Improve online performance of a terminal

MFS also improves online performance of a terminal-oriented IMS by using control blocks designed for online processing. The MFS control blocks are compiled offline, when the IMS Transaction Manager system 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 to reduce CPU and 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 real storage indexing and anticipatory fetching of the control blocks can also reduce response time. Further performance improvements can be gained when IMS is generated for z/OS, since multiple I/O operations can execute concurrently to load the format blocks from the MFS format library.

In addition, MFS uses z/OS paging services; this helps to reduce page faults by the IMS control region task.

MFS can reduce use of communication lines by compressing data and transmitting only required data. This reduces line load and improves both response time and device performance.

MFS control blocks

There are four types of MFS control blocks that you specify to format input and output for the application program and the terminal or remote program.

The four types are:

Message Output Descriptors (MODs)

Define the layout of messages MFS receives from the application program.

Device Output Formats (DOFs)

Describe how MFS formats messages for each of the devices the program communicates with.

Device Input Formats (DIFs)

Describe the formats of messages MFS receives from each of the devices the program communicates with.

Message Input Descriptors (MIDs)

Describe how MFS further formats messages so that the application program can process them.

Throughout this information, the term "message descriptors" refers to both MIDs and MODs. The term "device formats" refers to both DIFs and DOFs.

Each MOD, DOF, DIF and MID deals with a specific message. There must be a MOD and DOF for each unique message a program sends, and a DIF and MID for each unique message a program receives.

MFS examples

One way to understand the relationship between the MFS control blocks is to look at a message from the time a user enters it at the terminal to the time the application program processes the message and sends a reply back to the terminal. Though MFS can be used with both display terminals and printer devices, for clarity in this example, a display terminal is being used.

The following figure shows the relationships between the MFS control blocks.




Figure 83. MFS control block relationships

Looking at payroll records

Suppose your installation has a message processing program used to view employee payroll records. From a display terminal, issue the IMS format command (/FORMAT), and the MOD name. This formats the screen in the way defined by the MOD written by the MFS programmer. When you enter the MOD name, the screen contains only literals and no output data from the application program. At this stage, no application program is involved. (For more information about /FORMAT, see *IMS Version 15.3 Commands, Volume 1: IMS Commands A-M*.)

In this example, suppose the name of the MOD that formats the screen for this application is PAYDAY. Enter this command:

```
/FORMAT PAYDAY
```

IMS locates the MFS MOD control block with the name PAYDAY and arranges the screen in the format defined by the DOF.

The following figure shows how this screen looks.

```

*EMPLOYEE PAYROLL*
*****

FIRST NAME:
EMPLOYEE NO:

LAST NAME:

INPUT:

```

Figure 84. PAYDAY screen, formatted by DOF

The DOF defines a terminal format that asks you to give the employee's name and employee number. PAYUP is the transaction code associated with the application that processes this information. When you enter the MOD name, the transaction code is included in the first screen format displayed. This means that you do not need to know the name of the program that processes the data; you only need the name of the MOD that formats the screen.

After the screen format is displayed, you can enter the information. There are four stages to sending a message to the program and receiving the reply:

1. Enter the information at the terminal. For this example, enter the prompted information.

The following figure shows how this screen looks after information is entered.

```

                                *EMPLOYEE PAYROLL*
                                *****

FIRST NAME: Joe                LAST NAME: Blutzen
EMPLOYEE NO: 60249

INPUT:
```

Figure 85. PAYDAY screen, with filled input fields

2. When IMS receives this data, MFS uses the DIF and the MID control blocks to translate the data from the way it was entered on the terminal screen to the way that the application program is expecting to receive it. The DIF control block tells MFS the format of the data to come in from the terminal. The MID control block tells MFS how the application program expects to receive the data. When the application program issues a message call, IMS places the "translated" message in the program's I/O area.

When the application receives the message in its I/O area, the message looks like this:

```
PAYUP JOE BLUTZEN 60249
```

"PAYUP" is the transaction code. The name of the logical terminal does not appear in the message itself; IMS places it in the first field of the I/O PCB.

3. The application program processes the message, including any required database access, and builds the output message in the application program's I/O area. After retrieving the information from the database, the program builds the output message segment for the employee, with social security and rate of pay information. The application program's I/O area contains:

```
LLZZJOE BLUTZEN 60249532596381150.00
```

The LL is a 2-byte field in MFS messages that indicates the length of the field. How the LL field is defined depends on what programming language used to write the application program. For the AIBTDLI, ASMTDLI, CEETDLI, or PASTDLI interfaces, the LL field must be defined as a binary half word. For the PLITDLI interface, the LL field must be defined as a binary fullword. The value provided in the PLITDLI interface must represent the actual segment length minus 2 bytes.

The ZZ is a 2-byte length field in MFS messages that contains the MFS formatting option that is being used to format the messages to and from the application program. MFS options are discussed in further detail in the topic "Input Message Formatting Options" in *IMS Version 15.3 Application Programming APIs*.

4. When the application program sends the message back to the terminal, MFS translates the message again, this time from the application program format to the format in which the terminal expects the data.

The MOD tells MFS the format that the message will be in when it comes from the application program's I/O area. The DOF tells MFS how the message is supposed to look on the terminal screen. MFS translates the message and IMS displays the translated message on the terminal screen.

The following figure shows how the screen looks.

```
*EMPLOYEE PAYROLL*
*****
```

```
FIRST NAME: Joe                      LAST NAME: Blutzen
EMPLOYEE NO: 60249
SOC SEC NO: 532-59-6381
RATE OF PAY: $150.00
```

```
INPUT:
```

Figure 86. PAYDAY screen, output formatted by DOF and displayed

Listing a subset of employees

Suppose you have an MPP that answers this request:

List the employees who have the skill "ENGINEER" with a skill level of "3." List only those employees who have been with the firm for at least 4 years.

To enter the request from a display terminal, issue the format command (/FORMAT) and the MOD name. This formats the screen in the way defined by the MOD you supply. When you enter the MOD name, the screen contains only literals and no output data from an application program. At this stage, an MPP is not involved. Suppose the name of the MOD that formats the screen for this request is LE, for "locate employee." Enter this:

```
/FORMAT LE
```

IMS locates the MFS MOD control block with the name LE and arranges your screen in the format defined by the DOF. Your screen then looks like this:

```
SKILL
LEVEL
YEARS
      LOCEMP
```

The DOF defines a terminal format that asks you to qualify your request for an employee by giving the skill, level, and number of years of service of the employee you want. LOCEMP is the transaction code that is associated with the MPP that can process this request. When you enter the MOD name, the transaction code is included in the first screen format that is displayed for you. This means that you do not need the name of the program that processes your request; you only need the name of the MOD that formats the screen.

After the screen format is displayed, you can enter your request. There are four stages in sending a message to the program and receiving the reply.

1. Enter the information at the terminal. In this example, enter the values of the qualifications that IMS has given you on the screen: the skill is "eng" (engineer), the skill level is "3," and the number of years with the firm is "4".

After you enter your request, your screen contains this data:

```
SKILL ENG
LEVEL 3
YEARS 4
      LOCEMP
```

2. When IMS receives this data, MFS uses the DIF and the MID control blocks to translate the data from the way you entered it on the terminal screen to the way that the application program is expecting to receive it. The DIF control block tells MFS how the data is going to come in from the terminal. The MID control block tells MFS how the application program is expecting to receive the data. When

the application program issues a GU call to the I/O PCB, IMS places the "translated" message in the program's I/O area.

When the MPP receives the message in its I/O area, the message looks like this:

LOCEMP ENG0304

"LOCEMP" is the transaction code. The name of the logical terminal does not appear in the message itself; IMS places it in the first field of the I/O PCB.

3. The MPP processes the message, including any required database access, and builds the output message in the MPP's I/O area.

Suppose more than one employee meets these qualifications. The MPP can use one message segment for each employee. After retrieving the information from the database, the program builds the output message segment for the first employee. The program's I/O area contains:

```
LLZZJONES,CE    3294
```

When the program sends the second segment, the I/O area contains:

```
LLZZBAKER,KT    4105
```

4. When the application program sends the message back to the terminal, MFS translates the message again, this time from the application program format to the format in which the terminal expects the data.

The MOD tells MFS the format that the message will be in when it comes from the application program's I/O area. The DOF tells MFS how the message is supposed to look on the terminal screen. MFS translates the message and IMS displays the translated message on the terminal screen. The screen then contains the following data:

```
SKILL    ENG
NAME     NO
JONES,CE 3294
BAKER,KT 4105
```

Related concepts

["Relationship between MFS control blocks and screen format" on page 442](#)

Use the control blocks in the MFS source language to define the formats that you see at the device.

Relationship between MFS control blocks and screen format

Use the control blocks in the MFS source language to define the formats that you see at the device.

The standard way for an end-user or operator to receive an initial format is to request it with a **/FORMAT** command, specifying the name of a MOD. In the following code example, the label on the MOD is PAYDAY. This MOD contains the parameter SOR=PAYF, which points to a device output format, or DOF, with the same label.

The initial DOF also becomes the format for device input. Therefore, if you specify DIV TYPE=INOUT in the DOF, a device input format (DIF) is also generated. In the sample code, PAYF is both a DOF and a DIF, since it also describes the format of the next input. The final output message can be displayed with a format that is specified for output only and no DIF is generated.

Both the MOD and the MID point to the same DOF, thus establishing the relationship between device-related and message-related control blocks.

For output, MFS moves fields defined in a MOD to fields on the screen defined by a DOF. When a field definition is coded (MFLD) in a MOD, it is given a label. The same label is used in the coding of the device field (DFLD) in the DOF, defining where the field appears on the screen.

MFS moves data fields from output messages to screen fields; this is referred to as *mapping*. For input, MFS moves modified screen fields to data fields in the input message for the program by mapping identically labeled fields in the DIF and MID.

For more detailed information on specifying these control blocks, see *IMS Version 15.3 Database Utilities*.

The MFS control blocks are generated from the source statements like those in the following code example during execution of the MFS Language utility. The control blocks are stored in the various MFS libraries.

The sample code is designed for a 3270 display.

Sample MFS control block coding

DOF/DIF

```
PAYF      FMT      TYPE=(3270,2),FEAT=IGNORE,DSCA=X'00A0'  
          DEV      TYPE=INOUT  
          DIV      CURSOR=((5,15))  
          DPAGE     '*****',POS=(1,21)  
          DFLD      '* EMPLOYEE PAYROLL *',POS=(2,21)  
          DFLD      '*****',POS=(3,21)  
          DFLD      'FIRST NAME:',POS=(5,2)  
FNAME     DFLD      POS=(5,15),LTH=16  
          DFLD      'LAST NAME:',POS=(5,36)  
LNAME     DFLD      POS=(5,48),LTH=16  
          DFLD      'EMPLOYEE NO:',POS=(7,2)  
EMPNO     DFLD      POS=(7,16),LTH=6  
          DFLD      'SOC SEC NO:',POS=(9,2)  
SSN       DFLD      POS=(9,15),LTH=11  
          DFLD      'RATE OF PAY: $',POS=(11,2)  
RATE      DFLD      POS=(11,17),LTH=9  
          DFLD      'INPUT:',POS=(16,2)  
INPUT     DFLD      POS=(16,10),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
```

Related reference

[“MFS examples” on page 438](#)

One way to understand the relationship between the MFS control blocks is to look at a message from the time a user enters it at the terminal to the time the application program processes the message and sends a reply back to the terminal. Though MFS can be used with both display terminals and printer devices, for clarity in this example, a display terminal is being used.

Overview of MFS components

IMS Message Format Service (MFS) components include three utilities, a message editor, and two pool managers.

MFS utilities

You can use the MFS utilities for multiple service and generation purposes:

- MFS Device Characteristics Table utility (DFSUTB00): Define new screen sizes in a descriptor member of the IMS.PROCLIB library without completing an IMS system definition.
- MFS Language utility (DFSUPAA0): Create and store the MFS control blocks.
- MFS Service utility (DFSUTSA0): Control and maintain MFS intermediate text blocks and control blocks after they are processed and stored by the MFS Language utility (DFSUPAA0).

In addition to the using the MFS utilities to update MFS libraries, you can also use the IMS online change function. You can modify control block libraries while the IMS control region is executing.

MFS message editor

Use the MFS message editor to formats messages according to the control block specifications generated by the MFS Language utility from control statement definitions that you enter.

MFS pool managers

You can customize the functions of the following MFS pool managers:

- MFS pool manager: MFS tries to minimize I/O to the format library by keeping referenced blocks in storage. This storage is managed by the MFS pool manager. You can use the INDEX function of the MFS Service utility to customize this function by constructing a list of the directory addresses for specified format blocks. This list eliminates the need for IMS to read the data set directory before it fetches a block.
- MFSTEST pool manager: If you use the MFSTEST facility, MFS control blocks are managed by the MFSTEST pool manager. The communication line buffer pool space allowed for MFS testing is specified during system definition, but the space can be changed when the IMS control region is initialized. This space value is the maximum amount used for MFSTEST blocks at any one time. The space value is not a reserved portion of the pool.

Related concepts

[The online change function \(System Administration\)](#)

[MFS components \(Communications and Connections\)](#)

[Use of the message format buffer pool \(System Definition\)](#)

Related reference

[MFS Language utility \(DFSUPAA0\) \(System Utilities\)](#)

[MFS Service utility \(DFSUTSA0\) \(System Utilities\)](#)

[MFS Device Characteristics Table utility \(DFSUTB00\) \(System Utilities\)](#)

Devices and logical units that operate with MFS

In addition to 3270 devices, MFS operates with the 3600 and 4700 Finance Communication System (FIN), the 3770 Data Communication System, the 3790 Communication System, and with Secondary Logical Unit (SLU) types 1, 2, 6, and P. Network Terminal Option (NTO) devices are supported as secondary logical unit type 1 consoles.

The following table shows which devices or logical units can be defined for MFS operation in the IMS system by their number (3270, for example), and which can be defined by the type of logical unit to which they are assigned (SLU 1, for example).

Although the 3600 devices are included in the FIN series, you can specify them with their 36xx designations; MFS messages use the FIXx designations regardless of which form of designation you specify. In general, however, application designers and programmers using this information need to know only how the devices they are defining control blocks for have been defined to the IMS system in their installation.

Table 85. Terminal devices that operate with MFS

Device	Devices defined by number¹	NTD devices²	SLU 1	SLU 2	SLU P	LU 6.1
3180	X ³			X ³		
3270	X ³			X ³		
3290	X ³			X ³		
5550	X ³			TYPE: 3270- An 3270- Ann		
3270 printers; 5553, 5557	X ³		COMPT _n = MFS-SCS1			
3730					X	
3767			COMPT _n = MFS-SCS1			
3770 console, printers, print data set			COMPT _n = MFS-SCS1		X	
3770 readers, punches, transmit data set			COMPT _n = MFS-SCS2		X	
3790 print data set (bulk)			COMPT _n = MFS-SCS1		COMPT _n = MFS-SCS1 DPM-An	
3790 transmit data set			COMPT _n = MFS-SCS2			
3790 attached 3270				X ³		
6670						
8100					X	
8100 attached 3270			X	X ³		
8100 attached Series/1					X	
8100 attached S/32			X			
8100 attached S/34					X	
8100 attached S/38			X			
Finance	X				COMPT _n = MFS-SCS1 DPM-An	
TTY		X				
3101		X				

Table 85. Terminal devices that operate with MFS (continued)

Device	Devices defined by number ¹	NTO devices ²	SLU 1	SLU 2	SLU P	LU 6.1
Other systems (IMS to IMS or IMS to other)						COMPTn= DPM=Bn

Notes:

1. With options= (...,MFS,...) in the TERMINAL or TYPE macro.
2. Defined with UNITYPE= on the TYPE macro and PU= on the TERMINAL macro.
3. Defaults to operate with MFS.

Logical units are defined by logical unit type or logical unit type with COMPTn= or TYPE= in the TERMINAL macro or ETO logon descriptor. The LU 6.1 definition refers to ISC subsystems.

The definition for SLU 1 can specify an MFS operation with SNA character strings (SCS) 1 or 2. SCS1 designates that messages are sent to a printer or the print data set or received from a keyboard in the 3770 Programmable or 3790 controller disk storage; SCS2 designates that messages are sent to or received from card I/O or a transmit data set.

Terminals defined as SLU 2 have characteristics like the 3270, and like the 3270, can be defined to operate with MFS. In general, a 3290 terminal operates like a 3270 terminal, and references to 3270 terminals in this information are applicable to 3290 devices. However, 3290 partitioning and scrolling support is provided only for 3290 devices defined to IMS as SLU 2.

Generally, the 3180 and 5550 terminals operate like a 3270 terminal, and references to 3270 terminals also apply to these devices. Likewise, the 5553 and 5557 printer devices operate like a 3270P.

Restriction: 5550 Kanji support is provided only for the 5550 terminal defined as an SLU 2 and for the 5553 and 5557 defined as SCS1 printers.

If IMS is to communicate with the user-written remote program in a 3790 or an FIN controller, the device must be defined as an SLU P. Definitions for SLU P must specify MFS operation as either MFS-SCS1 or DPM-An, where DPM means distributed presentation management and An is a user-assigned number (A1 through A15).

Most of the MFS formatting functions currently available to other devices, except specific device formatting, are available to the user-written program. Under user control, these formatting functions (such as paging) can be divided between MFS and the remote program.

Using distributed presentation management (DPM)

With distributed presentation management (DPM), formatting functions usually performed by MFS are distributed between MFS and a user-written program for SLU P devices or ISC nodes. If the 3790 or FIN controller has previously been defined to IMS by unit number, some changes must be made to convert to DPM.

With DPM, the physical terminal characteristics of the secondary logical unit do not have to be defined to MFS. MFS has to format only the messages for transmission to the user program in the remote controller or ISC node, which must assume responsibility for completing the device formatting, if necessary, and present the data to the physical device it selects.

For remote programs using DPM, the data stream passing between MFS and the remote programs can be device independent. The messages from the IMS application program can include some device control characters. If so, the IMS application program and the data stream to the remote program might lose their device independence.

If IMS is to communicate with other subsystems (such as IMS, CICS or user-written), the other subsystem must be defined as an ISC subsystem. Definitions for ISC must:

- Specify MFS operation as DPM-Bn, where Bn is a user-assigned number (B1 through B15).
- Define TYPE:LUTYPE6 on the TERMINAL macro during system definition.

DPM with ISC provides:

- Output paging on demand that allows paging to be distributed between IMS and another system
- Automatically paged output that allows MFS pages to be transmitted to another system without intervening paging requests
- Transaction routing that allows application programs to view the routing information when it is provided in the input message

Chapter 28. Callout requests for services or data

IMS applications can issue callout requests for services or data, and optionally receive responses back in the same or a different transaction, through IMS Connect and OTMA. The request for services or data is a *callout request*.

If the IMS application, after issuing the request, waits for a response in the dependent region, the request is a *synchronous callout request*. If the IMS application terminates after the request is issued and does not wait for a response in the dependent region, the request is an *asynchronous callout request*. IMS applications can also issue synchronous callout requests to an IMS transaction with a *synchronous program switch request*.

Each type of request is processed as follows:

Synchronous callout request

An IMS application program that runs in an IMS dependent region issues a DL/I ICAL call and waits in the dependent region to process the response. The application program can use the optional *control data area* of the DL/I ICAL call to pass routing, security, or other data to IMS Connect and its clients. When the DL/I ICAL call is issued, IMS generates a correlation token for synchronous callout requests. This 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.

Asynchronous callout request

An IMS application program that runs in an IMS dependent region inserts the callout request to an ALTPCB queue (the ISRT ALTPCB call) and then terminates to free the dependent region. IMS does not generate a correlation token for asynchronous callout requests. If a response to the callout request is required, the correlation of the response to the callout request 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.

Synchronous program switch request

A synchronous program switch uses the ICAL call, but the request is routed to an IMS transaction rather than an external server. Synchronous program switch requests do not use a correlation token because IMS automatically correlates the response back to the waiting application program.

The following table summarizes the differences among synchronous callout requests, asynchronous callout requests, and synchronous program switch requests.

Table 86. Comparison of synchronous and asynchronous callout requests

Callout process	Synchronous callout request	Asynchronous callout request	Synchronous program switch request
Placing the request in the OTMA hold queue	The requesting IMS application issues an ICAL call with or without control data.	The requesting IMS application issues an ISRT ALTPCB call.	The requesting IMS application issues an ICAL call.
Status of the IMS application after the request is issued	The application waits in the dependent region for the response. Dependent regions are blocked.	The application terminates.	The application waits in the dependent region for the response. Dependent regions are blocked.
Message processing handling	The message processing is handled by IMS OTMA.	The message processing is handled by the IMS message queue.	The message processing is handled by the IMS message queue.

Table 86. Comparison of synchronous and asynchronous callout requests (continued)

Callout process	Synchronous callout request	Asynchronous callout request	Synchronous program switch request
Response handling	The response is correlated back to the requesting IMS application, based on the correlation token, during the same unit of work.	If there is a response, the requesting or a different IMS application must be coded to handle the response that is returned in a different transaction. The unit of work for the transaction has to commit for the asynchronous output to flow.	The response is returned to the requesting IMS application in the same unit of work. However, the target transaction executes in a separate unit of work, so it is not eligible for two-phase commit and it is not part of the RRS commit scope for the original application.

Related reference

[ICAL call \(Application Programming APIs\)](#)

Callout request approaches

You can issue a callout request to IMS Enterprise Suite SOAP Gateway, IMS TM Resource Adapter, IBM MQ, to your own user-supplied IMS Connect client applications, or to another IMS application program.

For synchronous callout requests, you can optionally include *control data* in a callout message. By using control data, you can pass the URL for the port, UUID, user token, security information, or any other information to the IMS Connect and its client.

Using SOAP Gateway

Use SOAP Gateway to issue callout requests from IMS applications to any generic web service.

SOAP Gateway enables IMS applications as either web service providers or consumers. SOAP Gateway supports both asynchronous and synchronous callout approaches for IMS applications as web service consumers. Tooling support for SOAP Gateway is available in IBM Developer for System z® for generating the required web service artifacts based on connection and interaction information for communicating with IMS Connect, and the language structure of the IMS applications. SOAP Gateway also provides a deployment utility to support the deployment of IMS applications as either providers or consumers of web services.

For SOAP Gateway messages, you can specify an XML converter name in the control data by using the tags <DFSCNVTR>CONVERTER_NAME</DFSCNVTR>. The converter name and the tags must be in uppercase EBCDIC. If the converter name is present, it overrides the current converter name that IMS Connect would have used to process the message. As soon as the converter name is extracted, no further scanning of the tags will take place for that message. The tags that contain the XML converter name are not removed from the control data and are sent to SOAP Gateway with any other tags that are present in the control data.

The following example illustrates control data area with two control data items:

```
Total Length = 4 + 10 + 8 + 11 + 4 + 9 + 25 + 10 = 81 = X'51'
AIBOPLN = X'00000051'

Control Data = X'00000021' <DFSCNVTR>CONVERT1</DFSCNVTR> X'00000030'
<USERTAG>USER DATA CAN BE ANYTHING</USERTAG>
```

For more information about enabling IMS application callout through SOAP Gateway, see [Enabling an IMS application as a web service consumer](#).

Using IMS TM Resource Adapter

Use IMS TM Resource Adapter Version 10 or later to issue synchronous or asynchronous callout requests from IMS applications to any message-driven bean (MDB), Enterprise JavaBeans (EJB) component, Java EE (previously known as J2EE) application, or web service.

IMS TM Resource Adapter enables Java EE applications to access IMS transactions over the Internet, as well as to issue callout requests to external Java EE applications from IMS applications that run in IMS dependent regions. The IMS TM Resource Adapter includes a runtime component for WebSphere Application Server. Tooling support for the IMS TM Resource Adapter is available in IBM Rational® Application Developer for WebSphere Software, as well as various Rational and WebSphere integrated development environments (IDEs) that include the J2EE Connector (J2C) wizard.

For more information about callout support in IMS TM Resource Adapter, see [Callout programming models \(TM Resource Adapter\)](#).

Using IBM MQ

You can write applications that make asynchronous callout requests to other applications via IBM MQ. You must configure the IBM MQ to IMS bridge.

You must also create a destination descriptor for IBM MQ with TYPE=MQSERIES. You can either add the descriptor to the DFSYDTx member of the IMS PROCLIB data set, or add it with the **CREATE OTMADESC** command. The OTMA routing exits (DFSYPX0 and DFSYDRU0) are not required.

The destination descriptor type for IBM MQ includes parameters that you can use to customize the fields of the MQMD data structure. The MQMD structure controls how IBM MQ processes messages that use the descriptor.

Using a user-written IMS Connect TCP/IP application

You can write your own IMS Connect TCP/IP applications or use a vendor-supplied solution that uses TCP/IP and the IMS Connect protocol to retrieve callout requests. Your custom IMS Connect client application must issue a RESUME TPIPE call to an OTMA routing destination, also known as a transaction pipe (tpipe), that is defined in an OTMA destination descriptor. This tpipe holds the callout requests. Your custom IMS Connect TCP/IP application must poll the tpipe to retrieve the callout requests.

Using IMS synchronous program switch requests

You can use the ICAL call to send a request to another IMS application by creating a destination descriptor with TYPE=IMSTRAN. This type of callout request is a *synchronous program switch*.

The IMSTRAN descriptor type specifies the destination transaction, and can also be used to create a late response message queue. You can either add the descriptor to the DFSYDTx member of the IMS PROCLIB data set, or add it with the **CREATE OTMADESC** command. Multiple synchronous program switch requests can be chained together if the destination application also issues an ICAL call for synchronous program switch.

OTMA is not required to use the ICAL call for synchronous program switching. You must define an OTMA destination descriptor for the destination transaction, but IMS schedules the transaction whether or not OTMA is active.

Control data is not supported for synchronous program switch requests.

Using the Java dependent region resource adapter

The Java dependent region resource adapter supports the Java Message Service (JMS) interface that can be used by Java applications for the same functionality in Java dependent regions.

IMS™ provides a set of Java™ APIs called the IMS Java dependent region resource adapter to develop Java applications to run on the IMS Java dependent regions.

The IMS Java dependent region resource adapter provides Java application programs running in JMP or JBP regions with similar DL/I functionality to that provided in message processing program (MPP) and non-message driven BMP regions, such as:

- Accessing IMS message queues to read and write messages
- Performing program switches
- Commit and rollback processing
- Accessing GSAM databases
- Database recovery (CHKP/XRST)

IMS transactions written in Java and leveraging the IMS Universal Java dependent region (JDR) resource adapter can issue ICAL calls that include control data. The JDR resource adapter invokes the Universal Drivers C library (DFSCLIBU) through Java Native Interface (JNI) calls to issue the calls from C to the AIBTDLI interface with the ICAL information.

Use the IMS Java dependent region resource adapter together with the type-2 IMS Universal JDBC driver or type-2 IMS Universal DL/I driver to perform database operations, including GSAM database access.

Related concepts

[Callout requests from IMS application programs \(Communications and Connections\)](#)

[OTMA destination descriptors \(Communications and Connections\)](#)

Related reference

[DFSYDTx member of the IMS PROCLIB data set \(System Definition\)](#)

[ICAL call \(Application Programming APIs\)](#)

Resume tpipe protocol

The resume tpipe protocol retrieves asynchronous and synchronous callout messages from IMS.

An IMS Connect client signals how long to wait for output from IMS by specifying an IRM timeout value with the IRM_TIMER field. The IRM timeout value affects the RESUME TPIPE call that is sent to IMS Connect and the ACK or NAK response message that is sent to IMS Connect.

When you use IMS TM Resource Adapter or IMS Enterprise Suite SOAP Gateway to handle the callout request from your IMS application, the communication with IMS Connect is handled for you.

Both the IMS TM Resource Adapter and SOAP Gateway listen for synchronous 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 (for SOAP Gateway if using the IMS Connect XML adapter function), and then sends the message to the IMS TM Resource Adapter or SOAP Gateway.

If you have a custom IMS Connect client, you must code the client to issue a RESUME TPIPE call to retrieve the callout messages

Resume tpipe security

You can protect callout messages from unauthorized use of the RESUME TPIPE call by using either the Resource Access Control Facility (RACF), the OTMA Resume TPIPE Security user exit (OTMARTUX), or both.

When security is enabled, the user ID that issues the RESUME TPIPE call must be authorized to access the tpipe name that is contained in the RESUME TPIPE call message before any messages are sent to an OTMA client.

The security checking performed by RACF and the security checking performed by the OTMARTUX user exit are optional. If both RACF and the OTMARTUX are used, RACF is called first before giving control to the OTMARTUX user exit, in which case, the OTMARTUX user exit can override RACF, depending on your needs.

Implementing the synchronous callout function

To issue a synchronous callout request from your IMS application, issue the ICAL call and specify the OTMA descriptor name.

The ICAL call can also be issued through a REXXTDLI call, or from a Java application that runs in a JMP or JBP region. Optionally, you can also specify a timeout value (the maximum time to wait for the response to return).

Input and output messages from IMS can be 32 KB or larger per segment for a synchronous callout request. However, the maximum segment size is 32 KB for a synchronous program switch request.

The following diagram shows the message flow of the synchronous callout function. The request starts with an IMS application that issues an ICAL call. The response is returned to the requesting IMS application.




Figure 87. Message flow of the synchronous callout function

You can make concurrent ICAL calls from your IMS applications that are routed to the same or different IMS Connect destinations. By using the RESUME TPIPE call and the send-only protocol, the synchronous callout function allows the requests to be sent and the responses to be received in different connections and threads. You can have one connection for retrieving the callout requests and use other connections to return the response messages simultaneously.

A correlation token is created by IMS to correlate the response message back to the correct IMS transaction instance. The capability of having different threads and connections for pulling callout requests and for returning response messages provides maximum concurrency.

Correlation tokens are not used for synchronous program switch requests.

Your application can make an ICAL call with the RECEIVE subfunction code to get the complete response data for a previously incomplete ICAL call with the SENDRECV subfunction code. However,

your application cannot retrieve the complete response data with a RECEIVE subfunction call after issuing a subsequent call with the SENDRECV subfunction. Issuing a subsequent ICAL call with the SENDRECV subfunction code clears the ICAL response buffer in the IMS control region for the previous ICAL call.

The following high-level steps provide an overview of implementing and deploying your synchronous callout application and function.

1. Create or modify an IMS application for the ICAL call.
2. Define the OTMA destination descriptor for one of the following callout request approaches:
 - Send the callout request to an external application viaIMS Enterprise SuiteIMS TM Resource Adapter.
 - Send the callout request to an external application viaIMS Enterprise Suite SOAP Gateway.
 - Send the callout request to a user-written IMS Connect client application.
 - Send a synchronous program switch request to another IMS application.

The destination descriptor can be defined in the DFSYDTx member of the IMS.PROCLIB data set, or with the **CREATE OTMADESC** command.

3. Restart IMS for the newly defined OTMA descriptor.
This step is not required if the descriptor was dynamically added with the **CREATE OTMADESC** command.
4. Run the IMS application that was created or modified in step 1 to issue the synchronous callout request.

The IMS application receives the response message from the synchronous callout target. If the response message does not fit in the allocated response area, the application can expand the available response data area and then issue an ICAL call with the RECEIVE subfunction code to get the complete response.

Example COBOL program implementation of the synchronous callout function

To issue the ICAL call in a COBOL program, use the CALL statement.

```
CALL 'AIBTDLI' USING ICAL, AIB, CA-REQUEST, SCA-RESPONSE.
```

The following example demonstrates the required AIB field declaration for the ICAL call in the COBOL program. A complete COBOL example (with part name DFSSSCBL) is provided with the callout IVP sample in the SDFSSMPL sample library.

```
01 AIB.
  02 AIBRID PIC x(8) VALUE 'DFS AIB '.
  02 AIBRLN PIC 9(9) USAGE BINARY.
  02 AIBSFUNC PIC x(8) VALUE 'SENDRECV'.
  02 AIBRSNM1 PIC x(8) VALUE 'OTMDEST1'.
  02 AIBRSNM2 PIC x(8).
  02 AIBRESV1 PIC x(8).
  02 AIBOALEN PIC 9(9) USAGE BINARY VALUE 28.
  02 AIBOAUSE PIC 9(9) USAGE BINARY VALUE 30.
  02 AIBRSFLD PIC 9(9) USAGE BINARY VALUE 5000.
  02 AIBRESV2 PIC x(8).
  02 AIBRETRN PIC 9(9) USAGE BINARY.
  02 AIBREASN PIC 9(9) USAGE BINARY.
  02 AIBERRXT PIC 9(9) USAGE BINARY.
  ...
```

The following example shows the CA-REQUEST and SCA-RESPONSE declarations in the COBOL program.

```
* ICAL Request Area
01 CA-REQUEST.
  02 CA-MESSAGE PIC X(45) VALUE SPACES.

* ICAL Response Area
01 SCA-RESPONSE.
  02 SCA-MESSAGE PIC X(100) VALUE SPACES.
```

Related concepts

[OTMA destination descriptors \(Communications and Connections\)](#)

[Callout programming models \(TM Resource Adapter\)](#)

Related tasks

[Modifying an IMS application for callout requests](#)

Related reference

[ICAL call \(Application Programming APIs\)](#)

[Examples of DL/I call functions \(Application Programming APIs\)](#)

[“Programming with the Callout API” on page 692](#)

Use the `IMSCallout` API to issue IMS synchronous callout requests from JMP or JBP applications that are running in a Java dependent region. Another option is to use the Java Message Service (JMS) API, which provides limited support.

Control data in synchronous callout requests

When the DL/I ICAL call is used for synchronous callout requests, IMS application programs can specify the endpoint information or other routing specification for the callout message in the control data area of the ICAL call during run time when they issue the call.

IMS application programs can make a synchronous callout request by using the ICAL DL/I call. The ICAL call sends a callout request to an IMS Connect client application and receives a response. The routing specification of a callout message is defined in the message's application interface block (AIB) and an OTMA destination descriptor. You can define up to 4095 destination descriptors in a DFSYDTx member of the IMS PROCLIB data set. When there are thousands of callout end points for the messages, using a limited number of OTMA destination descriptor entries for routing specification creates a challenge. Additionally, there is no easy way to include multiple Universally Unique Identifiers (UUID), SOAP headers, security token, or even user-provided routing information in the callout message.

To solve this problem, the ICAL call format accepts optional *control data*, which can be the URL for a port, UUID, user token, security information, or any other information. By using the control data field, IMS application programs can specify the routing information or other control data at run time when they issue the ICAL call.

The control data can consist of 1 to many control data items so that a number of services or operations can be specified on the same synchronous callout call. Each control data item starts with 4 bytes length field followed by a tag, data, and an end tag.

The ICAL control data can consist of one to many control data items so that a number of services or operations can be specified on the same ICAL call. Each control data item starts with a 4-byte length field followed by a tag, data, and the end tag. Tags can be of any length. The start tag consists of a less than sign (<), the tag name, and a greater than sign (>). The end tag consists of a less than sign (<), a slash (/), the tag name that matches the start tag, and a greater than sign (>).

The format of a control data item in ICAL control data is as follows:

```
LLLL <tag1> data1 </tag1> { LLLL <tag2> data2 </tag2> ... }
```

An IBM-initiated control data item starts with DFS in the tag, such as <DFSCNVTR>.

The tag name and data contents are treated as binary data and passed "as is" to the target client. The <, /, and > signs, and IBM-initiated control data tag names, which begin with DFS, must be EBCDIC.

OTMA does the “well-form” checking for the control data to make sure that it follows the supported format with the correct length. The total length of the control data with control data items needs to be specified in the AIBOPLN field in the AIB. OTMA puts the control data in front of the application data and update the OTMA prefix to indicate the number of segments of control data in the application data section for a callout message. The OTMA resume tpipe protocol command supports the callout with control data option.

The IMS Connect user message exits look for a flag in the IRM that indicates whether the client supports control data. If a flag is set, an appropriate flag is set in the state data section of the OTMA headers for the resume tpipe. Upon receiving the message that contains control data from IMS, the user message exits extract the control data from the application data section of the message and build a segment similar to the callout correlation token segment. This is done for messages other than IMS TM Resource Adapter messages for IMS Connect. The resulting message from the user message exit will then be transmitted to the client by IMS Connect.

If you are using SOAP Gateway, you can use control data to specify the name of the XML converter you would like to use to process the request. In the control data, use the <DFSCNVTR> tags to specify an XML converter name.

The IMS Java dependent region (JDR) support provides an API for control data of the ICAL call.

The control data is designed for a callout request for an outbound message. The control data in a response message is not supported.

Implementing the asynchronous callout function

To issue an asynchronous callout request from your IMS application, issue the ISRT ALTPCB call and specify the OTMA destination descriptor name or the DFSYPRX0 and DFSYDRU0 routing exit routines.

Any response to the callout request that is returned to IMS is handled as a new incoming transaction. If there is a response, the requesting application or a different IMS application must be coded to handle the response that is returned in a separate transaction.

Unlike synchronous callout requests, asynchronous callout requests do not require the IMS application program that issues the request to wait for a response in the dependent region. After it issues 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 handled 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 that correlation.

The following high-level steps provide an overview of implementing and deploying your asynchronous callout application and function.

1. Plan for the correlation of asynchronous callout responses.
2. Create or modify an IMS application to issue an ISRT ALTPCB call for asynchronous callout requests.
3. Define the callout routing information. There are two options to define the required information:

- Define an OTMA routing descriptor.
- Code the DFSYPRX0 and DFSYDRU0 exit routines.

You can use either the routing descriptor, the exit routines, or a combination of both to specify how the callout request is routed.

4. Optional: Restart IMS for the newly defined OTMA descriptor.

A restart is required only if you create or modify an OTMA routing descriptor in the DFSYDTx member of the IMS.PROCLIB data set. You do not need to restart IMS if you use the **CREATE OTMADESC** or **UPDATE OTMADESC** commands.

5. Run the IMS application that issues the callout request.

The IMS application is usually triggered through an initiating client, such as a terminal, or an IMS Connect or OTMA client.

Related concepts

[Asynchronous callout request \(Communications and Connections\)](#)

[OTMA destination descriptors \(Communications and Connections\)](#)

Related reference

[OTMA User Data Formatting exit routine \(DFSYDRU0\) \(Exit Routines\)](#)

[OTMA Destination Resolution user exit \(DFSYPX0 and other OTMAYPRX type exits\) \(Exit Routines\)](#)

Part 4. Application programming for EXEC DLI

IMS provides support for writing applications to access IMS resources using EXEC DLI.

Chapter 29. Writing your application programs for EXEC DLI

You can write programs in assembler language, COBOL, PL/I, C, and C++ that execute EXEC DLI commands to access IMS.

Programming guidelines

Use the following guidelines to write efficient and error-free EXEC DL/I programs.

The number, type, and sequence of the DL/I requests your program issues affect the efficiency of your program. A program that is poorly designed runs if it is coded correctly. The suggestions that follow can help you develop the most efficient design possible for your application program. Inefficiently designed programs can adversely affect performance and are hard to change. Being aware of how certain combinations of commands or calls affects performance helps you to avoid these problems and design a more efficient program.

After you have a general sequence of calls mapped out for your program, use these guidelines to improve the sequence. Usually an efficient sequence of requests causes efficient internal DL/I processing.

- Use the simplest call. Qualify your requests to narrow the search for DL/I, but do not use more qualification than required.
- Use the request or sequence of requests that gives DL/I the shortest path to the segment you want.
- Use the fewest number of requests possible in your program. Each DL/I request your program issues uses system time and resources. You may be able to eliminate unnecessary calls by:
 - Using path requests if you are replacing, retrieving, or inserting more than one segment in the same path. If you are using more than one request to do this, you are issuing unnecessary requests.
 - Changing the sequence so that your program saves the segment in a separate I/O area, and then gets it from that I/O area the second time it needs the segment. If your program retrieves the same segment more than once during program execution, you are issuing an unnecessary request.
 - Anticipating and eliminating needless and nonproductive requests, such as requests that result in GB, GE, and II status codes. For example, if you are issuing GNs for a particular segment type and you know how many occurrences of that segment type exist, do not issue the GN that results in a GE status code. You can keep track of the number of occurrences your program retrieves, and then continue with other processing when you know you have retrieved all the occurrences of that segment type.
 - Issuing an insert request with a qualification for each parent instead of issuing Get requests for the parents to make sure that they exist. When you are inserting segments, you cannot insert dependents unless the parents exist. If DL/I returns a GE status code, at least one of the parents does not exist.
- Keep the main section of the program logic together. For example, branch to conditional routines, such as error and print routines, in other parts of the program, instead of having to branch around them to continue normal processing.
- Use call sequences that make good use of the physical placement of the data. Access segments in hierarchical sequence as much as possible. Avoid moving backward in the hierarchy.
- Process database records in order of the key field of the root segments. (For HDAM databases, this order depends on the randomizing routine that is used. Check with your DBA for this information.)
- Try to avoid constructing the logic of the program and the structure of commands or calls in a way that depends heavily on the database structure. Depending on the current structure of the hierarchy reduces the program's flexibility.

Coding a program in assembler language

The following sample assembler language program shows how the different parts of a command-level program fit together, and how the EXEC DLI commands are coded in a CICS online program.

Except for a few commands, this program applies to batch, BMP, and CICS programs. Any differences are highlighted in the notes for the sample assembler code. The numbering on the right of the sample code references these notes.

```
*ASM XOPTS(CICS,DLI)
*
R2      EQU 2
R3      EQU 3
R4      EQU 4
R11     EQU 11
R12     EQU 12
R13     EQU 13
DFHEISTG DSECT
SEGKEYA DS CL4
SEGKEYB DS CL4
SEGKEYC DS CL4
SEGKEY1 DS CL4
SEGKEY2 DS CL4
CONKEYB DS CL8
SEGNAME DS CL8
SEGLN   DS H
PCBNUM  DS H
AREAA   DS CL80
AREAB   DS CL80
AREAC   DS CL80
AREAG   DS CL250
AREASTAT DS CL360
*      COPY MAPSET
*
*****
*  INITIALIZATION
*  HANDLE ERROR CONDITIONS IN ERROR ROUTINE
*  HANDLE ABENDS (DLI ERROR STATUS CODES) IN ABEND ROUTINE
*  RECEIVE INPUT MESSAGE
*****
SAMPLE  DFHEIENT CODEREG=(R2,R3),DATAREG=(R13,R12),EIBREG=R11
*
      EXEC CICS HANDLE CONDITION ERROR(ERRORS)
*
      EXEC CICS HANDLE ABEND LABEL(ABENDS)
*
      EXEC CICS RECEIVE MAP('SAMPMAP') MAPSET('MAPSET')
*      ANALYZE INPUT MESSAGE AND PERFORM NON-DLI PROCESSING
*
*****
*  SCHEDULE PSB NAMED 'SAMPLE1'
*****
      EXEC DLI SCHD PSB(SAMPLE1)
BAL     R4,TESTDIB          CHECK STATUS
*
*****
*  RETRIEVE ROOT SEGMENT AND ALL ITS DEPENDENTS
*****
*
      MVC     SEGKEYA,=C'A300'
      EXEC DLI GU USING PCB(1) SEGMENT(SEGA) INTO(AREAA)
      SEGLNTH(80) WHERE(KEYA=SEGKEYA) FIELDLENGTH(4)
      BAL     R4,TESTDIB          CHECK STATUS
GNPLOOP EQU *
      EXEC DLI GNP USING PCB(1) INTO(AREAG) SEGLNTH(250)
      CLC     DIBSTAT,=C'GE'      LOOK FOR END
      BE      LOOPDONE           DONE AT 'GE'
      BAL     R4,TESTDIB          CHECK STATUS
      B       GNPLOOP
LOOPDONE EQU *
*
*****
*  INSERT NEW ROOT SEGMENT
*****
*
      MVC     AREAA,=CL80'DATA FOR NEW SEGMENT INCLUDING KEY'
      EXEC DLI ISRT USING PCB(1) SEGMENT(SEGA) FROM(AREAA)
```

1

2

3

4

5

6

6

6

7

8

X

9

X

```

      SEGLENGTH(80)
BAL   R4,TESTDIB          CHECK STATUS
*
*****
* RETRIEVE 3 SEGMENTS IN PATH AND REPLACE THEM
*****
*
      MVC   SEGKEYA,=C'A200'
      MVC   SEGKEYB,=C'B240'
      MVC   SEGKEYC,=C'C241'
      EXEC  DLI GU USING PCB(1)
      SEGMENT(SEGA) WHERE(KEYA=SEGKEYA)
      FIELDLENGTH(4)
      INTO(AREA)
      SEGLENGTH(80)
      SEGMENT(SEGB) WHERE(KEYB=SEGKEYB) FIELDLENGTH(4)
      INTO(AREAB)
      SEGLENGTH(80)
      SEGMENT(SEGC) WHERE(KEYC=SEGKEYC) FIELDLENGTH(4)
      INTO(AREAC)
      SEGLENGTH(80)
BAL   R4,TESTDIB
*
      UPDATE FIELDS IN THE 3 SEGMENTS
      EXEC  DLI REPL USING PCB(1)
      SEGMENT(SEGA) FROM(AREA) SEGLENGTH(80)
      SEGMENT(SEGB) FROM(AREAB) SEGLENGTH(80)
      SEGMENT(SEGC) FROM(AREAC) SEGLENGTH(80)
BAL   R4,TESTDIB          CHECK STATUS
*
*****
* INSERT NEW SEGMENT USING CONCATENATED KEY TO QUALIFY PARENT
*****
*
      MVC   AREAC,=CL80'DATA FOR NEW SEGMENT INCLUDING KEY'
      MVC   CONKEYB,=C'A200B240'
      EXEC  DLI ISRT USING PCB(1)
      SEGMENT(SEGB) KEYS(CONKEYB) KEYLENGTH(8)
      SEGMENT(SEGC) FROM(AREAC) SEGLENGTH(80)
BAL   R4,TESTDIB          CHECK STATUS
*
*****
* RETRIEVE SEGMENT DIRECTLY USING CONCATENATED KEY
* AND THEN DELETE IT AND ITS DEPENDENTS
*****
*
      MVC   CONKEYB,=C'A200B230'
      EXEC  DLI GU USING PCB(1)
      SEGMENT(SEGB)
      KEYS(CONKEYB) KEYLENGTH(8)
      INTO(AREAB) SEGLENGTH(80)
BAL   R4,TESTDIB          CHECK STATUS
      EXEC  DLI DLET USING PCB(1)
      SEGMENT(SEGB) SEGLENGTH(80) FROM(AREAB)
BAL   R4,TESTDIB          CHECK STATUS
*
*****
* RETRIEVE SEGMENT BY QUALIFYING PARENT WITH CONCATENATED KEY,
* OBJECT SEGMENT WITH WHERE OPTION USING A LITERAL,
* AND THEN SET PARENTAGE
*
* USE VARIABLES FOR PCB INDEX, SEGMENT NAME, AND SEGMENT LENGTH
*****
*
      MVC   CONKEYB,=C'A200B230'
      MVC   SEGNAME,=CL8'SEGA'
      MVC   SEGLEN,=H'80'
      MVC   PCBNUM,=H'1'
      EXEC  DLI GU USING PCB(PCBNUM)
      SEGMENT((SEGNAME))
      KEYS(CONKEYB) KEYLENGTH(8) SETPARENT
      SEGMENT(SEGC) INTO(AREAC) SEGLENGTH(SEGLEN)
      WHERE(KEYC='C520')
BAL   R4,TESTDIB          CHECK STATUS
*
*****
* RETRIEVE DATABASE STATISTICS
*****
*
      EXEC  DLI STAT USING PCB(1) INTO(AREASTAT)
      VSAM FORMATTED LENGTH(360)
BAL   R4,TESTDIB          CHECK STATUS
*

```

```

*****
* RETRIEVE ROOT SEGMENT USING BOOLEAN OPERATORS
*****
*
    MVC    SEGKEY1,=C'A050'
    MVC    SEGKEY2,=C'A150'
    EXEC DLI GU USING PCB(1) SEGMENT(SEGA) INTO(AREAA)           X
                                SEGLength(80) FIELDLENGTH(4,4,4,4) X
                                WHERE(KEYA > SEGKEY1 AND KEYA < SEGKEY2
                                KEYA > 'A275' AND KEYA < 'A350')
    BAL    R4,TESTDIB                                CHECK STATUS

*
*****
* TERMINATE PSB WHEN DLI PROCESSING IS COMPLETED
*****
*
    EXEC DLI TERM                                           11
*
*****
* SEND OUTPUT MESSAGE
*****
*
    EXEC CICS SEND MAP('SAMPMap') MAPSET('MAPSET')           6
    EXEC CICS WAIT TERMINAL
*
*****
* COMPLETE TRANSACTION AND RETURN TO CICS
*****
*
    EXEC CICS RETURN                                           12
*
*****
* CHECK STATUS IN DIB
*****
*
TESTDIB EQU    *
CLC     DIBSTAT,=C' ' IS STATUS BLANK
BER     R4                                YES - RETURN
*
HANDLE DLI STATUS CODES REPRESENTING EXCEPTIONAL CONDITIONS
*
BR      R4                                RETURN
ERRORS EQU    *
*
HANDLE ERROR CONDITIONS
*
ABENDS EQU    *
*
HANDLE ABENDS INCLUDING DLI ERROR STATUS CODES
*
END

```

Notes for the sample assembler code:

- 1** For a CICS online program containing EXEC DLI commands, you must specify the DLI and CICS options. For a batch or BMP program containing EXEC DLI, you must specify only the DLI option.
- 2** For reentry, define each of the areas the program uses—I/O areas, key feedback areas, and segment name areas in DFHEISTG.
- 3** Define an I/O area for each segment you retrieve, add, or replace (in a single command).
- 4** For a batch or BMP program containing EXEC DLI, you must save registers on entry and restore registers on exit according to z/OS register-saving conventions.
- 5** In a batch or BMP program, a DFHEIRET with an optional DFHEIENT saves the registers on entry. Do not specify the EIBREG parameter in a batch program.
- 6** Do not code EXEC CICS commands in a batch or BMP program.
- 7** In a CICS online program, use the SCHD PSB command to obtain a PSB for the use of your program. Do *not* schedule a PSB in a batch or BMP program.
- 8** This GU command retrieves the first occurrence of SEGA with a key of A300. You do not have to provide the KEYLENGTH or SEGLength options in an assembler language program.
- 9** This GNP command retrieves all dependents under segment SEGA. The GE status code indicates that no more dependents exist.
- 10** This GU command is an example of a path command. Use a separate I/O area for each segment you retrieve.
- 11** In a CICS online program, the TERM command terminates the PSB scheduled earlier. You do *not* terminate the PSB in a batch or BMP program.

12 For a batch or BMP program, code RCREG parameter instead of EXEC CICS RETURN. The RCREG parameter identifies a register containing the return code.

13 After issuing each command, you should check the status code in the DIB.

Coding a program in COBOL

The following sample COBOL program shows how the different parts of a command-level program fit together, and how the EXEC DLI commands are coded in a CICS online program.

Except for a few commands, this program applies to batch, BMP, and CICS programs. Any differences are highlighted in the notes for the sample COBOL code. The numbering on the right of the sample code references the notes.

```
CBL LIB,APOST,XOPTS(CICS,DLI)          IDENTIFICATION DIVISION.
PROGRAM-ID. SAMPLE.
ENVIRONMENT DIVISION.
CONFIGURATION SECTION.
.* SOURCE-COMPUTER. IBM-370.
.* OBJECT-COMPUTER. IBM-370.
DATA DIVISION.
WORKING-STORAGE SECTION.
77 SEGKEYA          PIC X(4).
77 SEGKEYB          PIC X(4).
77 SEGKEYC          PIC X(4).
77 SEGKEY1          PIC X(4).
77 SEGKEY2          PIC X(4).
77 SEGKEY3          PIC X(4).
77 SEGKEY4          PIC X(4).
77 CONKEYB          PIC X(8).
77 SEGNAME          PIC X(8).
77 SEGLEN           COMP PIC S9(4).
77 PCBNUM           COMP PIC S9(4).
01 AREAA            PIC X(80).
*   DEFINE SEGMENT I/O AREA
01 AREAB            PIC X(80).
01 AREAC            PIC X(80).
01 AREAG            PIC X(250).
01 AREASTAT         PIC X(360).
*   COPY MAPSET.
PROCEDURE DIVISION.
*
* *****
*   INITIALIZATION
*   HANDLE ERROR CONDITIONS IN ERROR ROUTINE
*   HANDLE ABENDS (DLI ERROR STATUS CODES) IN ABEND ROUTINE
*   RECEIVE INPUT MESSAGE
* *****
*
*   EXEC CICS HANDLE CONDITION ERROR(ERRORS) END-EXEC.
*
*   EXEC CICS HANDLE ABEND LABEL(ABENDS) END-EXEC.
*
*   EXEC CICS RECEIVE MAP ('SAMPMAP') MAPSET('MAPSET') END-EXEC.
*   ANALYZE INPUT MESSAGE AND PERFORM NON-DLI PROCESSING
*
* *****
*   SCHEDULE PSB NAMED 'SAMPLE1'
* *****
*
*   EXEC DLI SCHD PSB(SAMPLE1) END-EXEC.
*   PERFORM TEST-DIB THRU OK.
*
* *****
*   RETRIEVE ROOT SEGMENT AND ALL ITS DEPENDENTS
* *****
*
*   MOVE 'A300' TO SEGKEYA.
*   EXEC DLI GU USING PCB(1) SEGMENT(SEGA) INTO(AREAA)
*       SEGLENGTH(80) WHERE (KEYA=SEGKEYA)
*       FIELDLENGTH(4)
*   END-EXEC.
*   PERFORM TEST-DIB THRU OK.
GNPLOOP.
*   EXEC DLI GNP USING PCB(1) INTO(AREAG) SEGLENGTH(250)
*   END-EXEC.
*   IF DIBSTAT EQUAL TO 'GE' THEN GO TO LOOPDONE.
*   PERFORM TEST-DIB THRU OK.
```

```

GO TO GNPLOOP.
LOOPDONE.
*
* *****
* INSERT NEW ROOT SEGMENT
* *****
*
    MOVE 'DATA FOR NEW SEGMENT INCLUDING KEY' TO AREAA.
    EXEC DLI ISRT USING PCB(1) SEGMENT(SEGA) FROM(AREAA)
      SEGLENGTH(80) END-EXEC.
    PERFORM TEST-DIB THRU OK.
*
* *****
* RETRIEVE 3 SEGMENTS IN PATH AND REPLACE THEM
* *****
*
    MOVE 'A200' TO SEGKEYA.
    MOVE 'B240' TO SEGKEYB.
    MOVE 'C241' TO SEGKEYC.
    EXEC DLI GU USING PCB(1)
      SEGMENT(SEGA) WHERE(KEYA=SEGKEYA) FIELDLENGTH(4)
      INTO(AREAA)
      SEGLENGTH(80)
      SEGMENT(SEGB) WHERE(KEYB=SEGKEYB) FIELDLENGTH(4)
      INTO(AREAB)
      SEGLENGTH(80)
      SEGMENT(SEGC) WHERE(KEYC=SEGKEYC) FIELDLENGTH(4)
      INTO(AREAC)
      SEGLENGTH(80)
    END-EXEC.
    PERFORM TEST-DIB THRU OK.
*
    UPDATE FIELDS IN THE 3 SEGMENTS
    EXEC DLI REPL USING PCB(1)
      SEGMENT(SEGA) FROM(AREAA) SEGLENGTH(80)
      SEGMENT(SEGB) FROM(AREAB) SEGLENGTH(80)
      SEGMENT(SEGC) FROM(AREAC) SEGLENGTH(80)
    END-EXEC.
    PERFORM TEST-DIB THRU OK.
*
* *****
* INSERT NEW SEGMENT USING CONCATENATED KEY TO QUALIFY PARENT
* *****
*
    MOVE 'DATA FOR NEW SEGMENT INCLUDING KEY' TO AREAC.
    MOVE 'A200B240' TO CONKEYB.
    EXEC DLI ISRT USING PCB(1)
      SEGMENT(SEGB) KEYS(CONKEYB) KEYLENGTH(8)
      SEGMENT(SEGC) FROM(AREAC) SEGLENGTH(80)
    END-EXEC.
    PERFORM TEST-DIB THRU OK.
*
* *****
* RETRIEVE SEGMENT DIRECTLY USING CONCATENATED KEY
* AND THEN DELETE IT AND ITS DEPENDENTS
* *****
*
    MOVE 'A200B230' TO CONKEYB.
    EXEC DLI GU USING PCB(1)
      SEGMENT(SEGB)
      KEYS(CONKEYB) KEYLENGTH(8)
      INTO(AREAB) SEGLENGTH(80)
    END-EXEC.
    PERFORM TEST-DIB THRU OK.
    EXEC DLI DLET USING PCB(1)
      SEGMENT(SEGB) SEGLENGTH(80) FROM(AREAB) END-EXEC.
    PERFORM TEST-DIB THRU OK.
*
* *****
* RETRIEVE SEGMENT BY QUALIFYING PARENT WITH CONCATENATED KEY,
* OBJECT SEGMENT WITH WHERE OPTION,
* AND THEN SET PARENTAGE
*
* USE VARIABLES FOR PCB INDEX, SEGMENT NAME, AND SEGMENT LENGTH
* *****
*
    MOVE 'A200B230' TO CONKEYB.
    MOVE 'C520' TO SEGKEYC.
    MOVE 'SEGA' TO SEGNAME.
    MOVE 80 TO SEGLEN.
    MOVE 1 TO PCBNUM.
    EXEC DLI GU USING PCB(PCBNUM)
      SEGMENT((SEGNAME))

```

```

        KEYS(CONKEYB) KEYLENGTH(8) SETPARENT
        SEGMENT(SEGC) INTO(AREAC) SEGLENGTH(SEGLEN)
        WHERE(KEYC=SEGKEYC) FIELDLENGTH(4) END-EXEC.
        PERFORM TEST-DIB THRU OK.
*
* *****
* RETRIEVE DATABASE STATISTICS
* *****
*
        EXEC DLI STAT USING PCB(1) INTO(AREASTAT)
        VSAM FORMATTED LENGTH(360) END-EXEC.
        PERFORM TEST-DIB THRU OK.
*
* *****
* RETRIEVE ROOT SEGMENT USING BOOLEAN OPERATORS
* *****
*
        MOVE 'A050' TO SEGKEY1.
        MOVE 'A150' TO SEGKEY2.
        MOVE 'A275' TO SEGKEY3.
        MOVE 'A350' TO SEGKEY4.
        EXEC DLI GU USING PCB(1) SEGMENT(SEGA) INTO(AREAA)
        SEGLENGTH(80) FIELDLENGTH(4,4,4,4)
        WHERE(KEYA > SEGKEY1 AND KEYA < SEGKEY2 OR
        KEYA > SEGKEY3 AND KEYA < SEGKEY4)
        END-EXEC.
        PERFORM TEST-DIB THRU OK.
*
* *****
* TERMINATE PSB WHEN DLI PROCESSING IS COMPLETED
* *****
*
        EXEC DLI TERM END-EXEC.
*
* *****
* SEND OUTPUT MESSAGE
* *****
*
        EXEC CICS SEND MAP('SAMPMP') MAPSET('MAPSET') END-EXEC.
        EXEC CICS WAIT TERMINAL END-EXEC.
*
* *****
* COMPLETE TRANSACTION AND RETURN TO CICS
* *****
*
        EXEC CICS RETURN END-EXEC.
*
* *****
* CHECK STATUS IN DIB
* *****
*
        TEST-DIB.
        IF DIBSTAT EQUAL TO ' ' THEN GO TO OK.
        OK.
        ERRORS.
* HANDLE ERROR CONDITIONS
        ABENDS.
* HANDLE ABENDS INCLUDING DLI ERROR STATUS CODES

```

8

9

Notes for the sample COBOL code:

- 1** For a CICS online program containing EXEC DLI commands, you must specify the DLI and CICS options. For a batch or BMP program containing EXEC DLI, you must specify only the DLI option.
- 2** Define each of the areas the program uses—I/O areas, key feedback areas, and segment name areas—as 77- or 01-level working storage entries.
- 3** Define an I/O area for each segment you retrieve, add, or replace (in a single command).
- 4** Do not code EXEC CICS commands in a batch or BMP program.
- 5** For CICS online programs, you use a SCHD PSB command to obtain a PSB. You do *not* schedule a PSB in a batch or BMP program.
- 6** This GU command retrieves the first occurrence of SEGA with a key of A300. KEYLENGTH and SEGLENGTH are optional for IBM COBOL for z/OS & VM (and VS COBOL II). For COBOL V4 and OS/VS COBOL, KEYLENGTH and SEGLENGTH are required.
- 7** This GU command is an example of a path command. You must use a separate I/O area for each segment you retrieve.

8 For a CICS online program, the TERM command terminates the PSB scheduled earlier. You do *not* terminate the PSB in a batch or BMP program.

9 After issuing each command, you should check the status code in the DIB.

Coding a program in PL/I

The following sample PL/I program shows how the different parts of a command-level program fit together, and how the EXEC DLI commands are coded in a CICS online program.

Except for a few commands, this program applies to batch, BMP, and CICS programs. Any differences are highlighted in the notes for the sample PL/I code. The numbering on the right of the sample code references those notes.

```
*PROCESS INCLUDE,GN,XOPTS(CICS,DLI);
SAMPLE: PROCEDURE OPTIONS(MAIN);
DCL      SEGKEYA          CHAR (4);
DCL      SEGKEYB          CHAR (4);
DCL      SEGKEYC          CHAR (4);
DCL      SEGKEY1          CHAR (4);
DCL      SEGKEY2          CHAR (4);
DCL      SEGKEY3          CHAR (4);
DCL      SEGKEY4          CHAR (4);
DCL      CONKEYB          CHAR (8);
DCL      SEGNAME          CHAR (8);
DCL      PCBNUM           FIXED BIN (15);
DCL      AREAA            CHAR (80);
DCL      /* DEFINE SEGMENT I/O AREA
DCL      AREAB            CHAR (80);
DCL      AREAC            CHAR (80);
DCL      AREAG            CHAR (250);
DCL      AREASTAT         CHAR (360);
%INCLUDE MAPSET
/*
/*
/* *****
/*  INITIALIZATION
/*  HANDLE ERROR CONDITIONS IN ERROR ROUTINE
/*  HANDLE ABENDS (DLI ERROR STATUS CODES) IN ABEND PROGRAM
/*  RECEIVE INPUT MESSAGE
/*  *****
/*
EXEC CICS HANDLE CONDITION ERROR(ERRORS);
/*
EXEC CICS HANDLE ABEND PROGRAM('ABENDS');
/*
EXEC CICS RECEIVE MAP ('SAMPMP') MAPSET('MAPSET');
/* ANALYZE INPUT MESSAGE AND PERFORM NON-DLI PROCESSING
/*
/* *****
/*  SCHEDULE PSB NAMED 'SAMPLE1'
/*  *****
/*
EXEC DLI SCHD PSB(SAMPLE1);
CALL TEST_DIB;

/* *****
/*  RETRIEVE ROOT SEGMENT AND ALL ITS DEPENDENTS
/*  *****
/*
SEGKEYA = 'A300';
EXEC DLI GU USING PCB(1) SEGMENT(SEGA) INTO(AREAA)
WHERE(KEYA=SEGKEYA);
CALL TEST_DIB;
GNPLOOP:
EXEC DLI GNP USING PCB(1) INTO(AREAG);
IF DIBSTAT = 'GE' THEN GO TO LOOPDONE;
CALL TEST_DIB;
GO TO GNPLOOP;
LOOPDONE:
/*
/* *****
/*  INSERT NEW ROOT SEGMENT
/*  *****
/*
AREAA = 'DATA FOR NEW SEGMENT INCLUDING KEY';
EXEC DLI ISRT USING PCB(1) SEGMENT(SEGA) FROM(AREAA);
CALL TEST_DIB;
```



```

/* */
/* ***** */
/* RETRIEVE 3 SEGMENTS IN PATH AND REPLACE THEM */
/* ***** */
/* */
SEGKEYA = 'A200';
SEGKEYB = 'B240';
SEGKEYC = 'C241';
EXEC DLI GU USING PCB(1)
    SEGMENT(SEGA) WHERE(KEYA=SEGKEYA)
        INTO(AREAA)
    SEGMENT(SEGB) WHERE(KEYB=SEGKEYB)
        INTO(AREAB)
    SEGMENT(SEGC) WHERE(KEYC=SEGKEYC)
        INTO(AREAC);
CALL TEST_DIB;
/* UPDATE FIELDS IN THE 3 SEGMENTS */
EXEC DLI REPL USING PCB(1)
    SEGMENT(SEGA) FROM(AREAA)
    SEGMENT(SEGB) FROM(AREAB)
    SEGMENT(SEGC) FROM(AREAC);
CALL TEST_DIB;
/* */
/* ***** */
/* INSERT NEW SEGMENT USING CONCATENATED KEY TO QUALIFY PARENT */
/* ***** */
/* */
AREAC = 'DATA FOR NEW SEGMENT INCLUDING KEY';
CONKEYB = 'A200B240';
EXEC DLI ISRT USING PCB(1)
    SEGMENT(SEGB) KEYS(CONKEYB)
    SEGMENT(SEGC) FROM(AREAC);
CALL TEST_DIB;
/* */
/* ***** */
/* RETRIEVE SEGMENT DIRECTLY USING CONCATENATED KEY */
/* AND THEN DELETE IT AND ITS DEPENDENTS */
/* ***** */
/* */
CONKEYB = 'A200B230';
EXEC DLI GU USING PCB(1)
    SEGMENT(SEGB)
    KEYS(CONKEYB)
    INTO(AREAB);
CALL TEST_DIB;
EXEC DLI DLET USING PCB(1)
    SEGMENT(SEGB) FROM(AREAB);
CALL TEST_DIB;
/* */
/* ***** */
/* RETRIEVE SEGMENT BY QUALIFYING PARENT WITH CONCATENATED KEY, */
/* OBJECT SEGMENT WITH WHERE OPTION */
/* AND THEN SET PARENTAGE */
/* ***** */
/* USE VARIABLES FOR PCB INDEX, SEGMENT NAME */
/* ***** */
/* */
CONKEYB = 'A200B230';
SEGNAME = 'SEGA';
SEGKEYC = 'C520';
PCBNUM = 1;
EXEC DLI GU USING PCB(PCBNUM)
    SEGMENT((SEGNAME))
    KEYS(CONKEYB) SETPARENT
    SEGMENT(SEGC) INTO(AREAC)
    WHERE(KEYC=SEGKEYC);
CALL TEST_DIB;
/* */
/* ***** */
/* RETRIEVE DATABASE STATISTICS */
/* ***** */
/* */
EXEC DLI STAT USING PCB(1) INTO(AREASTAT) VSAM FORMATTED;
CALL TEST_DIB;
/* */
/* ***** */
/* RETRIEVE ROOT SEGMENT USING BOOLEAN OPERATORS */
/* ***** */
/* */
SEGKEY1 = 'A050';
SEGKEY2 = 'A150';
SEGKEY3 = 'A275';

```

```

SEGKEY4 = 'A350';
EXEC DLI GU USING PCB(1) SEGMENT(SEGA) INTO(AREAA)
      WHERE(KEYA &A; SEGKEY1 AND KEYA &A1; SEGKEY2 OR
      KEYA &A; SEGKEY3 AND KEYA &A1; SEGKEY4);
CALL TEST_DIB;
/*
/* *****
/* TERMINATE PSB WHEN DLI PROCESSING IS COMPLETED
/* *****
/*
EXEC DLI TERM;

/*
/* *****
/* SEND OUTPUT MESSAGE
/* *****
/*
EXEC CICS SEND MAP('SAMPMAP') MAPSET('MAPSET');
EXEC CICS WAIT TERMINAL;
/*
/* *****
/* COMPLETE TRANSACTION AND RETURN TO CICS
/* *****
/*
EXEC CICS RETURN;
/*
/* *****
/* CHECK STATUS IN DIB
/* *****
/*
TEST_DIB: PROCEDURE;
  IF DIBSTAT = ' ' RETURN;

  /* HANDLE DLI STATUS CODES REPRESENTING EXCEPTIONAL CONDITIONS
  /*
OK:
END TEST_DB;
ERRORS:
  /* HANDLE ERROR CONDITIONS
  /*
END SAMPLE;

```

9

4

4

10

Notes to the sample PL/I code:

- 1** For a CICS online program containing EXEC DLI commands, you must specify the DLI and CICS options. For a batch or BMP program containing EXEC DLI, you must specify only the DLI option.
- 2** Define, in automatic storage, each of the areas; I/O areas, key feedback areas, and segment name areas.
- 3** Define an I/O area for each segment you retrieve, add, or replace in a single command.
- 4** Do not code EXEC CICS commands in a batch or BMP program.
- 5** For CICS online programs, you use a SCHD PSB command to obtain a PSB. You do *not* schedule a PSB in a batch or BMP program.
- 6** This GU command retrieves the first occurrence of SEGA with a key of A300. Notice that you do not need to include the KEYLENGTH and SEGLENGTH options.
- 7** This GNP command retrieves all dependents under segment SEGA. The GE status code indicates that no more dependents exist.
- 8** This GU command is an example of a path command. You must use a separate I/O area for each segment you retrieve.
- 9** For a CICS online program, the TERM command terminates the PSB scheduled earlier. You do *not* terminate the PSB in a batch or BMP program.
- 10** After issuing each command, you should check the status code in the DIB.

Coding a program in C

the following sample C program shows how the different parts of a command-level program fit together, and how the EXEC DLI commands are coded in a CICS online program.

Except for a few commands, this program applies to batch, BMP, and CICS programs. Any differences are highlighted in the notes for the sample C code. The numbering on the right of the sample code references those notes.

```
#include < string.h>
#include < stdio.h >

char DIVIDER[120] = "-----\
-----\
";
char BLANK[120] = "\
\0";
char BLAN2[110] = "\
\0";

char SCHED[120] = "Schedule PSB(PC3COCHD) " \0";
char GN1[120] = "GN using PCB(2) Segment(SE2ORDER) check dibstat \
is blank";
char GNP1[120] = "GNP using PCB(2) check dibstat = GK or blank \
(or GE for last GNP)";
char GU1[120] = "GU using PCB(2) Segment(SE2ORDER) where(\
FE20GREF=000000') check dibstat blank";
char GU2[120] = "GU using PCB(2) Segment(SE2ORDER) where(\
FE20GREF=000999') check dibstat blank";
char REP1[120] = "REPLACE using PCB(2) Segment(SE2ORDER) check \
dibstat is blank";
char DEL1[120] = "DELETE using PCB(2) Segment(SE2ORDER) check \
dibstat is blank";
char INS1[120] = "INSERT using PCB(2) Segment(SE2ORDER) where(\
FE20GREF='000999') check dibstat is blank";
char TERM[120] = "TERM - check dibstat is blank";
char STAT[120] = "STAT USING PCB(2) VSAM FORMATTED";
char DATAB[6] = "000999";
char DATAC[114] = " REGRUN TEST INSERT NO1.";
char START[120] = "PROGXIV STARTING";
char OKMSG[120] = "PROGXIV COMPLETE";
int TLINE = 120;
int L11 = 11;
int L360 = 11;
struct {
    char NEWSEGB[6];
    char NEWSEGC[54];
} NEWSEG;
char OUTLINE[120];
struct {
    char OUTLINA[9];
    char OUTLINB[111];
} OUTLIN2;
struct {
    char OUTLINX[9];
    char OUTLINY[6];
    char OUTLINZ[105];
} OUTLIN3;
char GUIOA[60];
char GNIOA[60];
struct {
    char ISRT1[6];
    char ISRT2[54];
} ISRTIOA;
struct {
    char REPLIO1[6];
    char REPLIO2[54];
} REPLIOA;
struct {
    char DLET1[6];
    char DLET2[54];
} DLETIOA;
struct {
    char STATA1[120];
    char STATA2[120];
    char STATA3[120];
} STATAAREA;
struct {
    char DHPART[2];
    char RETCODE[2]
} DHABCODE;
```

1
2

3

4

```

main()
{
    EXEC CICS ADDRESS EIB(dfheiptr);
    strcpy(OUTLINE,DIVIDER);
    SENDLINE();
    strcpy(OUTLINE,START);
    SENDLINE();
    /*
    /* SCHEDULE PSB
    /*
    strcpy(OUTLINE,SCHED);
    SENDLINE();
    EXEC DLI SCHEDULE PSB(PC3COCHD);
    SENDSTAT();
    TESTDIB();
    /*
    /* ISSUE GU REQUEST
    /*
    strcpy(OUTLINE,GU1);
    SENDLINE();
    EXEC DLI GET UNIQUE USING PCB(2)
        SEGMENT(SE2ORDER)
        WHERE(FE2OGREF>="000000")
        INTO(&GUIOA) SEGLLENGTH(60);
    strcpy(OUTLIN2.OUTLINA,"SE2ORDER=");
    strcpy(OUTLIN2.OUTLINB,GUIOA);
    SENDLIN2();
    SENDSTAT();
    TESTDIB();
    /*
    /* ISSUE GNP REQUEST
    /*
do {
    strcpy(OUTLINE,GNP1);
    SENDLINE();
    EXEC DLI GET NEXT IN PARENT USING PCB(2)
        INTO(&GNIOA) SEGLLENGTH(60);
    strcpy(OUTLIN2.OUTLINA,"SEGMENT=");
    strcpy(OUTLIN2.OUTLINB,GNIOA);
    SENDLIN2();
    SENDSTAT();
    if (strncmp(dibptr->dibstat,"GE",2) != 0)
        TESTDIB();
} while (strncmp(dibptr->dibstat,"GE",2) != 0);
/*
/* ISSUE GN REQUEST
/*
    strcpy(OUTLINE,GN1);
    SENDLINE();
    EXEC DLI GET NEXT USING PCB(2)
        SEGMENT(SE2ORDER)
        INTO(&GNIOA) SEGLLENGTH(60);
    strcpy(OUTLIN2.OUTLINA,"SE2ORDER=");
    strcpy(OUTLIN2.OUTLINB,GNIOA);
    SENDLIN2();
    SENDSTAT();
    TESTDIB();
/*
/* INSERT SEGMENT
/*
    strcpy(OUTLINE,INS1);
    SENDLINE();
    strcpy(NEWSEG.NEWSEGB,DATAB);
    strcpy(NEWSEG.NEWSEGC,DATAC);
    strcpy(ISRTIOA.ISRT1,NEWSEG.NEWSEGB);
    strcpy(ISRTIOA.ISRT2,NEWSEG.NEWSEGC);
    strcpy(OUTLIN3.OUTLINX,"ISRT SEG=");
    strcpy(OUTLIN3.OUTLINY,ISRTIOA.ISRT1);
    strcpy(OUTLIN3.OUTLINZ,ISRTIOA.ISRT2);
    SENDLIN3();
    EXEC DLI ISRT USING PCB(2)
        SEGMENT(SE2ORDER)
        FROM(&ISRTIOA) SEGLLENGTH(60);
    SENDSTAT();
    if (strncmp(dibptr->dibstat,"II",2) == 0)
        strcpy(dibptr->dibstat," ",2);
    TESTDIB();
/*
/* ISSUE GN REQUEST
/*
    strcpy(OUTLINE,GN1);

```

```

SENDLINE();
EXEC DLI GET NEXT USING PCB(2)
    SEGMENT(SE2ORDER)
    INTO(&GNIOA) SEGLNGTH(60);
strcpy(OUTLIN2.OUTLINA,"SE2ORDER=");
strcpy(OUTLIN2.OUTLINB,GNIOA);
SENDLIN2();
SENDSTAT();
TESTDIB();
/*
/* GET INSERTED SEGMENT TO BE REPLACED
/*
    strcpy(OUTLINE,GU2);
    SENDLINE();
    EXEC DLI GET UNIQUE USING PCB(2)
        SEGMENT(SE2ORDER)
        WHERE(FE20GREF="000999")
        INTO(&ISRTIOA) SEGLNGTH(60);
    strcpy(OUTLIN3.OUTLINX,"ISRT SEG=");
    strcpy(OUTLIN3.OUTLINY,ISRTIOA.ISRT1);
    strcpy(OUTLIN3.OUTLINZ,ISRTIOA.ISRT2);
    SENDLIN3();
    SENDSTAT();
    TESTDIB();
/*
/* REPLACE SEGMENT
/*
    strcpy(OUTLINE,REP1);
    SENDLINE();
    strcpy(REPLIOA.REPLIO1,DATAB);
    strcpy(REPLIOA.REPLIO2,"REGRUN REPLACED SEGMENT NO1.");
    strcpy(OUTLIN3.OUTLINX,"REPL SEG=");
    strcpy(OUTLIN3.OUTLINY,REPLIOA.REPLIO1);
    strcpy(OUTLIN3.OUTLINZ,REPLIOA.REPLIO2);
    SENDLIN3();
    EXEC DLI REPLACE USING PCB(2)
        SEGMENT(SE2ORDER)
        FROM(&REPLIOA) SEGLNGTH(60);
    SENDSTAT();
    TESTDIB();
/*
/* ISSUE GN REQUEST
/*
    strcpy(OUTLINE,GN1);
    SENDLINE();
    EXEC DLI GET NEXT USING PCB(2)
        SEGMENT(SE2ORDER)
        INTO(&GNIOA) SEGLNGTH(60);
    strcpy(OUTLIN2.OUTLINA,"SE2ORDER=");
    strcpy(OUTLIN2.OUTLINB,GNIOA);
    SENDLIN2();
    SENDSTAT();
    TESTDIB();
/*
/* GET REPLACED SEGMENT
/*
    strcpy(OUTLINE,GU2);
    SENDLINE();
    EXEC DLI GET UNIQUE USING PCB(2)
        SEGMENT(SE2ORDER)
        WHERE(FE20GREF="000999")
        INTO(&REPLIOA) SEGLNGTH(60);
    strcpy(OUTLIN3.OUTLINX,"REPL SEG=");
    strcpy(OUTLIN3.OUTLINY,REPLIOA.REPLIO1);
    strcpy(OUTLIN3.OUTLINZ,REPLIOA.REPLIO2);
    SENDLIN3();
    SENDSTAT();
    TESTDIB();
/*
/* ISSUE DELETE REQUEST
/*
    strcpy(OUTLINE,DEL1);
    SENDLINE();
    strcpy(DLETIOA.DLET1,REPLIOA.REPLIO1);
    strcpy(DLETIOA.DLET2,REPLIOA.REPLIO2);
    strcpy(OUTLIN3.OUTLINX,"DLET SEG=");
    strcpy(OUTLIN3.OUTLINY,DLETIOA.DLET1);
    strcpy(OUTLIN3.OUTLINZ,DLETIOA.DLET2);
    SENDLIN3();
    EXEC DLI DELETE USING PCB(2)
        SEGMENT(SE2ORDER)
        FROM(&DLETIOA) SEGLNGTH(60);

```

```

        SENDSTAT();
        TESTDIB();
    /*                                     */
    /* ISSUE STAT REQUEST                 */
    /*                                     */
        strcpy(OUTLINE,STAT);
        SENDLINE();
        EXEC DLI STAT USING PCB(2)
            VSAM FORMATTED
            INTO(&STATAREA);
        SENDSTT2();
        TESTDIB();
    /*                                     */
    /* ISSUE TERM REQUEST                 */
    /*                                     */
        strcpy(OUTLINE,TERM);
        SENDLINE();
        EXEC DLI TERM;
        SENDSTAT();
        TESTDIB();
        strcpy(OUTLINE,DIVIDER);
        SENDLINE();
        SENDOK();
    /*                                     */
    /* RETURN TO CICS                     */
    /*                                     */
        EXEC CICS RETURN;
    }
    /*                                     */
    /*                                     */
    /*                                     */
    SENDLINE()
    {
        EXEC CICS SEND FROM(OUTLINE) LENGTH(120);
        EXEC CICS WRITEQ TD QUEUE("PRIM") FROM(OUTLINE) LENGTH(TLINE);
        strcpy(OUTLINE,BLANK);
        return;
    }

    SENDLIN2()
    {
        EXEC CICS SEND FROM(OUTLIN2) LENGTH(120);
        EXEC CICS WRITEQ TD QUEUE("PRIM") FROM(OUTLIN2) LENGTH(TLINE);
        strcpy(OUTLIN2.OUTLINA,BLANK,9);
        strcpy(OUTLIN2.OUTLINB,BLANK,111);
        return;
    }

    SENDLIN3()
    {
        EXEC CICS SEND FROM(OUTLIN3) LENGTH(120);
        EXEC CICS WRITEQ TD QUEUE("PRIM") FROM(OUTLIN3) LENGTH(TLINE);
        strcpy(OUTLIN3.OUTLINX,BLANK,9);
        strcpy(OUTLIN3.OUTLINY,BLANK,6);
        strcpy(OUTLIN3.OUTLINZ,BLANK,105);
        return;
    }

    SENDSTAT()
    {
        strncpy(OUTLIN2.OUTLINA,BLANK,9);
        strncpy(OUTLIN2.OUTLINB,BLANK,110);
        strcpy(OUTLIN2.OUTLINA," DIBSTAT=");
        strcpy(OUTLIN2.OUTLINB,dibptr->dibstat);
        EXEC CICS SEND FROM(OUTLIN2) LENGTH(11);
        EXEC CICS WRITEQ TD QUEUE("PRIM") FROM(OUTLIN2) LENGTH(L11);
        strcpy(OUTLINE,DIVIDER);
        SENDLINE();
        return;
    }

    SENDSTT2()
    {
        strncpy(OUTLIN2.OUTLINA,BLANK,9);
        strncpy(OUTLIN2.OUTLINB,BLANK,110);
        strcpy(OUTLIN2.OUTLINA," DIBSTAT=");
        strcpy(OUTLIN2.OUTLINB,dibptr->dibstat);
        EXEC CICS SEND FROM(STATAREA) LENGTH(360);
        EXEC CICS WRITEQ TD QUEUE("PRIM") FROM(STATAREA)
            LENGTH(L360);
        return;
    }

```

18

19

20

```

SENDOK()
{
    EXEC CICS SEND FROM(OKMSG) LENGTH(120);
    EXEC CICS WRITEQ TD QUEUE("PRIM") FROM(OKMSG) LENGTH(TLINE);
    return;
}

TESTDIB()
{
    if (strcmp(dibptr->dibstat," ",2) == 0)
        return;
    else if (strcmp(dibptr->dibstat,"GK",2) == 0)
        return;
    else if (strcmp(dibptr->dibstat,"GB",2) == 0)
        return;
    else if (strcmp(dibptr->dibstat,"GE",2) == 0)
        return;
    else
    {
        EXEC CICS ABEND ABCODE("PETE");
        EXEC CICS RETURN;
    }
    return;
}

```

21

22

Notes for the sample C code:

- 1 You must include a standard header file `string.h` to gain access to string manipulation facilities.
- 2 You must include standard header file `stdio.h` to access the standard I/O library.
- 3 Define DL/I messages.
- 4 Define the I/O areas.
- 5 Program start.
- 6 Define PSB PC3COCHD.
- 7 Issue the first command. Retrieves the first occurrence of segment SE2ORDER and puts it into array OUTLIN2.
- 8 Issue the GNP command to get the next segment and put it into array OUTLIN2.
- 9 GE status codes indicate no more segments to get.
- 10 Get next segment SE2ORDER and put it into the array OUTLIN2.
- 11 Insert segment into array OUTLIN3.
- 12 Issue GN to retrieve next segment and put it into array OUTLIN2.
- 13 Get next segment that will be replaced and put it into OUTLIN3.
- 14 Replace the segment and put it into array OUTLIN3.
- 15 Get next segment and put it into array OUTLIN2.
- 16 Get the replaced segment and put it into array OUTLIN3.
- 17 Issue DELETE command after putting content of segment into array OUTLIN3.
- 18 Issue STAT REQUEST command.
- 19 Issue TERM command.
- 20 Output processing.
- 21 Check return code.
- 22 Do not code EXEC CICS commands in a batch or BMP program.

Preparing your EXEC DLI program for execution

You must translate, compile, and bind your EXEC DLI program before it can be executed.

You can use CICS-supplied procedures to translate, compile, and bind your program. The procedure you use depends on the type of program (batch, BMP, or CICS online) and the language it is written in (COBOL, PL/I, or assembler language).

The steps for preparing your program for execution are as follows:

1. Run the CICS command language translator to translate the EXEC DLI and EXEC CICS commands. COBOL, PL/I, and assembler language programs have separate translators.
2. Compile your program.
3. Bind:
 - An online program with the appropriate CICS interface module
 - A batch or BMP program with the IMS interface module.

Translator, compiler, and binder options required for EXEC DLI

To execute your EXEC DLI program, you must set the required translator, compile, and binder options.

Translator options required for EXEC DLI

Even when you use the CICS-supplied procedures for preparing your program, you must supply certain translator options.

For a CICS online program containing EXEC DLI commands, you must specify the DLI and CICS options. For a batch or BMP program containing EXEC DLI commands, you must specify the DLI option.

You can also specify the options on the EXEC job control statement that invokes the translator; if you use both methods, the CBL and *PROCESS statement overrides those in the EXEC statement. For more information on the translator options, see *CICS Transaction Server for z/OS CICS Application Programming Guide*.

You must ensure that the translator options you use in a COBOL program do not conflict with the COBOL compiler options.

Compiler options required for EXEC DLI

To compile your batch COBOL program, you may have to use different compiler options, depending on which COBOL compiler was chosen. For information on which compiler options should be used for a CICS program, see *CICS Transaction Server for z/OS CICS Application Programming Guide*.

Binder options required for EXEC DLI

If the compiler being used supports it, you can link a program written with EXEC commands as AMODE(31) RMODE(ANY).

Chapter 30. Defining application program elements

Define application program elements by using the EXEC DLI commands with the application interface block (AIB) and DL/I interface block (DIB), and by defining feedback and I/O areas.

Specifying an application interface block (AIB)

EXEC DLI commands can use the AIB interface.

For example, using the AIB interface, the format for the GU command would be EXEC DLI GU AIB(aib), instead of EXEC DLI GU USING PCB(n) using the PCB format.

With IBM CICS Transaction Server for z/OS, the EXEC DLI commands are supported in the AIB format (as well as the PCB format). The AIB-only commands ICMD, RCMD, and GMSG are supported by using the EXEC DLI interface.

The CICS EDF (Execution Diagnostic Facility) debugging transaction supports AIB EXEC DLI requests, just as it handles PCB type requests.

AIB mask

The AIB mask must be supplied by the application and referenced in the EXEC call instead of the PCB number (for example, EXEC DLI GU AIB(aib)).

The DIBSTAT field is set with a valid STATUS code when AIBRETRN = X'00000000' or X'00000900'. Applications should test AIBRETRN for any other values and respond accordingly.

CICS restrictions with AIB support

Restrictions due to function shipping include:

- The AIBLEN field must be between 128 and 256 bytes. 128 bytes is recommended.
- LIST=NO must not be specified on any PCBs in the PSB.

Related reference

[“EXEC DLI commands for an application program” on page 483](#)

The EXEC DLI commands in your application program is used together with a program specification block (PSB) and different kinds of program communication blocks (PCBs).

Specifying the DL/I interface block (DIB)

Each time your program executes a DL/I command, DL/I returns a status code and other information to your program through the DL/I interface block (DIB), which is a subset of IMS PCB. Your program should check the status code to make sure the command executed successfully.

Each program's working storage contains its own DIB. The contents of the DIB reflect the status of the last DL/I command executed in that program. If the information in your program's DIB is required by another program used by your transaction, you must pass the information to that program.

To access fields in the DIB, use labels that are automatically generated in your program by the translator.

Restriction: These labels are reserved; you must not redefine them.

In your COBOL, PL/I, assembler language, and C programs, some variable names are mandatory.

For a COBOL program:

```
DIBVER      PICTURE X(2)
DIBSTAT     PICTURE X(2)
DIBSEGM     PICTURE X(8)
DIBFIL01    PICTURE X
DIBFIL02    PICTURE X
```

```

DIBSEGLV    PICTURE X(2)
DIBKFBL     PICTURE S9(4) COMPUTATIONAL
DIBDBDNM    PICTURE X(8)
DIBDBORG    PICTURE X(8)
DIBFIL03    PIC X(6)

```

For a PL/I program:

```

DIBVER      CHAR(2)
DIBSTAT     CHAR(2)
DIBSEGM     CHAR(8)
DIBFIL01    CHAR(1)
DIBFIL02    CHAR(1)
DIBSEGLV    CHAR(2)
DIBKFBL     FIXED BINARY (15,0)
DIBDBDNM    CHAR(8)
DIBDBORG    CHAR(8)
DIBFIL03    CHAR(6)

```

For an assembler language program:

```

DIBVER      CL2
DIBSTAT     CL2
DIBSEGM     CL8
DIBFIL01    CL1
DIBFIL02    CL1
DIBSEGLV    CL2
DIBKFBL     H
DIBDBDNM    CL8
DIBDBORG    CL8
DIBFILR3    CL6

```

For a C program:

```

unsigned char    dibver    {2} ;
unsigned char    dibstat   {2} ;
unsigned char    dibsegm   {8} ;
unsigned char    dibfic01  ;
unsigned char    dibfic02  ;
unsigned char    dibseglv  {2} ;
signed short int dibkfb1   ;
unsigned char    dibdbdnm  {8} ;
unsigned char    dibdborg  {8} ;
unsigned char    dibfic03  {6} ;

```

The following notes explain the contents of each variable name. The name in parenthesis is the label used to access the contents.

1. **Translator Version (DIBVER)**

This is the version of the DIB format your program is using. (DIBVER is used for documentation and problem determination.)

2. **Status Codes (DIBSTAT)**

DL/I places a 2-character status code in this field after executing each DL/I command. This code describes the results of the command.

After processing a DL/I command, DL/I returns control to your program at the next sequential instruction following the command. The first thing your program should do after each command is to test the status code field and take appropriate action. If the command was completely successful, this field contains blanks.

The status codes that can be returned to this field (they are the only status codes returned to your program) are:

bb

(Blanks) The command was completely successful.

BA

For GU, GN, GNP, DLET, REPL, and ISRT commands. Data was unavailable.

BC

For DLET, REPL, and ISRT commands. A deadlock was detected.

FH

For GU, GN, GNP, DLET, REPL, ISRT, POS, CHKP, and SYMCHKP commands. The DEDB was inaccessible.

FW

For GU, GN, GNP, DLET, REPL, ISRT, and POS commands. More buffer space is required than normally allowed.

GA

For unqualified GN and GNP commands. DL/I returned a segment, but the segment is at a higher level in the hierarchy than the last segment that was returned.

GB

For GN commands. DL/I reached the end of the database trying to satisfy your GN command and did not return a segment to your program's I/O area.

GD

For ISRT commands. The program issued an ISRT command that did not have SEGMENT options for all levels above that of the segment being inserted.

GE

For GU, GN, GNP, ISRT, and STAT commands. DL/I was unable to find the segment you requested, or one or more of the parents of the segment you are trying to insert.

GG

For Get commands. DL/I returns a GG status code to a program with a processing option of GOT or GON when the segment that the program is trying to retrieve contains an invalid pointer.

GK

For unqualified GN and GNP commands. DL/I returned a segment that satisfies an unqualified GN or GNP request, but the segment is of a different segment type (but at the same level) than the last segment returned.

II

For ISRT commands. The segment you are trying to insert already exists in the database. This code can also be returned if you have not established a path for the segment before trying to insert it. The segment you are trying to insert might match a segment with the same key in another hierarchy or database record.

LB

For load programs only after issuing a LOAD command. The segment you are trying to load already exists in the database. DL/I returns this status code only for segments with key fields.

NI

For ISRT and REPL commands. The segment you are trying to insert or replace requires a duplicate entry to be inserted in a secondary index that does not allow duplicate entries. This status code is returned for batch programs that write log records to direct access storage. If a CICS program that does not log to disk encounters this condition, the program (transaction) is abnormally terminated.

TG

For TERM commands. The program tried to terminate a PSB when one was not scheduled.

The listed status codes (DIBSTAT) indicate exceptional conditions, and are the only status codes returned to your program. All other status codes indicate error conditions and cause your transaction or batch program to abnormally terminate. If you want to pass control to an error routine from your CICS program, you can use the CICS HANDLE ABEND command; the last 2 bytes of theabend code are the IMS status code that caused the abnormal termination. For more information on the HANDLE ABEND command, see the application programming reference manual for your version of CICS. Batch BMP programsabend withabend 1041.

3. Segment Name (DIBSEGM)

This is the name of the lowest-level segment successfully accessed. When a retrieval is successful, this field contains the name of the retrieved segment. If the retrieval is unsuccessful, this field contains the last segment, along the path to the requested segment, that satisfies the command.

After issuing an XRST command, this field is either set to blanks (indicating a successful normal start), or a checkpoint ID (indicating the checkpoint ID from which the program was restarted).

You should test this field after issuing any of the following commands:

- GN
- GNP
- GU
- ISRT
- LOAD
- RETRIEVE
- XRST

4. **Segment Level Number (DIBSEGLV)**

This is the hierarchic level of the lowest-level segment retrieved. When IMS DB retrieves the segment you have requested, IMS DB places, in character format, the level number of that segment in this field. If you are issuing a path command, IMS DB places the number of the lowest-level segment retrieved. If IMS DB is unable to find the segment you have requested, it gives the level number of the last segment it encountered that satisfied your command. This is the lowest segment on the last path that IMS DB encountered while searching for the segment you requested.

You should test this field after issuing any of the listed commands:

- GN
- GNP
- GU
- ISRT
- LOAD
- RETRIEVE

5. **Key Feedback Length (DIBKFBL)**

This is a halfword field that contains the length of the concatenated key when you use the KEYFEEDBACK option with get commands. If your key feedback area is not long enough to contain the concatenated key, the key is truncated, and this area indicates the actual length of the full concatenated key.

6. **Database Description Name (DIBDBDNM)**

This is the fullword field that contains the name of the DBD. The DBD is the DL/I control block that contains all information used to describe a database. The DIBDBDNM field is returned only on a QUERY command.

7. **Database Organization (DIBDBORG)**

This is the fullword field that names the type of database organization (HDAM, HIDAM, HISAM, HSAM, GSAM, SHSAM, INDEX, or DEDB) padded to the right with blanks. The DIBDBORG field is returned only on a QUERY command.

Defining a key feedback area

To retrieve the concatenated key of a segment, you must define an area into which the key is placed.

The concatenated key returned is that of the lowest-level segment retrieved. (The segment retrieved is indicated in the DIB by the DIBSEGM and DIBSEGLV fields.)

Specify the name of the area using the KEYFEEDBACK option on a GET command.

A concatenated key is made up of the key of a segment, plus the keys for all of its parents. For example, say you requested the concatenated key of the ILLNESS segment for January 2, 1988, for patient number 05142. **0514219880102** would be returned to your key feedback field. This number includes the key field of the ILLNESS segment, ILLDATE, concatenated to the key field of the PATIENT segment, PATNO.

If you define an area that is not long enough to contain the entire concatenated key, the key is truncated.

Defining I/O areas

Use I/O areas to pass segments back and forth between your program and the database.

The contents of an I/O area depends on the kind of command you are issuing:

- When you retrieve a segment, DL/I places the segment you requested in the I/O area.
- When you add a new segment, you build the new segment in the I/O area before issuing an ISRT command.
- Before you modify a segment, you first retrieve the segment into the I/O area then issue the DLET or REPL command.

Restriction: The I/O area must be long enough to contain the longest segment you retrieve from or add to the database. (Otherwise, you might experience storage overlap.) If you are retrieving, adding, or replacing multiple segments in one command, you must define an I/O area for each segment.

As an example of what a segment looks like in your I/O area, say that you retrieved the ILLNESS segment for Robert James, who came to the clinic on March 3, 1988. He was treated for strep throat. The data returned to your I/O area would look like this:

```
19880303STREPTHROA
```

COBOL I/O area

The I/O area in a COBOL program should be defined as a 01 level working storage entry. You can further define the area with 02 entries.

```
IDENTIFICATION DIVISION.  
:  
DATA DIVISION.  
WORKING-STORAGE SECTION.  
01   INPUT-AREA.  
    02 KEY PICTURE X(6).  
    02 FIELD PICTURE X(84).
```

PL/I I/O area

In PL/I, the name for the I/O area used in the DL/I call can be the name of a fixed-length character string, a major structure, a connected array, or an adjustable character string.

Restriction: The PL/I I/O area *cannot* be the name of a minor structure or a character string with the attribute VARYING. If you want to define it as a minor structure, you can use a pointer to the minor structure as the parameter.

Your program should define the I/O area as a fixed-length character string and pass the name of that string, or define it in one of the other ways described previously and then pass the pointer variable that points to that definition. If you want to use substructures or elements of an array, use the DEFINED or BASED attribute.

```
DECLARE    1 INPUT_AREA,  
          2 KEY CHAR(6),  
          2 FIELD CHAR(84);
```

Assembler language I/O area

The I/O area in an assembler language program is formatted as follows:

```
IOAREA  DS  0CL90  
KEY      DS  CL6  
FIELD    DS  CL84
```

Chapter 31. EXEC DLI commands for an application program

The EXEC DLI commands in your application program is used together with a program specification block (PSB) and different kinds of program communication blocks (PCBs).

Related reference

[“Specifying an application interface block \(AIB\)” on page 477](#)
EXEC DLI commands can use the AIB interface.

PCBs and PSB

A program specification block (PSB) used in a DBCTL environment can contain I/O PCBs, alternate PCBs, database PCBs (DB PCB), or GSAM PCBs.

I/O PCB

In a DBCTL environment, an I/O PCB is needed to issue DBCTL service requests. Unlike the other types of PCB, it is not defined with PSB generation, but if the application program is using an I/O PCB, this has to be indicated in the PSB scheduling request.

Alternate PCB

An alternate PCB defines a logical terminal and can be used instead of the I/O PCB when it is necessary to direct a response to a terminal. Alternate PCBs appear in PSBs used in a CICS-DBCTL environment, but are used only in an IMS DC environment. CICS applications using DBCTL cannot successfully issue commands that specify an alternate PCB, an MSDB PCB, or a GSAM PCB. However, a PSB that contains PCBs of these types can be scheduled successfully in a CICS-DBCTL environment.

Alternate PCBs are included in the PCB address list returned to a call level application program. In an EXEC DLI application program, the existence of alternate PCBs in the PSB affects the PCB number used in the PCB keyword.

DB PCB

A DB PCB is the PCB that defines an application program's interface to a database. One DB PCB is needed for each database view used by the application program. It can be a full-function PCB, a DEDB PCB, or an MSDB PCB.

GSAM PCB

A GSAM PCB defines an application program's interface for GSAM operations.

When using DBCTL, a CICS program receives, by default, a DB PCB as the first PCB in the parameter list passed to it after scheduling. However, when your application program can handle an I/O PCB, you indicate this using the SYSSERVE keyword on the SCHD command. The I/O PCB is then the first PCB in the parameter address list passed back to your application program.

I/O PCBs and alternate PCBs in various types of application programs

DB batch programs

Alternate PCBs are always included in the list of PCBs supplied to the program by DL/I irrespective of whether you have specified CMPAT=Y. The I/O PCB is returned depending on the CMPAT option.

If you specify CMPAT=Y, the PCB list contains the address of the I/O PCB, followed by the addresses of any alternate PCBs, followed by the addresses of any DB PCBs.

If you do not specify CMPAT=Y, the PCB list contains the addresses of any alternate PCBs followed by the addresses of the DB PCBs.

BMP programs, MPPs, and IFPs

I/O PCBs and alternate PCBs are always passed to BMP programs. I/O PCBs and alternate PCBs are also always passed to MPPs and to IFP application programs.

The PCB list contains the address of the I/O PCB, followed by the addresses of any alternate PCBs, followed by the addresses of the DB PCBs.

CICS programs with DBCTL

The first PCB always refers to the first DB PCB whether you specify the SYSSERVE keyword.

The following table summarizes the I/O PCB and alternate PCB information. The first column lists different DB environments, the second and third column specify if the I/O PCB or alternate PCB, respectively, is valid in the specified environment.

Table 87. Summary of PCB information

Environment	EXEC DLI: I/O PCB count included in PCB(n)	EXEC DLI: Alternate PCB count included in PCB(n)
CICS DBCTL ¹	No	No
CICS DBCTL ²	No	No
BMP	Yes	Yes
Batch ³	No	Yes
Batch ⁴	Yes	Yes

Notes:

1. SCHD command issued without the SYSSERVE option.
2. SCHD command issued with the SYSSERVE option for a CICS DBCTL command or for a function-shipped command which is satisfied by a remote CICS system using DBCTL.
3. CMPAT=N specified on the PSBGEN statement.
4. CMPAT=Y specified on the PSBGEN statement.

Format of a PSB

The following is the format of a PSB.

```
[IOPCB]
[Alternate PCB ... Alternate PCB]
[DBPCB ... DBPCB]
[GSAMPCB ... GSAMPCB]
```

Each PSB must contain at least one PCB. The I/O PCB must be addressable in order to issue a system service command. An alternate PCB is used only for IMS online programs, which can run only with the Transaction Manager. Alternate PCBs can be present even though your program does not run under the Transaction Manager. A DB PCB can be a full-function PCB, a DEDB PCB, or an MSDB PCB.

Chapter 32. Recovering databases and maintaining database integrity

You can issue these commands to recover data accessed by your program and maintain data integrity.

- The Basic Checkpoint command, CHKP, which you can use to issue checkpoints from a batch or BMP program
- The Symbolic Checkpoint command, SYMCHKP, which you can use to issue checkpoints from a batch or BMP program and to specify data areas that can be restored when you restart your program
- The Extended Restart command, XRST, which you can use along with symbolic checkpoints to start or restart your batch or BMP program
- The rollback commands, ROLL and ROLB, which you can use to dynamically back out database changes from a batch or BMP program
- The managing-backout-points commands, SETS and ROLS, which you can use to set multiple backout points and then return to these points later
- The Dequeue command, DEQ, which releases previously reserved segments

To use any of the commands, you must have defined an I/O PCB for your program, except for the DEDB DEQ calls, which are issued against a DEDB PCB.

Issuing checkpoints in a batch or BMP program

The two kinds of commands that allow you to make checkpoints are: the CHKP, or Basic Checkpoint command, and the SYMCHKP, or Symbolic Checkpoint command.

Batch programs can use either the Symbolic Checkpoint or the Basic Checkpoint command.

Both checkpoint commands make it possible for you to commit your program's changes to the database and to establish places from which the batch or BMP program can be restarted, in cases of abnormal termination.

Requirement: You must not use the CHKPT=EOV parameter on any DD statement to take an IMS checkpoint.

Because both checkpoint commands cause a loss of database position at the time the command is issued, you must reestablish position with a GU command or other methods.

You cannot reestablish position in the midst of nonunique keys or nonkeyed segments.

Issuing the CHKP command

When you issue a CHKP command, you must provide the code for restarting your program and you must specify the ID for the checkpoint. You can supply either the name of a data area in your program that contains the ID, or you can supply the actual ID, enclosed in single quotation marks. For example, either of the following commands is valid:

```
EXEC DLI CHKP ID(chkpid);
```

```
EXEC DLI CHKP ID('CHKP0007');
```

Issuing the SYMCHKP command

The SYMCHKP command in batch and BMP programs:

- Works with the Extended Restart (XRST) command to restart your program if it terminates abnormally.

- Can save as many as seven program data areas, which are restored when your program is restarted. You can save variables, counters, and status information.

For examples of how to specify the SYMCHKP command, see the topic "SYMCHKP Command" in *IMS Version 15.3 Application Programming APIs*.

Restarting your program and checking for position

Programs that issue Symbolic Checkpoint commands must also issue the Extended Restart (XRST) command. You must issue XRST once, as the first command in the program. You can use the XRST command to start your program normally, or to restart it in case of an abnormal termination.

You can restart your program from one of the following:

- A specific checkpoint ID
- A time/date stamp

Because the XRST command attempts to reposition the database, your program also needs to check for correct position.

Backing out database updates dynamically: the ROLL and ROLB commands

When a batch program determines that some of its processing is invalid, the ROLL and ROLB commands make it possible for the program to remove the effects of its inaccurate processing.

You can use both ROLL and ROLB in batch programs. You can only use the ROLB command in batch programs if the system log is stored on direct access storage and if you have specified BKO=Y in the parm field of your JCL.

Issuing either of these commands causes DL/I to back out any changes your program has made to the database since its last checkpoint, or since the beginning of the program if your program has not issued a checkpoint.

Using intermediate backout points: the SETS and ROLS commands

Use the SETS and ROLS commands to define multiple points at which to preserve the state of DL/I full-function databases and to return to these points later. For example, you can use them to allow your program to handle situations that can occur when PSB scheduling complete without all of the referenced DL/I databases being available.

The SETS and ROLS commands apply only to DL/I full-function databases. Therefore, if a logical unit of work (LUW) is updating recoverable resources other than full-function databases (VSAM files, for example), the SETS and ROLS requests have no effect on the non-DL/I resources. The backout points are *not* CICS commit points; they are intermediate backout points that apply only to DBCTL resources. Your program must ensure the consistency of all the resources involved.

Before initiating a set of DL/I requests to perform a function, you can use a SETS command to define points in your application at which to preserve the state of DL/I databases. Your application can issue a ROLS command later if it cannot complete the function. You can use the ROLS command to back out to the state all full-function databases were in before either a specific SETS request or the most recent commit point.

Chapter 33. Processing Fast Path databases

Using EXEC DLI commands under DBCTL, a CICS program or a batch-oriented BMP program can access DEDBs. Parameters allow your program to use facilities of the DEDBs such as subset pointers.

A DEDB contains a root segment and as many as 127 types of dependent segment. One of these types can be a sequential dependent; the other 126 are direct dependents. Sequential dependent segments are stored in chronological order. Direct dependent segments are stored hierarchically.

DEDBs provide high data availability. Each DEDB can be partitioned, or divided into multiple *areas*. Each area contains a different set of database records. In addition, you can make up to seven copies of each area data set. If an error exists in one copy of an area, application programs can access the data by using another copy of that area. This is transparent to the application program. When an error occurs to data in a DEDB, IMS does not stop the database. It makes the data in error unavailable, but continues to schedule and process application programs. Programs that do not need the data in error are unaffected.

DEDBs can be shared among application programs in separate IMS systems. Sharing DEDBs is virtually the same as sharing full-function databases, and most of the same rules apply. IMS systems can share DEDBs at the area level (instead of at the database level as with full-function databases), or at the block level.

Processing Fast Path DEDBs with subset pointer options

Subset pointers and the options you use with them are optimization tools that significantly improve the efficiency of your program when you need to process long segment chains.

Subset pointers divide a chain of segment occurrences under the same parent into two or more groups, or subsets. You can define as many as eight subset pointers for any segment type. You then define the subset pointers from within an application program. Each subset pointer points to the start of a new subset. For example, in the following figure, suppose you defined one subset pointer that divided the last three segment occurrences from the first four. Your program can then refer to that subset pointer through options, and directly retrieve the last three segment occurrences.




Figure 88. Processing a long chain of segment occurrences with subset pointers

You can use subset pointers at any level of the database hierarchy, except at the root level. Subset pointers used for the root level are ignored.

The next two figures show some of the ways you can set subset pointers. Subset pointers are independent of one another, which means that you can set one or more pointers to any segment in

the chain. For example, you can set more than one subset pointer to a segment, as shown in the following figure.




Figure 89. Examples of setting multiple subset pointers

Alternatively, you can define a one-to-one relationship between the pointers and the segments, as shown in [Figure 90 on page 488](#) where each segment occurrence has one subset pointer.




Figure 90. More examples of setting subset pointers

The following figure shows how the use of subset pointers divides a chain of segment occurrences under the same parent into subsets. Each subset ends with the last segment in the entire chain. For example, the last segment in the subset defined by subset pointer 1 is B7.




Figure 91. How subset pointers divide a chain into subsets

Preparing to use subset pointers

For your program to use subset pointers, the pointers must be defined in the DBD for the DEDB, and in your program's PSB.

In the DBD, you specify the number of pointers for a segment chain. You can specify as many as eight pointers for any segment chain.

In the PSB, you specify which pointers your program uses; you define this on the SENSEG statement. (Each pointer is defined as an integer from 1 to 8.) You also specify on the SENSEG statement whether your program can set the pointers it uses. If your program has read-only sensitivity, it cannot set pointers, but can only retrieve segments using subset pointers already set. If your program has update sensitivity, it can update subset pointers by using the SET, SETCOND, MOVENEXT, and SETZERO options.

After the pointers are defined in the DBD and the PSB, an application program can set the pointers to segments in a chain. When an application program finishes executing, the subset pointers used by that program remain as they were set by the program and are not reset.

Designating subset pointers

To use subset pointers in your program, you must know the numbers for the pointers as they were defined in the PSB.

Then, when you use the subset pointer options, you specify the number for each subset pointer you want to use immediately after the option; for example, you would use P3 to indicate that you want to retrieve the first segment occurrence in the subset defined by subset pointer 3. No default exists, so if you do not include a number between 1 and 8, IMS considers your qualification statement invalid and returns an AJ status code to your program.

Subset pointer options

To take advantage of subsets, application programs use five different options.

The options are:

GETFIRST

Allows you to retrieve the first segment in a subset.

SETZERO

Sets a subset pointer to zero.

MOVENEXT

Sets a subset pointer to the segment following the current segment. Current position is at the current segment.

SET

Unconditionally sets a subset pointer to the current segment. Current position is at the current segment.

SETCOND

Conditionally sets a subset pointer to the current segment. Current position is at the current segment.

Banking transaction application example

The examples in this chapter are based on a sample application, the recording of banking transactions for a passbook account. The transactions are written to a database as either posted or unposted, depending on whether they were posted to the customer's passbook. For example, when Bob Emery does business with the bank, but forgets to bring in his passbook, an application program writes the transactions to the database as unposted. The application program sets a subset pointer to the first unposted transaction, so it can be easily accessed later. The next time Bob remembers to bring in his passbook, a program posts the transactions. The program can directly retrieve the first unposted transaction using the subset pointer that was previously set. After the program has posted the transactions, it sets the subset pointer to zero; an application program that subsequently updates the database can determine that no unposted transactions exist. The following figure summarizes the processing performed when the passbook is unavailable.




Figure 92. Processing performed for the sample passbook example when the passbook is unavailable

When the passbook is available, an application program adds the unposted transactions to the database, setting subset pointer 1 to the first unposted transaction. The following figure summarizes the processing performed when the passbook is available.




Figure 93. Processing performed for the sample passbook example when the passbook is available

When the passbook is available, an application program retrieves the first unposted transaction using the program, then posts all unposted transactions, setting subset pointer 1 to zero.

GETFIRST option: retrieving the first segment of a subset

To retrieve the first segment occurrence in the subset, your program issues a Get command with the GETFIRST option. The GETFIRST option does not set or move the pointer, but indicates to IMS that you want to establish position on the first segment occurrence in the subset. The GETFIRST option is like the FIRST option, except that the GETFIRST option applies to the subset instead of to the entire segment chain.

Using the previous example, imagine that Bob Emery visits the bank with his passbook and you want to post all of the unposted transactions. Because subset pointer 1 was previously set to the first unposted transaction, your program can use the following command to retrieve that transaction:

```
EXEC DLI GU SEGMENT(A) WHERE(AKEY = 'A1')
      SEGMENT(B) INTO(BAREA) GETFIRST('1');
```

As shown in following figure, this command retrieves segment B5. To continue processing segments in the chain, you can issue Get Next commands, as you would if you were not using subset pointers.




Figure 94. Retrieving the first segment in a chain of segments

If the subset does not exist (subset pointer 1 has been set to zero), IMS returns a GE status code, and your position in the database immediately follows the last segment in the chain. Using the passbook example, the GE status code indicates that no unposted transactions exist.

You can specify only one GETFIRST option per qualification statement; if you use more than one GETFIRST in a qualification statement, IMS returns an AJ status code to your program. The rules for using the GETFIRST option are:

1. You can use GETFIRST with all options except:
 - FIRST
 - LOCKCLASS
 - LOCKED
2. Other options take effect after the GETFIRST option has, and position has been established on the first segment in the subset.
3. If you use GETFIRST with LAST, the last segment in the segment chain is retrieved.
4. If the subset pointer specified with GETFIRST is not set, IMS returns a GE status code, not the last segment in the segment chain.
5. Do not use GETFIRST with FIRST. This causes you to receive an AJ status code.
6. GETFIRST overrides all insert rules, including LAST.

SETZERO, MOVENEXT, SET, and SETCOND options: setting the subset pointers

The SETZERO, MOVENEXT, SET, and SETCOND options allow you to redefine subsets by modifying the subset pointers. Before your program can set a subset pointer, it must establish a position in the database. A command must be fully satisfied before a subset pointer is set. The segment a pointer is set to depends on your current position at the completion of the command. If a command to retrieve a segment is not completely satisfied, and a position is not established, the subset pointers remain as they were before the command was issued.

• Setting the subset pointer to zero: SETZERO

The SETZERO option sets the value of the subset pointer to zero. After your program issues a command with the SETZERO option, the pointer is no longer set to a segment; the subset defined by that pointer no longer exists. (IMS returns a status code of GE to your program if you try to use a subset pointer having a value of zero.)

Using the previous example, say that you used the GETFIRST option to retrieve the first unposted transaction. You would then process the chain of segments, posting the transactions. After posting the transactions and inserting any new ones into the chain, you would use the SETZERO option to set the subset pointer to zero as shown in the following command:

```
EXEC DLI ISRT SEGMENT(A) WHERE(AKEY = 'A1')
      SEGMENT(B) FROM(BAREA) SETZERO('1');
```

After this command, subset pointer 1 would be set to zero, indicating to a program updating the database later on that no unposted transactions exist.

• Moving the subset pointer forward to the next segment after your current position: MOVENEXT

To move the subset pointer forward to the next segment after your current position, your program issues a command with the MOVENEXT option. Using the previous example, say that you wanted to post some of the transactions, but not all, and that you wanted the subset pointer to be set to the first unposted transaction. The following command sets subset pointer 1 to segment B6.

```
EXEC DLI GU SEGMENT(A) WHERE(AKEY = 'A1')
      SEGMENT(B) INTO(BAREA) GETFIRST('1') MOVENEXT('1');
```

The process of moving the subset pointer with this command is shown in the following figure. If the current segment is the last in the chain, and you use a MOVENEXT option, IMS sets the pointer to zero.




Figure 95. Moving the subset pointer to the next segment after your current position

- **Setting the subset pointer unconditionally: SET**

You use the SET option to set a subset pointer. The SET option sets a subset pointer unconditionally, regardless of whether or not it is already set. When your program issues a command that includes the SET option, IMS sets the pointer to your current position.

For example, to retrieve the first B segment occurrence in the subset defined by subset pointer 1, and to reset pointer 1 at the next B segment occurrence, you would issue the following commands:

```
EXEC DLI GU SEGMENT(A) WHERE(AKEY = 'A1')
      SEGMENT(B) INTO(BAREA) GETFIRST('1');
EXEC DLI GN SEGMENT(B) INTO(BAREA) SET('1');
```

After you have issued these commands, instead of pointing to segment B5, subset pointer 1 points to segment B6, as shown in the following figure.




Figure 96. Unconditionally setting the subset pointer to your current position

- **Setting the subset pointer conditionally: SETCOND**


Your program uses the SETCOND option to conditionally set the subset pointer. The SETCOND option is similar to the SET option; the only difference is that, with the SETCOND option, IMS updates the subset pointer only if the subset pointer is *not* already set to a segment.

Using the passbook example, say that Bob Emery visits the bank and forgets to bring his passbook; you add the unposted transactions to the database. You want to set the pointer to the first unposted transaction so that when you post the transactions later, you can immediately access the first one. The following command sets the subset pointer to the transaction you are inserting, if it is the first unposted one:

```
EXEC DLI ISRT SEGMENT(A) WHERE(AKEY = 'A1')
      SEGMENT(B) FROM(BAREA) SETCOND('1');
```

As shown by the following figure, this command sets subset pointer 1 to segment B5. If unposted transactions already existed, the subset pointer is not changed.

Before the command:



After the command:




Figure 97. Conditionally setting the subset pointer to your current position

Inserting segments in a subset

When you use the GETFIRST option to insert an unkeyed segment in a subset, the new segment is inserted before the first segment occurrence in the subset. However, the subset pointer is not automatically set to the new segment occurrence. For example, the following command inserts a new B segment occurrence in front of segment B5, but does not set subset pointer 1 to point to the new B segment occurrence:

```
EXEC DLI ISRT SEGMENT(A) WHERE(AKEY = 'A1')  
      SEGMENT(B) FROM(BAREA) GETFIRST('1');
```

To set subset pointer 1 to the new segment, you use the SET option along with the GETFIRST option, as shown in the following example:

```
EXEC DLI ISRT SEGMENT(A) WHERE(AKEY = 'A1')  
      SEGMENT(B) FROM(BAREA) GETFIRST('1') SET ('1');
```

If the subset does not exist (subset pointer 1 has been set to zero), the segment is added to the end of the segment chain.

Deleting the segment pointed to by a subset pointer

If you delete the segment pointed to by a subset pointer, the subset pointer points to the next segment occurrence in the chain. If the segment you delete is the last in the chain, the subset pointer is set to zero.

Combining options

You can use the SET, MOVENEXT, and SETCOND options with other options, and you can combine subset pointer options with each other, provided they do not conflict. For example, you can use GETFIRST and SET together, but you cannot use SET and SETZERO together because their functions conflict. If you combine options that conflict, IMS returns an AJ status code to your program.

You can use one GETFIRST option per qualification statement, and one update option (SETZERO, MOVENEXT, SET, or SETCOND) for each subset pointer.

Subset pointer status codes

If you make an error in a qualification statement that contains subset pointer options, IMS can return these status codes to your program.

AJ

The qualification statement used a GETFIRST, SET, SETZERO, SETCOND, or MOVENEXT option for a segment for which there are no subset pointers defined in the DBD.

The subset options included in the qualification statement are in conflict; for example, if one qualification statement contained a SET option and a SETZERO option for the same subset pointer, IMS would return an AJ status code. S means to set the pointer to current position; Z means to set the pointer to zero. You cannot use these options together in one qualification statement.

The qualification statement included more than one GETFIRST option.

The pointer number following a subset pointer option is invalid. You either did not include a number, or included an invalid character. The number following the option must be between 1 and 8, inclusive.

AM

The subset pointer referenced in the qualification statement was not specified in the program's PSB. For example, if your program's PSB specifies that your program can use subset pointers 1 and 4, and your qualification statement referenced subset pointer 5, IMS would return an AM status code to your program.

Your program tried to use an option that updates the pointer (SET, SETCOND, or MOVENEXT) but the program's PSB did not specify pointer update sensitivity.

Your program attempted to open a GSAM database without specifying an IOAREA.

The POS command

You can use the Position (POS) command (only with DEDBs) to perform the following functions.

- Retrieve the location of a specific sequential dependent segment, or retrieves the location of the last inserted sequential dependent segment.
- Tell you the amount of unused space within each DEDB area. For example, you can use the position information that IMS returns for a POS command to scan or delete the sequential dependent segments for a particular time period.

For the syntax of the POS command, see the topic "POS Command" in *IMS Version 15.3 Application Programming APIs*.

If the area the POS command specifies is unavailable, the I/O area is unchanged and the status code FH is returned.

Related reference

[POS command \(Application Programming APIs\)](#)

Locating a specific sequential dependent segment

When you have position on a particular root segment, you can retrieve the position information and the area name of a specific sequential dependent segment of that root.

If you have a position established on a sequential dependent segment, the search starts from that position. IMS returns the position information for the first sequential dependent segment that satisfies the command.

To retrieve this information, you issue a POS command with a qualification statement containing the segment name of the sequential dependent. The current position after this kind of POS command is in the same place that it is after a GNP command.

After a successful POS command, the I/O area contains:

LL

A 2-byte field giving the total length of the data in the I/O area, in binary.

Area Name

An 8-byte field giving the ddname from the AREA statement.

Position

An 8-byte field containing the position information for the requested segment.

If the sequential dependent segment that is the target of the POS command is inserted in the same synchronization interval, no position information is returned. Bytes 11-18 contain X'FF'; other fields contain normal data.

Unused CIs

A 4-byte field containing the number of unused CIs in the sequential dependent part.

Unused CIs

A 4-byte field containing the number of unused CIs in the independent overflow part.

Locating the last inserted sequential dependent segment

You can also retrieve the position information for the most recently inserted sequential dependent segment of a given root segment.

To do this, you issue a POS command with a qualification statement containing the root segment as the segment name. The current position after this type of command follows the same rules as position after a GU.

After a successful command, the I/O area contains:

LL

A 2-byte field containing the total length of the data in the I/O area, in binary.

Area Name

An 8-byte field giving the ddname from the AREA statement.

Position

An 8-byte field containing the position information for the most recently inserted sequential dependent segment. This field contains zeros provided no sequential dependent for this root exist.

Unused CIs

A 4-byte field containing the number of unused CIs in the sequential dependent part.

Unused CIs

A 4-byte field containing the number of unused CIs in the independent overflow part.

Identifying free space with the POS command

To retrieve the area name and the next available position within the sequential dependent part from all online areas, you can issue an unqualified POS command. This type of command also retrieves the free space in the independent overflow and sequential dependent parts.

After a successful unqualified POS command, the I/O area contains the length (LL) followed by the same number of entries as areas within the database. Each entry contains field two through five shown below:

LL

A 2-byte field containing the total length of the data in the I/O area, in binary. The length includes the 2 bytes for the LL field, plus 24 bytes for each entry.

Area Name

An 8-byte field giving the ddname from the AREA statement.

Position

An 8-byte field giving the next available position within the sequential dependent part.

Unused CIs

A 4-byte field containing the number of unused CIs in the sequential dependent part.

Unused CIs

A 4-byte field containing the number of unused CIs in the independent overflow part.

The P processing option

If the P processing option has been specified (with the PROCOPT parameter) in the PCB for your program, a GC status code is returned to your program whenever a command to retrieve or insert a segment causes a Unit of Work (UOW) boundary to be crossed.

Although crossing the UOW boundary probably has no particular significance for your program, the GC status code indicates that this is a good time to issue a CHKP command. The advantages of doing this are:

- Your position in the database is kept. Issuing a CHKP normally causes position in the database to be lost, and the application program has to reestablish position before it can resume processing.
- Commit points occur at regular intervals.

When a GC status code is returned, no data is retrieved or inserted. In your program, you can either:

- Issue a CHKP command, and resume database processing by reissuing the command that caused the GC status code.
- Ignore the GC status code and resume database processing by reissuing the command that caused the status code.

Chapter 34. Comparing command-level and call-level programs

Call-level and command-level programs exhibit different behavior.

DL/I calls for IMS and CICS

The following table provides a reference for using DL/I calls in a batch, batch-oriented BMP, or CICS with DBCTL environment.

Table 88. DL/I calls available to IMS and CICS command-level application programs

Request type	Batch	Batch-oriented BMP	CICS with DBCTL¹
CHKP call (symbolic)	Yes	Yes	No
CHKP call (basic)	Yes	Yes	No
GSCD call ²	Yes	No	No
INIT call	Yes	Yes	Yes
ISRT call (initial load)	Yes	No	No
ISRT call	Yes	Yes	Yes
LOG call	Yes	Yes	Yes
SCHD call	No	No	Yes
ROLB call	Yes	Yes	No
ROLL call	Yes	Yes	No
ROLS call (Roll Back to SETS) ³	Yes	Yes	Yes
ROLS call (Roll Back to Commit)	Yes	Yes	Yes
SETS call ³	Yes	Yes	Yes
STAT call ⁴	Yes	Yes	Yes
TERM call	No	No	Yes
XRST call	Yes	Yes	No

1. In a CICS remote DL/I environment, CALLs in the CICS-DBCTL column are supported if you are shipping a function to a remote CICS that uses DBCTL.
2. GSCD is a Product-sensitive Programming Interface.
3. SETS and ROLS calls are not valid when the PSB contains a DEDB.
4. STAT is a Product-sensitive Programming Interface.

Comparing EXEC DLI commands and DL/I calls

Use the appropriate EXEC DLI commands and DL/I calls in your program.

The following table compares EXEC DLI commands with DL/I calls. For example, in a command-level program, you use the LOAD command instead of the ISRT call to initially load a database.

Table 89. Comparing call-level and command-level programs: commands and calls

Call-level	Command-level	Purpose
INIT call	ACCEPT command	Initialize for data availability status codes.
CHKP call (basic)	CHKP command	Issue a basic checkpoint.
DEQ call	DEQ command	Release segments retrieved using LOCKCLASS option or Q command code.
DLET call	DLET command	Delete segments from a database.
GU, GN, and GNP calls	GU, GN, and GNP commands ¹	Retrieve segments from a database.
GHU, GHN, and GHNP calls ¹	GU, GN, and GNP commands ¹	Retrieve segments from a database for updating.
GSCD call	GSCD call ²	Retrieve system addresses.
ISRT call	ISRT command	Add segments to a database.
ISRT call	LOAD command	Initially load a database.
LOG call	LOG command	Write a message to the system log.
POS call	POS command	Retrieve positioning or space usage or positioning and space usage in a DEDB area.
INIT call	ACCEPT command	Initialize for data availability status.
INIT call	QUERY command	Obtain information of initial data availability.
INIT call	REFRESH command	Availability information after using a PCB.
REPL call	REPL command	Replace segments in a database.
XRST call	RETRIEVE command	Issue an extended restart.
ROLL or ROLB call	ROLL or ROLB command	Dynamically back out changes.
ROLS call	ROLS command	Back out to a previously set backout point.
PCB call	SCHD command	Schedule a PSB.
SETS call	SETS command	Set a backout point.
SETU call	SETU command	Set a backout point even if unsupported PCBs (like DEDBs or MSDBs) are present.
STAT call ³	STAT command	Obtain system and buffer pool statistics.
CHKP call (extended)	SYMCHKP command	Issue a symbolic checkpoint.
TERM call	TERM command	Terminate a PSB.
XRST call	XRST command	Issue an extended restart.

Notes:

1. Get commands are just like Get Hold calls, and the performance of Get commands and Get calls is the same.
2. You can use the GSCD call in a batch command-level program. GSCD is a Product-sensitive Programming Interface.
3. STAT is a Product-sensitive Programming Interface.

Comparing command codes and options

The following table compares the options you use with EXEC DLI commands with the command codes you use with DL/I calls. For example, the LOCKED option performs the same function as a Q command code.

Table 90. Comparing call-level and command-level programs: command codes and options

Call- Level	Command-Level	Allows You to . . .
C	KEYS option	Use the concatenated key of a segment to identify the segment.
D	INTO or FROM specified on segment level to be retrieved or inserted.	Retrieve or insert a sequence of segments in a hierarchic path using only one request, instead of having to use a separate request for each segment. (Path call or command).
F	FIRST option	Back up to the first occurrence of a segment under its parent when searching for a particular segment occurrence. Disregarded for a root segment.
L	LAST option	Retrieve the last occurrence of a segment under its parent.
M	MOVENEXT option	Set a subset pointer to the segment following the current segment.
N	Leave out the SEGMENT option for segments you do not want replaced.	Designate segments you do not want replaced, when replacing segments after a get hold request. Usually used when replacing a path of segments.
P	SETPARENT	Set parentage at a higher level than what it usually is (the lowest hierarchic level of the request).
Q	LOCKCLASS, LOCKED	Reserve a segment so that other programs are not able to update it until you have finished processing it.
R	GETFIRST option	Retrieve the first segment in a subset.
S	SET option	Unconditionally set a subset pointer to the current segment.
U	No equivalent for command level programs.	Limit the search for a segment to the dependents of the segment occurrence on which position is established.
V	CURRENT option	Use the hierarchic level of and levels above the current position as qualifications for the segment.
W	SETCOND option	Conditionally set a subset pointer to the current segment.
Z	SETZERO option	Set a subset pointer to zero.
–	No command-level equivalent.	Null. Use an SSA in command code format without specifying the command code. Can be replaced during execution with the command codes you want.

Chapter 35. Data availability enhancements

Your program might fail when it receives a status code indicating that a DL/I full-function database is unavailable. To avoid this, you can use these data availability enhancements. After a PSB has been scheduled in DBCTL, your application program can issue requests to indicate to IMS that the program can handle data availability status codes and to obtain information about the availability of each database.

Accepting database availability status codes

These status codes occur because PSB scheduling was completed without all of the referenced databases being available. Use the ACCEPT command to tell DBCTL to return a status code instead of abending the program:

```
EXEC DLI ACCEPT STATUSGROUP('A');
```

Obtaining information about database availability

You can put data availability status codes into each of the DB PCBs if:

- In a CICS DBCTL environment, by using the PSB scheduling request command, SCHD.
- In a Batch or BMP environment, at initialization time.

You can obtain the data availability status codes within the DL/I interface block (DIB) by using the following QUERY command:

```
EXEC DLI QUERY USING PCB(n);
```

n specifies the PCB.

The QUERY command is used after scheduling the PSB but before making the first database call. If the program has already issued a call using a DB PCB, then the QUERY command must follow the REFRESH command:

```
EXEC DLI REFRESH DBQUERY
```

The REFRESH command updates the information in the DIB. You can only issue this command one time.

For full-function databases, the DIBSTAT should contain NA, NU, TH, or blanks. For MSDBs and DEDBs, the DIBSTAT always contains blanks.

If a CICS command language translator has been used to translate the EXEC DLI commands, then, in addition to data availability status, the DBDNAME will be returned in the DIB field DIBBDNM. Also, the name of the database organization will be returned in the DIB field DIBDBORG.

Part 5. Application programming for SQL

These topics provide detailed information about developing IMS application programs in Structured Query Language (SQL).

These topics also provide detailed information about using SQL queries to retrieve and modify IMS data. This information is intended for IMS application developers who are familiar with SQL and who know one or more programming languages that IMS supports.

Chapter 36. SQL considerations and restrictions for COBOL

SQL support for COBOL has the several considerations and restrictions with the current implementation.

- A subset of SQL keywords is supported. There are SQL keywords that are currently supported by the IMS Universal JDBC driver but not supported when the SQL keyword is used in a COBOL application. For example:
 - XML is not supported by COBOL SQL in SELECT statements.
 - SQL COMMIT and ROLLBACK keywords are not supported. You should use IMS DB system services call to commit or roll back your database changes

For information about the supported SQL statement and keywords, see [SQL statements \(Application Programming APIs\)](#).

- SQL statements are supported for COBOL applications that are running in a BMP, IFP, and MPR for the IMS TM/DB environment and BMP is supported for the DBCTL environment. Batch and DB Batch are not supported. Also, the use of the EXEC SQLIMS API from IBM CICS Transaction Server for z/OS and Db2 for z/OS stored procedures to IMS is not supported. The DL/I API should be used instead of the EXEC SQLIMS API.
- You can access Db2 for z/OS data using DB2 SQL support in the same COBOL application. IMS must connect to DB2 by using the SSM parameter, and the DFSLI000 module must be included.
- The IMS catalog must be enabled to use SQL support for COBOL. Make sure that you load the database metadata that would be needed by the COBOL SQL application into the IMS catalog.
- Ensure that you have sufficient storage for your IMS dependent region in your COBOL SQL applications. Specify at least 12MB for your IMS dependent region size for running a COBOL SQL application. You would encounter an 878 or other storage related abends if you run out of storage.
- Only one cursor and SQL statement can be active at a time in the application. If you must execute multiple SQL statements in the same application, you must first close the cursor for the previous statement and then open a new cursor or prepare a new statement.
- The set of SQL keywords that are supported in COBOL is only for database access calls. For IMS database services, GSAM, IMS TM, and message processing services, continue to use DL/I API.
- Dynamic SQL statement is supported. Static SQL is not supported currently.
- Only EBCDIC CCSID 37 and 1140 codepages for the COBOL CODEPAGE option are supported.
- The qualifiers for type Zoned Decimal or Packed Decimal can only be combined with other qualifiers by using the logical operator OR in a query
- For the predicate in a WHERE clause, when the data type is integer, no more than 10 consecutive AND statements are allowed due to potential extensive memory consumption. You can have more than 10 statements if the WHERE clause also contains one or more OR predicates, but not 10 consecutive AND statements.
- When you write a WHERE clause for an SQL statement and specify virtual foreign keys in the qualification statements, the following limitations apply when access to the database is random:
 - The virtual foreign key must be for one of the first child segments of the root segment (level 2). You cannot specify virtual foreign keys for segments that are at lower levels in the hierarchy.
 - The OR operator must be used to combine multiple qualification statements that use the virtual foreign key.
 - Only virtual foreign keys can be specified. You cannot specify other columns from the same segment or another segment.

For example, assume that the HOSPITAL table is the root segment and the HOSPCODE column is the unique key. The WARD table is a first child segment of HOSPITAL and HOSPITAL_HOSPCODE is the

virtual foreign key that references the HOSPCODE column in the HOSPITAL table. The WHERE clause in the following SELECT statement is valid:

```
SELECT WARD.WARDNAME, WARD.PATCOUNT FROM PCB01.WARD
WHERE HOSPITAL_HOSPCODE='ARS100100D' OR HOSPITAL_HOSPCODE='ARS100100D'
```

The WHERE clause in the following SELECT statement will fail:

```
SELECT WARD.WARDNAME, WARD.PATCOUNT FROM PCB01.WARD
WHERE HOSPITAL_HOSPCODE='ARS100100D' AND XINTEGER > '5'
```

Chapter 37. Writing application programs for SQL

Applications that interact with IMS must first connect to IMS. They can then read, add, or modify data or manipulate IMS databases.

Coding SQL statements in application programs: General information

IMS provides a way to issue SQL statements directly in a COBOL application to access your IMS data. IMS uses the same SQL coding techniques as Db2 for z/OS when programming IMS SQL for COBOL.

To include a SQL statement in an application program:

1. Choose one of the following methods for communicating with IMS:

- Embedded dynamic SQL
- JDBC application support

If you are writing your applications in Java, you can use JDBC application support to access IMS.

If you are writing your applications for COBOL, use embedded dynamic SQL.

2. Define an SQL communications area (SQLIMSCA) that for your COBOL program can use to check whether an SQL statement executed successfully.
3. Define at least one SQL descriptor area (SQLIMSDA).
4. Declare any of the following data items for passing data between IMS and COBOL:
 - host variables
 - host structures

Ensure that you use the appropriate data types.

5. Code SQL statements to access IMS data.

If you are using a SELECT statement to query IMS data, use cursors to select a set of rows and then process one row at a time.

6. Check the execution of the SQL statements.
7. Handle any SQL error codes.

Related concepts

[“Dynamic SQL” on page 511](#)

Dynamic SQL statements are prepared and executed while the program is running. Use dynamic SQL when you do not know what SQL statements your application needs to execute before run time.

[“Programming with the IMS Universal JDBC driver” on page 594](#)

IMS provides a Java Database Connectivity (JDBC) driver for SQL-based database connectivity to access IMS databases over TCP/IP with the IMS Universal JDBC driver that is included in the IMS Universal drivers. The IMS Universal JDBC driver is based on the JDBC 4.0 standard.

Related tasks

[“Retrieving a set of rows by using a cursor” on page 545](#)

In an application program, you can retrieve a set of rows from IMS.

Defining the items that your program can use to check whether an SQL statement executed successfully

If your program contains SQL statements, the program should include an SQL communications area (SQLIMSCA), which contains SQLIMSCODE, SQLIMSSTATE, and SQLIMSERRMT variables, so that it can check whether the statements executed successfully.

Related tasks

[“Defining the SQL communications area in COBOL” on page 524](#)

COBOL programs that contain SQL statements can include an SQL communications area (SQLIMSCA) to check whether an SQL statement executed successfully.

Related reference

[Description of SQLIMSCA fields \(Application Programming APIs\)](#)

[INCLUDE \(Application Programming APIs\)](#)

Defining SQL descriptor areas

If your program includes certain SQL statements, you must include an SQL descriptor area (SQLIMSDA). Depending on the context in which it is used, the SQLIMSDA stores information about prepared SQL statements or host variables. This information can then be read by either the application program or IMS.

If your program includes any of the following statements, you must include an SQLIMSDA in your program:

- DESCRIBE *statement-name* INTO *descriptor-name*
- FETCH ... INTO DESCRIPTOR *descriptor-name*

Related tasks

[“Defining SQL descriptor areas in COBOL” on page 524](#)

If your program includes certain SQL statements such as DESCRIBE, you must define at least one SQL descriptor area (SQLIMSDA). Depending on the context in which it is used, the SQLIMSDA stores information about prepared SQL statements or host variables. This information can then be read by either the application program or IMS.

Related reference

[Description of SQLIMSCA fields \(Application Programming APIs\)](#)

[SQL descriptor area \(SQLIMSDA\) \(Application Programming APIs\)](#)

Declaring host variables and indicator variables

You can use host variables in SQL statements in your program to pass data between IMS and your application.

Use the techniques that are appropriate for COBOL.

Host variables

Use host variables to pass a single data item between IMS and your application.

A *host variable* is a single data item that is declared in the host language to be used within an SQL statement. You can use host variables in application programs that are written in COBOL to perform the following actions:

- Retrieve data into the host variable for your application program's use
- Use the data in the dynamic SQL statement host variables that have parameter marker during the EXECUTE, PREPARE, and OPEN calls

Related concepts

[“Rules for host variables in an SQL statement” on page 513](#)

Use host variables in embedded SQL statements to represent a single value. Host variables are useful for storing retrieved data or for passing values that are to be assigned or used for comparisons.

Related reference

[“Host variables in COBOL” on page 525](#)

In COBOL programs, you can specify numeric and character host variables.

Host structures

Use host structures to pass a group of host variables between IMS and your application.

A *host structure* is a group of host variables that can be referenced with a single name. You define host structures with statements in the host language. You can refer to a host structure in any context where you want to refer to the list of host variables in the structure. A host structure reference is equivalent to a reference to each of the host variables within the structure in the order in which they are defined in the structure declaration. You can also use indicator variables (or indicator structures) with host structures.

Related reference

[“Host structures in COBOL” on page 529](#)

A COBOL host structure is a named set of host variables that are defined in your program's WORKING-STORAGE SECTION or LINKAGE SECTION.

Indicator variables, arrays, and structures

An indicator variable is associated with a particular host variable. Each indicator variable contains a small integer value that indicates some information about the associated host variable. Indicator structures serve the same purpose for host variable structures.

You can use indicator variables to perform the following actions:

- Determine whether the value of an associated output host variable is null
- Determine the original length of a character string that was truncated when it was assigned to a host variable

You can use indicator structures to perform these same actions for individual items in host data structures.

If you provide an indicator variable for a nullable field of a variable-length segment, a negative value (-1) is set for the indicator variable when the field is null. Your program should check the indicator variable before using the field to determine whether the field is truly null. If the indicator variable contains a positive integer, the retrieved character string value is truncated, and the integer is the original length of the string.

An *indicator structure* is an array of halfword integer variables that supports a specified host structure. If the field values that your program retrieves into a host structure can be null, you can attach an indicator structure name to the host structure name. This name enables IMS to notify your program about each null value it returns to a host variable in the host structure.

Using SQL statements in your application

You can code SQL statements in a COBOL program using dynamic SQL.

Related concepts

[“SQL statements in COBOL programs” on page 533](#)

You can code SQL statements in certain COBOL program sections.

Dynamic SQL

Dynamic SQL statements are prepared and executed while the program is running. Use dynamic SQL when you do not know what SQL statements your application needs to execute before run time.

Dynamic SQL prepares and executes the SQL statements within a program, while the program is running. Two types of dynamic SQL are:

- Interactive SQL

A user enters SQL statements through IMS Explorer for Development. IMS prepares and executes those statements as dynamic SQL statements.

- Embedded dynamic SQL

Your application puts the SQL source in host variables and includes PREPARE and EXECUTE statements that tell IMS to prepare and run the contents of those host variables at run time. You must precompile and bind programs that include embedded dynamic SQL.

Dynamic SQL processing

A program that provides for dynamic SQL accepts as input, or generates, an SQL statement in the form of a character string. You can simplify the programming if you know a known number of values of known types. In the most general case, in which you do not know in advance about the SQL statements that will execute, the program typically takes these steps:

1. Translates the input data, including any parameter markers, into an SQL statement
2. Prepares the SQL statement to execute and acquires a description of the result segment
3. Obtains, for SELECT statements, enough main storage to contain retrieved data
4. Executes the statement or fetches the rows of data
5. Processes the information returned
6. Handles SQL return codes.

Dynamically executing SQL for fixed-list SELECT statements

A fixed-list SELECT statement returns rows that contain a known number of values of a known type. When you use this type of statement, you can specify a list of host variables to contain the fetched values.

The term "fixed-list" does not imply that you must know in advance how many rows of data will be returned. However, you must know the number of fields and the data types of those fields. A fixed-list SELECT statement returns a result segment that can contain any number of rows; your program looks at those rows one at a time, using the FETCH statement. Each successive fetch returns the same number of values as the last, and the values have the same data types each time.

To execute a fixed-list SELECT statement dynamically, your program must:

1. Include an SQLIMSCA.
2. Load the input SQL statement into a data area.
The preceding two steps are exactly the same including dynamic SQL for non-SELECT statements in your program.
3. Declare a cursor for the statement name.
4. Prepare the statement.
5. Open the cursor.
6. Fetch rows from the result segment.
7. Close the cursor.
8. Handle any resulting errors. This step is the same as for static SQL, except for the number and types of errors that can result.

Example: Suppose that your program retrieves hospital names and codes by dynamically executing SELECT statements of this form:

```
SELECT HOSPNAME, HOSPCODE FROM PCB01.HOSPITAL
```

Declaring a cursor for the statement name:

Use a cursor to put the results into host variables for a SELECT statement.

Example: When you declare the cursor, use the statement name (call it STMT), and give the cursor itself a name (for example, C1):

```
EXEC SQLIMS DECLARE C1 CURSOR FOR STMT  
END-EXEC.
```

Preparing the statement:

Prepare a statement (STMT) from STMTSTR.

Example: This is one possible PREPARE statement:

```
EXEC SQLIMS PREPARE STMT FROM :STMTSTR  
END-EXEC.
```

To execute STMT, your program must open the cursor, fetch rows from the result segment, and close the cursor.

Opening the cursor:

The OPEN statement evaluates the SELECT statement named STMT.

Example:

```
EXEC SQLIMS OPEN C1  
END-EXEC.
```

Fetching rows from the result table:

Example: Your program could repeatedly execute a statement such as this:

```
EXEC SQLIMS FETCH C1 INTO :HOSPNAM, :HOSPCODE  
END-EXEC.
```

The key feature of this statement is the use of a list of host variables to receive the values returned by FETCH. The list has a known number of items (in this case, two items, :HOSPNAM and :HOSPCODE) of known data types (both are character strings, of lengths 15 and 4, respectively).

You can use this list in the FETCH statement only because you planned the program to use only fixed-list SELECTs. Every row that cursor C1 points to must contain exactly two character values of appropriate length. If the program is to handle anything else, it must use the techniques for including dynamic SQL for varying-list SELECT statements in your program.

Closing the cursor:

Example: Close the cursor when your program is finished running the FETCH statement:

```
EXEC SQLIMS CLOSE C1  
END-EXEC.
```

Related concepts

[“SQL statements in COBOL programs” on page 533](#)

You can code SQL statements in certain COBOL program sections.

Related tasks

[“Dynamically executing SQL for non-SELECT statements” on page 518](#)

The easiest way to use dynamic SQL is to use non-SELECT statements such as the INSERT, UPDATE, or DELETE statement.

[“Dynamically executing SQL for varying-list SELECT statements” on page 514](#)

A varying-list SELECT statement returns rows that contain an unknown number of values of unknown type. When you use this type of statement, you do not know in advance exactly what kinds of host variables you need to declare for storing the results.

Rules for host variables in an SQL statement

Use host variables in embedded SQL statements to represent a single value. Host variables are useful for storing retrieved data or for passing values that are to be assigned or used for comparisons.

When you use host variables, adhere to the following requirements:

- You must declare the name of the host variable in the host program before you use it. Host variables follow the naming conventions of the host language.

- You can use a host variable to represent a data value, but you cannot use it to represent a segment, view, or field name. You can specify segment or field names at run time by using dynamic SQL.
- To use a host variable in an SQL statement, you can specify any valid host variable name that is declared according to the rules of the host language.
- A colon (:) must precede host variables that are used in SQL statements so that IMS can distinguish a variable name from a field name. When host variables are used outside of SQL statements, do not precede them with a colon.
- To optimize performance, make sure that the host language declaration maps as closely as possible to the data type of the associated data in the database.
- For assignments and comparisons between an IMS field and a host variable of a different data type or length, expect conversions to occur.

Related concepts

[“Dynamic SQL” on page 511](#)

Dynamic SQL statements are prepared and executed while the program is running. Use dynamic SQL when you do not know what SQL statements your application needs to execute before run time.

[Assignment and comparison \(Application Programming APIs\)](#)

Dynamically executing SQL for varying-list SELECT statements

A varying-list SELECT statement returns rows that contain an unknown number of values of unknown type. When you use this type of statement, you do not know in advance exactly what kinds of host variables you need to declare for storing the results.

What your application program must do for varying-list SELECT statements: To execute a varying-list SELECT statement dynamically, your program must follow these steps:

1. Include an SQLIMSCA.
2. Load the input SQL statement into a data area.
3. Prepare and execute the statement. This step is more complex than for fixed-list SELECTs. It involves the following steps:
 - a. Include an SQLIMSDA (SQL descriptor area).
 - b. Declare a cursor and prepare the statement.
 - c. Obtain information about each field of the result segment.
 - d. Determine the main storage needed to hold a row of retrieved data.
 - e. Put storage addresses in the SQLIMSDA to tell where each item of retrieved data should be stored.
 - f. Open the cursor.
 - g. Fetch a row.
 - h. Eventually close the cursor and free main storage.
4. Handle any errors that might result.

Preparing a varying-list SELECT statement:

Suppose that your program dynamically executes SQL statements.

Your program puts the statements into a varying-length character variable; call it STMTSTR. Your program goes on to prepare a statement from the variable and then give the statement a name; call it S1.

If the statement is a SELECT statement, the program must find out how many values are in each row, and what their data types are. The information comes from an SQL descriptor area (SQLIMSDA).

An SQL descriptor area:

The SQLIMSDA is a structure that is used to communicate with your program, and storage for it is usually allocated dynamically at run time.

For COBOL, use:

```
EXEC SQLIMS INCLUDE SQLIMSDA END-EXEC.
```

Obtaining information about the SQL statement:

An SQLIMSDA can contain a variable number of occurrences of SQLIMSVAR, each of which is a set of five fields that describe one field in the result segment of a SELECT statement. The included SQLIMSDA contains a maximum of 750 occurrences of SQLVARs, which means it can hold up to 750 of resulting columns.

Declaring a cursor for the statement:

As before, you need a cursor for the dynamic SELECT. For example, write:

```
EXEC SQLIMS  
  DECLARE C1 CURSOR FOR S1
```

Preparing the statement using the minimum SQLIMSDA:

To prepare a statement from the character string in STMTSTR and also enter its description into SQLIMSDA, write this:

```
EXEC SQLIMS PREPARE STMT FROM :STMTSTR END-EXEC.  
EXEC SQLIMS DESCRIBE STMT INTO :SQLIMSDA END-EXEC.
```

The following figure shows the contents of the minimum SQLIMSDA in use.




Figure 98. The minimum SQLIMSDA structure

SQLIMSN determines what SQLVAR gets:

The SQLIMSN field, which you must set before using DESCRIBE (or PREPARE INTO), tells how many occurrences of SQLIMSVAR the SQLIMSDA is allocated for.

- Base SQLIMSVAR information includes:

- Data type code
- Length attribute
- Column name or label
- Host variable address
- Indicator variable address

Whenever you execute DESCRIBE, IMS returns the following values, which you can use to build an SQLIMSDA of the correct size:

- SQLIMSD is 0 if the SQL statement is not a SELECT. Otherwise, SQLIMSD is the number of fields in the result segment.

If the statement is not a SELECT:

To find out if the statement is a SELECT, your program can query the SQLIMSD field in SQLIMSDA after the DESCRIBE statement. If the field contains 0, the statement is not a SELECT, the statement is already prepared, and your program can execute it. You can use:

```
EXEC SQLIMS EXECUTE STMT END-EXEC.
```

If the statement is a SELECT:

After the DESCRIBE statement executes, each occurrence of SQLIMSVAR contains a description of one field of the result segment in five fields.

The following table describes the values in the descriptor area.

<i>Table 91. Values inserted in the SQLIMSDA</i>		
Value	Field	Description
SQLIMSDA	SQLIMSDAID	An "eye-catcher"
750	SQLIMSN	The number of occurrences of SQLIMSVAR, set by the program
2	SQLIMSD	The number of occurrences of SQLIMSVAR actually used by the DESCRIBE statement
452	SQLIMSTYPE	The value of SQLIMSTYPE in the first occurrence of SQLIMSVAR. It indicates that the first field contains fixed-length character strings, and does not allow nulls.
3	SQLIMSLen	The length attribute of the column
Undefined	SQLIMSIND	
8	SQLIMSNAME	The number of characters in the field name
HOSPCODE	SQLIMSNAME+2	The field name of the first column

The following figure shows an SQLIMSDA that describes two fields.




Figure 99. Contents of SQLIMSDA after executing DESCRIBE

The first SQLIMSVAR pertains to the first field of the result segment (the HOSPCODE column). SQLIMSVAR element 1 contains fixed-length character strings and does not allow null values (SQLIMSTYPE=452); the length attribute is 3.

Acquiring storage to hold a row:

Before fetching rows of the result segment, your program must:

1. Analyze each SQLIMSVAR description to determine how much space you need for the field value.
2. Derive the address of some storage area of the required size.
3. Put this address in the SQLIMSDATA field.

If the SQLIMSTYPE field indicates that the value can be null, the program must also put the address of an indicator variable in the SQLIMSIND field. The following figures show the SQL descriptor area after you take certain actions.

Putting storage addresses in the SQLIMSDA:

After analyzing the description of each column, your program must replace the content of each SQLIMSDATA field with the address of a storage area large enough to hold values from that column. Similarly, for every field that allows nulls, the program must replace the content of the SQLIMSIND field. The content must be the address of a halfword that you can use as an indicator variable for the column. The program can acquire storage for this purpose, of course, but the storage areas used do not have to be contiguous.

The following figure shows the SQLIMSDA after your program acquires storage for the field values and their indicators, and puts the addresses in the SQLIMSDATA fields of the SQLIMSDA. It shows the content

of the descriptor area before the program obtains any rows of the result table. Addresses of fields and indicator variables are already in the SQLIMSVAR.




Figure 100. SQL descriptor area after analyzing descriptions and acquiring storage

Executing a varying-list SELECT statement dynamically:

You can easily retrieve rows of the result table using a varying-list SELECT statement. The statements differ only a little from those for the fixed-list example.

Open the cursor: If the SELECT statement contains no parameter marker, this step is simple enough. For example:

```
EXEC SQLIMS OPEN C1 END-EXEC.
```

Fetch rows from the result table: This statement differs from the corresponding one for the case of a fixed-list select. Write:

```
EXEC SQLIMS
  FETCH C1 USING DESCRIPTOR :SQLIMSDA END-EXEC.
```

The key feature of this statement is the clause **USING DESCRIPTOR :SQLIMSDA**. That clause names an SQL descriptor area in which the occurrences of SQLIMSVAR point to other areas. Those other areas receive the values that **FETCH** returns. It is possible to use that clause only because you previously set up SQLIMSDA to look like [Figure 99 on page 516](#).

The following figure shows the result of the **FETCH**. The data areas identified in the SQLIMSVAR fields receive the values from a single row of the result table.




Figure 101. SQL descriptor area after executing **FETCH**

Successive executions of the same **FETCH** statement put values from successive rows of the result table into these same areas.

Close the cursor: This step is the same as for the fixed-list case. When no more rows need to be processed, execute the following statement:

```
EXEC SQLIMS CLOSE C1 END-EXEC.
```

Related concepts

[“SQL statements in COBOL programs” on page 533](#)

You can code SQL statements in certain COBOL program sections.

Related reference

[DESCRIBE OUTPUT \(Application Programming APIs\)](#)

[SQL descriptor area \(SQLIMSDA\) \(Application Programming APIs\)](#)

[SQLIMSTYPE and SQLIMSLLEN \(Application Programming APIs\)](#)

[The SQLIMSDA header \(Application Programming APIs\)](#)

Dynamically executing SQL for non-SELECT statements

The easiest way to use dynamic SQL is to use non-SELECT statements such as the INSERT, UPDATE, or DELETE statement.

Your program must take the following steps:

1. Include an SQLIMSCA.
2. Load the input SQL statement into a data area.
3. Execute the statement.
 - PREPARE and EXECUTE
4. Handle any errors that might result. The return code from the most recently executed SQL statement appears in the host variables SQLIMSCODE and SQLIMSSTATE or corresponding fields of the SQLIMSCA.
5. Check SQLIMSERRD(3) for the number of rows being changed.

Example:

Suppose that your program updates the hospital name by dynamically executing the UPDATE statement of this form:

```
UPDATE HOSPITAL SET HOSPNAME = 'MISSION CREEK'
WHERE HOSPITAL.HOSPCODE = 'H001007'
```

In this example, the UPDATE statement is stored in a host variable STMTSTR.

Declare the statement:

```
EXEC SQLIMS
  DELCARE STMT STATEMENT
END-EXEC.
```

Declaring a varying-length character host variable: Before you prepare and execute an SQL statement, you have to assign it into a host variable. Declare varying-length character host variable for the SQL statement. The first two bytes must contain the length of the SQL statement. The maximum length of the SQL statement is 32K. For example:

```
01 STMTSTR.
   49 STMTSTR-LEN PIC S9(4) COMP VALUE +180.
   49 STMTSTR-TXT PIC X(180) VALUE SPACES.
```

Preparing the statement:

Prepare a statement (STMT) from the STMTSTR host variable.

```
EXEC SQLIMS
  PREPARE STMT FROM :STMTSTR
END-EXEC.
```

Executing the statement:

To execute STMT, your program use the EXECUTE call.

```
EXEC SQLIMS  
EXECUTE STMT  
END-EXEC.
```

Related concepts

[“SQL statements in COBOL programs” on page 533](#)

You can code SQL statements in certain COBOL program sections.

Related tasks

[“Checking the execution of SQL statements” on page 521](#)

After executing an SQL statement, your program should check for any error codes before you commit the data and handle the errors that they represent.

Dynamically executing a SELECT SQL statement with parameter markers

Use the SELECT statements with parameter markers.

Suppose that you want to execute SELECT statements repeatedly using a list of hospital numbers. Suppose further that users enter a list of hospital numbers to be retrieved. You must construct and execute the entire statement dynamically. Your program can:

- Use parameter markers instead of constant values in the SQL statement
- Use OPEN statement with the USING clause to set values for parameter markers
- Use FETCH to retrieve data

Statements with parameter markers:

Dynamic SQL statements cannot use host variables. Therefore, you cannot dynamically execute an SQL statement that contains host variables. Instead, use parameter marker. A parameter marker is a question mark (?) that represents a position in a dynamic SQL statement where the application will provide a value.

Example using parameter markers:

```
SELECT HOSPNAME FROM PCB01.HOSPITAL WHERE HOSPCODE = ?;
```

You associate the host variable HOSPCODE with the parameter marker when you fetch data with the prepared statement.

Declaring a varying-length character host variable: Before you prepare and execute an SQL statement, you have to assign it into a host variable. Declare varying-length character host variable for the SQL statement. The first two bytes must contain the length of the SQL statement. The maximum length of the SQL statement is 32K. For example:

```
01 STMTSTR.  
49 STMTSTR-LEN PIC S9(4) COMP VALUE +180.  
49 STMTSTR-TXT PIC X(180) VALUE SPACES.
```

Declaring the CURSOR:

Declare a cursor to put the results for a SELECT statement.

When you declare the cursor, use the statement name (call it S1), and give the cursor itself a name (for example, C1):

```
EXEC SQLIMS DECLARE C1 CURSOR FOR S1  
END-EXEC.
```

Using the PREPARE statement:

Assume that the character host variable :STMTSTR has the value `SELECT HOSPNAME FROM PCB01.HOSPITAL WHERE HOSPCODE = ?`. To prepare an SQL statement from that string and assign it the name S1, write:

```
EXEC SQLIMS PREPARE S1 FROM :STMTSTR;
```

The prepared statement still contains a parameter marker, for which you must supply a value when the statement executes. After the statement is prepared, the parameter marker enables you to execute the same statement many times with different values of the hospital code.

To execute STMT, your program must open the cursor, fetch rows from the result segment, and close the cursor.

Using the OPEN statement:

The OPEN statement open a cursor for the prepared SQL statement. If the SQL statement contains parameter markers, you must use the USING clause of OPEN to provide values for all of the parameter markers. The USING clause of the OPEN statement names a list of one or more host variables or a host structure. This list supplies values for all of the parameter markers. Suppose that C1 is the cursor and the parameter value is contained in the host variable HOSPCODE, write:

```
OPEN C1 USING :HOSPCODE
```

The OPEN statement can be executed using different values for HOSPCODE.

Using more than one parameter marker: The prepared statement (S1 in the example) can contain more than one parameter marker. If it does, the USING clause of EXECUTE specifies a list of variables or a host structure. The variables must contain values that match the number and data types of parameters in S1 in the proper order. You must know the number and types of parameters in advance and declare the variables in your program.

For example, `OPEN C1 USING :PARM1, :PARM2`

Fetching rows from the result table:

This example shows you how to fetch data into host variables:

```
EXEC SQLIMS FETCH C1 INTO :HOSPNAME, :HOSPCODE  
END-EXEC.
```

Closing the cursor:

Close the cursor when your program is finished running the FETCH statement:

```
EXEC SQLIMS CLOSE C1  
END-EXEC.
```

Related concepts

[“SQL statements in COBOL programs” on page 533](#)

You can code SQL statements in certain COBOL program sections.

Related reference

[PREPARE \(Application Programming APIs\)](#)

Dynamically executing a non-select SQL statement with parameter markers

Use PREPARE and EXECUTE for non-SELECT statements like INSERT, UPDATE, and DELETE with parameter markers.

Suppose that you want to execute UPDATE statements repeatedly using a list of hospital numbers. Suppose further that users enter a list of hospital numbers to update. You must construct and execute the entire statement dynamically. Your program must now do these things differently:

- Use parameter markers instead of host variables

- Use the PREPARE and EXECUTE statement

Parameter markers with PREPARE and EXECUTE: Dynamic SQL statements cannot use host variables. Therefore, you cannot dynamically execute an SQL statement that contains host variables. Instead, use parameter marker. A parameter marker is a question mark (?) that represents a position in a dynamic SQL statement where the application will provide a value.

Example using parameter markers:

```
DELETE FROM PCB01.HOSPITAL WHERE HOSPCODE = ?;
```

You associate the host variable HOSPCODE with the parameter marker when you execute the prepared statement.

Declaring a varying-length character host variable: Before you prepare and execute an SQL statement, you have to assign it into a host variable. Declare varying-length character host variable for the SQL statement. The first two bytes must contain the length of the SQL statement. The maximum length of the SQL statement is 32K. For example:

```
01 STMTSTR.
49 STMTSTR-LEN PIC S9(4) COMP VALUE +180.
49 STMTSTR-TXT PIC X(180) VALUE SPACES.
```

Using the PREPARE statement: Assume that the character host variable :STMTSTR has the value DELETE FROM PCB01.HOSPITAL WHERE HOSPCODE = ?. To prepare an SQL statement from that string and assign it the name S1, write:

```
EXEC SQLIMS PREPARE S1 FROM :STMTSTR;
```

The prepared statement still contains a parameter marker, for which you must supply a value when the statement executes. After the statement is prepared, the parameter marker enables you to execute the same statement many times with different values of the hospital code.

Using the EXECUTE statement: The EXECUTE statement executes a prepared SQL statement by naming a list of one or more host variables or a host structure. This list supplies values for all of the parameter markers. Suppose that S1 is the prepared statement and the parameter value is contained in the host variable HOSPCODE, write:

```
EXECUTE S1 USING :HOSPCODE
```

The EXECUTE statement can be executed using different values for HOSPCODE.

Using more than one parameter marker: The prepared statement (S1 in the example) can contain more than one parameter marker. If it does, the USING clause of EXECUTE specifies a list of variables or a host structure. The variables must contain values that match the number and data types of parameters in S1 in the proper order. You must know the number and types of parameters in advance and declare the variables in your program.

Example:

If two parameter markers are in STMT, you need the following statement:

```
EXEC SQLIMS
EXECUTE STMT USING :PARM1, :PARM2
END-EXEC.
```

Checking the execution of SQL statements

After executing an SQL statement, your program should check for any error codes before you commit the data and handle the errors that they represent.

You can check the execution of SQL statements in one of the following ways:

- By displaying specific fields in the SQLIMSCA.

- By testing SQLIMSCODE or SQLIMSSTATE for specific values.
- By using the WHENEVER statement in your application program.
- By testing indicator variables to detect numeric errors.

Related tasks

[“Defining the SQL communications area in COBOL” on page 524](#)

COBOL programs that contain SQL statements can include an SQL communications area (SQLIMSCA) to check whether an SQL statement executed successfully.

Checking the execution of SQL statements by using the SQLIMSCA

One way to check whether an SQL statement executed successfully is to use the SQL communication area (SQLIMSCA). This area is set apart for communication with IMS.

If you use the SQLIMSCA, include the necessary instructions to display information that is contained in the SQLIMSCA in your application program.

- When IMS processes an SQL statement, it places return codes that indicate the success or failure of the statement execution in SQLIMSCODE and SQLIMSSTATE.
- When IMS processes a FETCH statement, and the FETCH is successful, the contents of SQLIMSERRD(3) in the SQLIMSCA is set to the number of returned rows.
- When IMS processes a FETCH statement, the contents of SQLIMSCODE is set to +100 if the last row in the segment has been returned with the set of rows.
- When IMS processes an UPDATE, INSERT, or DELETE statement, and the statement execution is successful, the contents of SQLIMSERRD(3) in the SQLIMSCA is set to the number of rows that are updated, inserted, or deleted.

Related tasks

[“Checking the execution of SQL statements by using SQLIMSCODE and SQLIMSSTATE” on page 522](#)

Whenever an SQL statement executes, the SQLIMSCODE and SQLIMSSTATE fields of the SQLIMSCA receive a return code. The SQLIMSERRMT field of the SQLIMSCA will contain message text that describes the error.

[“Defining the SQL communications area in COBOL” on page 524](#)

COBOL programs that contain SQL statements can include an SQL communications area (SQLIMSCA) to check whether an SQL statement executed successfully.

Related reference

[Description of SQLIMSCA fields \(Application Programming APIs\)](#)

Checking the execution of SQL statements by using SQLIMSCODE and SQLIMSSTATE

Whenever an SQL statement executes, the SQLIMSCODE and SQLIMSSTATE fields of the SQLIMSCA receive a return code. The SQLIMSERRMT field of the SQLIMSCA will contain message text that describes the error.

SQLIMSCODE:

IMS returns the following codes in SQLIMSCODE:

- If SQLIMSCODE = 0, execution was successful.
- If SQLIMSCODE > 0, execution was successful with a warning.
- If SQLIMSCODE < 0, execution was not successful.

SQLIMSCODE 100 indicates that no data was found.

SQLIMSSTATE: SQLIMSSTATE enables an application program to check for errors in the same way for different IBM database management systems.

Related tasks

[“Defining the SQL communications area in COBOL” on page 524](#)

COBOL programs that contain SQL statements can include an SQL communications area (SQLIMSCA) to check whether an SQL statement executed successfully.

Related reference

[SQL codes \(Messages and Codes\)](#)

Checking the execution of SQL statements by using the WHENEVER statement

The WHENEVER statement causes IMS to check the SQLIMSCA and continue processing your program. If an error, exception, or warning occurs, IMS branches to another area in your program. The condition handling area of your program can then examine the SQLIMSCODE or SQLIMSSTATE to react specifically to the error or exception.

The WHENEVER statement enables you to specify what to do if a general condition is true. You can specify more than one WHENEVER statement in your program. When you do this, the first WHENEVER statement applies to all subsequent SQL statements in the source program until the next WHENEVER statement.

The WHENEVER statement looks like this:

```
EXEC SQLIMS
  WHENEVER condition action
END-EXEC
```

The *condition* of the WHENEVER statement is one of these three values:

SQLWARNING

Indicates what to do when SQLIMSWARN0 = W or SQLIMSCODE contains a positive value other than 100. IMS can set SQLIMSWARN0 for several reasons—for example, if a field value is truncated when moved into a host variable. Your program might not regard this as an error.

SQLERROR

Indicates what to do when IMS returns an error code as the result of an SQL statement (SQLIMSCODE < 0).

NOT FOUND

Indicates what to do when IMS cannot find a row to satisfy your SQL statement or when there are no more rows to fetch (SQLIMSCODE = 100).

The *action* of the WHENEVER statement is one of these two values:

CONTINUE

Specifies the next sequential statement of the source program.

GOTO or GO TO *host-label*

Specifies the statement identified by *host-label*. For *host-label*, substitute a single token, preceded by an optional colon. The form of the token depends on the host language. In COBOL, for example, it can be *section-name* or an unqualified *paragraph-name*.

The WHENEVER statement must precede the first SQL statement it is to affect. However, if your program checks SQLIMSCODE directly, you must check SQLIMSCODE after each SQL statement.

Related reference

[WHENEVER \(Application Programming APIs\)](#)

Coding SQL statements in COBOL application programs

When you code SQL statements in COBOL application programs, you should follow certain guidelines.

Defining the SQL communications area in COBOL

COBOL programs that contain SQL statements can include an SQL communications area (SQLIMSCA) to check whether an SQL statement executed successfully.

Use the following SQL INCLUDE statement to request a standard SQLIMSCA declaration:

```
EXEC SQLIMS INCLUDE SQLIMSCA
```

You can specify INCLUDE SQLIMSCA or a declaration for SQLIMSCODE wherever you can specify a 77 level or a record description entry in the WORKING-STORAGE SECTION.

IMS sets the SQLIMSCODE, SQLIMSSTATE, and SQLIMSERRMT values in the SQLIMSCA after each SQL statement executes. Your application should check these values to determine whether the last SQL statement was successful.

Related tasks

[“Checking the execution of SQL statements” on page 521](#)

After executing an SQL statement, your program should check for any error codes before you commit the data and handle the errors that they represent.

[“Checking the execution of SQL statements by using the SQLIMSCA” on page 522](#)

One way to check whether an SQL statement executed successfully is to use the SQL communication area (SQLIMSCA). This area is set apart for communication with IMS.

[“Checking the execution of SQL statements by using SQLIMSCODE and SQLIMSSTATE” on page 522](#)

Whenever an SQL statement executes, the SQLIMSCODE and SQLIMSSTATE fields of the SQLIMSCA receive a return code. The SQLIMSERRMT field of the SQLIMSCA will contain message text that describes the error.

[“Defining the items that your program can use to check whether an SQL statement executed successfully” on page 509](#)

If your program contains SQL statements, the program should include an SQL communications area (SQLIMSCA), which contains SQLIMSCODE, SQLIMSSTATE, and SQLIMSERRMT variables, so that it can check whether the statements executed successfully.

Defining SQL descriptor areas in COBOL

If your program includes certain SQL statements such as DESCRIBE, you must define at least one SQL descriptor area (SQLIMSDA). Depending on the context in which it is used, the SQLIMSDA stores information about prepared SQL statements or host variables. This information can then be read by either the application program or IMS.

To define SQL descriptor areas:

1. Use the following SQL INCLUDE statement to request a standard SQLIMSDA declaration:

```
EXEC SQLIMS INCLUDE SQLIMSDA
```

2. You must place SQLIMSDA declarations in the WORKING-STORAGE SECTION, LINKAGE SECTION or LOCAL-STORAGE SECTION of your program, wherever you can specify a record description entry in that section. You must place SQLIMSDA declarations before the first SQL statement that references the data descriptor.

Related tasks

[“Defining SQL descriptor areas” on page 510](#)

If your program includes certain SQL statements, you must include an SQL descriptor area (SQLIMSDA). Depending on the context in which it is used, the SQLIMSDA stores information about prepared SQL statements or host variables. This information can then be read by either the application program or IMS.

Declaring host variables and indicator variables in COBOL

You can use host variables and host structures in SQL statements in your program to pass data between IMS and your application.

To declare host variables and host structures:

1. Declare the variables according to the following rules and guidelines:
 - You must explicitly declare all host variables that are used in SQL statements in the WORKING-STORAGE SECTION or LINKAGE SECTION of your program's DATA DIVISION.
 - You must explicitly declare each host variable before using them in an SQL statement.
 - You cannot implicitly declare any host variables through default typing or by using the IMPLICIT statement.
 - Ensure that any SQL statement that uses a host variable is within the scope of the statement that declares that variable.
2. Optional: Define any associated indicator variables, arrays, and structures.

Related tasks

[“Declaring host variables and indicator variables” on page 510](#)

You can use host variables in SQL statements in your program to pass data between IMS and your application.

Host variables in COBOL

In COBOL programs, you can specify numeric and character host variables.

Restrictions:

- Only some of the valid COBOL declarations are valid host variable declarations. If the declaration for a variable is not valid, any SQL statement that references the variable might result in the message UNDECLARED HOST VARIABLE.
- One or more REDEFINES entries can follow any level 77 data description entry. However, you cannot use the names in these entries in SQL statements. Entries with the name FILLER are ignored.

Recommendations:


- Be careful of overflow. For example, suppose that you retrieve an INTEGER field value into a PICTURE S9(4) host variable and the field value is larger than 32767 or smaller than -32768. You get an overflow warning or an error, depending on whether you specify an indicator variable.
- Be careful of truncation. For example, if you retrieve an 80-character CHAR field value into a PICTURE X(70) host variable, the rightmost 10 characters of the retrieved string are truncated. Retrieving a double precision floating-point or decimal field value into a PIC S9(8) COMP host variable removes any fractional part of the value. Similarly, retrieving a field value with DECIMAL data type into a COBOL decimal variable with a lower precision might truncate the value.

Numeric host variables

You can specify the following forms of numeric host variables:

- Floating-point numbers
- Integers and small integers
- Decimal numbers


The following diagram shows the syntax for declaring floating-point or real host variables.



Notes:

- ¹ level-1 indicates a COBOL level between 2 and 48.
- ² COMPUTATIONAL-1 and COMP-1 are equivalent.
- ³ COMPUTATIONAL-2 and COMP-2 are equivalent.


The following diagram shows the syntax for declaring integer, small integer, and big integer host variables.



Notes:

- ¹ level-1 indicates a COBOL level between 2 and 48.
- ² The COBOL binary integer data types BINARY, COMPUTATIONAL, COMP, COMPUTATIONAL-4, and COMP-4 are equivalent.
- ³ COMPUTATIONAL-5 (and COMP-5) are equivalent to the other COBOL binary integer data types if you compile the other data types with TRUNC(BIN).
- ⁴ Any specification for scale is ignored.

The following diagram shows the syntax for declaring decimal host variables.



Notes:

¹ *level-1* indicates a COBOL level between 2 and 48.

² The *picture-string* that is associated with **SIGN LEADING SEPARATE** must have the form $S9(i)V9(d)$ (or $S9...9V9...9$, with i and d instances of 9 or $S9...9V$ with i instances of 9).

³ **PACKED-DECIMAL**, **COMPUTATIONAL-3**, and **COMP-3** are equivalent. The *picture-string* that is associated with these types must have the form $S9(i)V9(d)$ (or $S9...9V9...9$, with i and d instances of 9) or $S9(i)V$.

In COBOL, you declare the **SMALLINT** and **INTEGER** data types as a number of decimal digits. IMS uses the full size of the integers (in a way that is similar to processing with the **TRUNC(BIN)** compiler option) and can place larger values in the host variable than would be allowed in the specified number of digits in the COBOL declaration. If you compile with **TRUNC(OPT)** or **TRUNC(STD)**, ensure that the size of numbers in your application is within the declared number of digits.

For small integers that can exceed 9999, use $S9(4)$ **COMP-5** or compile with **TRUNC(BIN)**. For large integers that can exceed 999 999 999, use $S9(10)$ **COMP-3** to obtain the decimal data type. If you use COBOL for integers that exceed the COBOL **PICTURE**, specify the field as decimal to ensure that the data types match and perform well.

If you are using a COBOL compiler that does not support decimal numbers of more than 18 digits, use one of the following data types to hold values of greater than 18 digits:


- A decimal variable with a precision less than or equal to 18, if the actual data values fit. If you retrieve a decimal value into a decimal variable with a scale that is less than the source field in the database, the fractional part of the value might be truncated.
- An integer or a floating-point variable, which converts the value. If you use an integer variable, you lose the fractional part of the number. If the decimal number might exceed the maximum value for an integer or if you want to preserve a fractional value, use a floating-point variable. Floating-point numbers are approximations of real numbers. Therefore, when you assign a decimal number to a floating-point variable, the result might be different from the original number.

Character host variables

You can specify the following forms of character host variables:

- Fixed-length strings

The following diagram shows the syntax for declaring fixed-length character host variables.




Notes:

¹ *level-1* indicates a COBOL level between 2 and 48.


² The *picture-string* that is associated with these forms must be $X(m)$ (or $XX...X$, with m instances of X), where m is up to COBOL's limitation.

The following diagrams show the syntax for declaring varying-length character host variables.



Notes:

¹ *level-1* indicates a COBOL level between 2 and 48.




Notes:

¹ You cannot use an intervening REDEFINE at level 49.

² You cannot directly reference *var-1* as a host variable.

³ IMS uses the full length of the S9(4) BINARY variable even though COBOL with TRUNC(STD) recognizes values up to only 9999. This behavior can cause data truncation errors when COBOL statements execute and might effectively limit the maximum length of variable-length character strings to 9999. Consider using the TRUNC(BIN) compiler option or USAGE COMP-5 to avoid data truncation.



Notes:

- ¹ You cannot use an intervening REDEFINE at level 49.
- ² You cannot directly reference *var-2* as a host variable.
- ³ For fixed-length strings, the *picture-string* must be *X(m)* (or *XX*, with *m* instances of *X*), where *mis* up to COBOL's limitation.

Related concepts

[“Host variables” on page 510](#)

Use host variables to pass a single data item between IMS and your application.

Related tasks

[“Using SQL statements in your application” on page 511](#)

You can code SQL statements in a COBOL program using dynamic SQL.

Host structures in COBOL

A COBOL host structure is a named set of host variables that are defined in your program's WORKING-STORAGE SECTION or LINKAGE SECTION.

Requirements: Host structure declarations in COBOL must satisfy the following requirements:


- COBOL host structures can have a maximum of two levels, even though the host structure might occur within a structure with multiple levels.
- A host structure name can be a group name whose subordinate levels name elementary data items.
- If you are using the IMS coprocessor, do not declare host variables or host structures on any subordinate levels after one of the following items:
 - A COBOL item that begins in area A
 - Any SQL statement (except SQL INCLUDE)
 - Any SQL statement within an included member

When the IMS precompiler encounters one of the preceding items in a host structure, it considers the structure to be complete.


When you write an SQL statement that contains a qualified host variable name (perhaps to identify a field within a structure), use the name of the structure followed by a period and the name of the field. For example, for structure B that contains field C1, specify *B.C1* rather than *C1 OF B* or *C1 IN B*.

Host structures

The following diagram shows the syntax for declaring host structures.



A




Notes:

- ¹ *level-1* indicates a COBOL level between 1 and 47.
- ² *level-2* indicates a COBOL level between 2 and 48.
- ³ For elements within a structure, use any level 02 through 48 (rather than 01 or 77), up to a maximum of two levels.
- ⁴ Using a FILLER or optional FILLER item within a host structure declaration can invalidate the whole structure.


Numeric usage items

The following diagram shows the syntax for numeric-usage items that are used within declarations of host structures.




Integer and decimal usage items

The following diagram shows the syntax for integer and decimal usage items that are used within declarations of host structures.



CHAR inner variables

The following diagram shows the syntax for CHAR inner variables that are used within declarations of host structures.



Related concepts

[“Host structures” on page 511](#)

Use host structures to pass a group of host variables between IMS and your application.

Equivalent SQL and COBOL data types

When you declare host variables in your COBOL programs, the precompiler uses equivalent SQL data types. When you retrieve data of a particular SQL data type into a host variable, you need to ensure that the host variable is of an equivalent data type.

The following table describes the SQL data type and the base SQLIMSTYPE and SQLIMSLLEN values that the precompiler uses for host variables in SQL statements.

Table 92. SQL data types, SQLIMSLLEN values, and SQLIMSTYPE values that the precompiler uses for host variables in COBOL programs

COBOL host variable data type	SQLIMSTYPE of host variable ¹	SQLIMSLLEN of host variable	SQL data type
COMP-2	480	8	DOUBLE PRECISION, or FLOAT(<i>n</i>) 22≤ <i>n</i> ≤53
S9(<i>i</i>)V9(<i>d</i>) COMP-3 or S9(<i>i</i>)V9(<i>d</i>) PACKED-DECIMAL	484	<i>i</i> + <i>d</i> in byte 1, <i>d</i> in byte 2	DECIMAL(<i>i</i> + <i>d</i> , <i>d</i>) or NUMERIC(<i>i</i> + <i>d</i> , <i>d</i>)
S9(<i>i</i>)V9(<i>d</i>) DISPLAY SIGN LEADING SEPARATE	504	<i>i</i> + <i>d</i> in byte 1, <i>d</i> in byte 2	No exact equivalent. Use DECIMAL(<i>i</i> + <i>d</i> , <i>d</i>) or NUMERIC(<i>i</i> + <i>d</i> , <i>d</i>)
S9(4) COMP-4, S9(4) COMP-5, S9(4) COMP, or S9(4) BINARY	500	2	SMALLINT
S9(9) COMP-4, S9(9) COMP-5, S9(9) COMP, or S9(9) BINARY	496	4	INTEGER
S9(18) COMP-4, S9(18) COMP-5, S9(18) COMP, or S9(18) BINARY	492	8	BIGINT
Fixed-length character data	452	<i>n</i>	CHAR(<i>n</i>)
SQL TYPE is BINARY(<i>n</i>), 1≤ <i>n</i> ≤255	912	<i>n</i>	BINARY(<i>n</i>)

Notes:

1. If a host variable includes an indicator variable, the SQLIMSTYPE value is the base SQLIMSTYPE value plus 1.

The following table shows equivalent COBOL host variables for each SQL data type. Use this table to determine the COBOL data type for host variables that you define to receive output from the database. For example, if you retrieve TIMESTAMP data, you can define a fixed-length character string variable of length *n*

This table shows direct conversions between SQL data types and COBOL data types. However, a number of SQL data types are compatible. When you do assignments or comparisons of data that have compatible data types, IMS converts those compatible data types.

Table 93. COBOL host variable equivalents that you can use when retrieving data of a particular SQL data type

SQL data type	COBOL host variable equivalent	Notes
SMALLINT	S9(4) COMP-4, S9(4) COMP-5, S9(4) COMP, or S9(4) BINARY	
INTEGER	S9(9) COMP-4, S9(9) COMP-5, S9(9) COMP, or S9(9) BINARY	
DECIMAL(<i>p</i> , <i>s</i>)	S9(<i>p</i> - <i>s</i>)V9(<i>s</i>) COMP-3 or S9(<i>p</i> - <i>s</i>)V9(<i>s</i>) PACKED-DECIMAL DISPLAY SIGN LEADING SEPARATE	<i>p</i> is precision; <i>s</i> is scale. 0≤ <i>s</i> ≤ <i>p</i> ≤31. If <i>s</i> =0, use S9(<i>p</i>)V or S9(<i>p</i>). If <i>s</i> = <i>p</i> , use SV9(<i>s</i>). If the COBOL compiler does not support 31-digit decimal numbers, no exact equivalent exists. Use COMP-2.

Table 93. COBOL host variable equivalents that you can use when retrieving data of a particular SQL data type (continued)

SQL data type	COBOL host variable equivalent	Notes
DOUBLE PRECISION, DOUBLE or FLOAT (<i>n</i>)	COMP-2	22 ≤ <i>n</i> ≤ 53
BIGINT	S9(18) COMP-4, S9(18) COMP-5, S9(18) COMP, or S9(18) BINARY	
CHAR(<i>n</i>)	Fixed-length character string. For example, 01 VAR-NAME PIC X(<i>n</i>).	1 ≤ <i>n</i> ≤ 255
BINARY(<i>n</i>)	SQL TYPE IS BINARY(<i>n</i>)	1 ≤ <i>n</i> ≤ 255
DATE	Fixed-length character string of length <i>n</i> . For example, 01 VAR-NAME PIC X(<i>n</i>).	
TIME	Fixed-length character string of length <i>n</i> . For example, 01 VAR-NAME PIC X(<i>n</i>).	
TIMESTAMP	Fixed-length character string of length <i>n</i> . For example, 01 VAR-NAME PIC X(<i>n</i>).	<i>n</i> must be at least 19. To include microseconds, <i>n</i> must be 26; if <i>n</i> is less than 26, truncation occurs on the microseconds part.

SQL statements in COBOL programs

You can code SQL statements in certain COBOL program sections.

The allowable sections are shown in the following table.

Table 94. Allowable SQL statements for COBOL program sections

SQL statement	Program section
INCLUDE SQLIMSCA	WORKING-STORAGE SECTION ¹ or LINKAGE SECTION
INCLUDE text-file-name	PROCEDURE DIVISION or DATA DIVISION ²
DECLARE CURSOR	DATA DIVISION or PROCEDURE DIVISION
Other	PROCEDURE DIVISION

Notes:

1. If you use the IMS coprocessor, you can use the LOCAL-STORAGE SECTION wherever WORKING-STORAGE SECTION is listed in the table.
2. When including host variable declarations, the INCLUDE statement must be in the WORKING-STORAGE SECTION or the LINKAGE SECTION.

You cannot put SQL statements in the DECLARATIVES section of a COBOL program.

Each SQL statement in a COBOL program must begin with EXEC SQLIMS and end with END-EXEC. The EXEC and SQLIMS keywords can be on different lines. Do not include any tokens between the two keywords EXEC and SQL except for COBOL comments, including debugging lines. Do not include SQL comments between the keywords EXEC and SQLIMS.

If the SQL statement appears between two COBOL statements, the period after END-EXEC is optional and might not be appropriate. If the statement appears in an IF...THEN set of COBOL statements, omit the ending period to avoid inadvertently ending the IF statement.

Comments: You can include COBOL comment lines (* in field 7) in SQL statements wherever you can use a blank. You cannot include COBOL comment lines between the keywords EXEC and SQLIMS. The IMS coprocessor treats the debugging lines based on the COBOL rules, which depend on the WITH DEBUGGING mode setting.

For an SQL INCLUDE statement, the IMS coprocessor treats this text as part of the COBOL program syntax.

In addition, you can include SQL comments ('--') in any embedded SQL statement.

Debugging lines: The IMS coprocessor follows the COBOL language rules regarding debugging lines.

Continuation for SQL statements: The rules for continuing a character string constant from one line to the next in an SQL statement embedded in a COBOL program are the same as those for continuing a non-numeric literal in COBOL. However, you can use either a quote or an apostrophe as the first nonblank character in area B of the continuation line. The same rule applies for the continuation of delimited identifiers and does not depend on the string delimiter option.

To conform with SQL standard, delimit a character string constant with an apostrophe, and use a quote as the first nonblank character in area B of the continuation line for a character string constant.

Continued lines of an SQL statement can be in fields 12 through 72 when using the IMS coprocessor.

Declaring segments: Your COBOL program should include the statement DECLARE TABLE to describe each segment and view the program accesses. You can use the IMS declarations generator to generate the DECLARE TABLE statements. You should include the generated members in the DATA DIVISION.

Dynamic SQL in a COBOL program: In general, COBOL programs can easily handle dynamic SQL statements. COBOL programs can handle SELECT statements if the data types and the number of fields returned are fixed. If you want to use variable-list SELECT statements, use an SQLIMSDA.

Including code: To include SQL statements or COBOL host variable declarations from a member of a partitioned data set, use the following SQL statement in the source code where you want to include the statements:

```
EXEC SQLIMS INCLUDE member-name END-EXEC.
```

Use the 'EXEC SQLIMS' and 'END-EXEC' keyword pair to include SQL statements only. COBOL statements, such as COPY or REPLACE, are not allowed.

Margins: You must code SQL statements that begin with EXEC SQLIMS in fields 12 through 72.

Names: You can use any valid COBOL name for a host variable. Do not use entry names that begin with DFS or DQF, and do not use host variable names that begin with 'SQL' or 'SQLIMS'. These names are reserved for IMS.

Sequence numbers: The source statements that the IMS coprocessor generates do not include sequence numbers.

Statement labels: You can precede executable SQL statements in the PROCEDURE DIVISION with a paragraph name, if you wish.

WHENEVER statement: The target for the GOTO clause in an SQL statement WHENEVER must be a section name or unqualified paragraph name in the PROCEDURE DIVISION.

Special COBOL considerations: The following considerations apply to programs written in COBOL:

- In a COBOL program that uses elements in a multi-level structure as host variable names, the IMS coprocessor generates the lowest two-level names.
- To avoid truncating numeric values, use either of the following methods:
 - Use the COMP-5 data type for binary integer host variables.
 - Specify the COBOL compiler option:
 - TRUNC(OPT) if you are certain that the data being moved to each binary variable by the application does not have a larger precision than is defined in the PICTURE clause of the binary variable.
 - TRUNC(BIN) if the precision of data being moved to each binary variable might exceed the value in the PICTURE clause.

IMS assigns values to binary integer host variables as if you had specified the COBOL compiler option TRUNC(BIN) or used the COMP-5 data type.

- Do not use COBOL figurative constants (such as ZERO and SPACE), symbolic characters, reference modification, and subscripts within SQL statements.
- Observe the rules for naming SQL identifiers. However, for COBOL only, the names of SQL identifiers can follow the rules for naming COBOL words, if the names do not exceed the allowable length for the IMS object. For example, the name 1ST-TIME is a valid cursor name because it is a valid COBOL word, but the name 1_TIME is not valid because it is not a valid SQL identifier or a valid COBOL word.
- Observe these rules for hyphens:
 - Surround hyphens used as subtraction operators with spaces. IMS usually interprets a hyphen with no spaces around it as part of a host variable name.
- If you include an SQL statement in a COBOL PERFORM ... THRU paragraph and also specify the SQL statement WHENEVER ... GO, the COBOL compiler returns the warning message IGYOP3094. That message might indicate a problem. This usage is not recommended.

Related concepts

[SQL identifiers \(Application Programming APIs\)](#)

Related tasks

[“Using SQL statements in your application” on page 511](#)

You can code SQL statements in a COBOL program using dynamic SQL.

[“Defining SQL descriptor areas” on page 510](#)

If your program includes certain SQL statements, you must include an SQL descriptor area (SQLIMSDA). Depending on the context in which it is used, the SQLIMSDA stores information about prepared SQL statements or host variables. This information can then be read by either the application program or IMS.

Delimiters in SQL statements in COBOL programs

You must delimit SQL statements in your COBOL program so that IMS knows when a particular SQL statement ends.

Delimit an SQL statement in your COBOL program with the beginning keyword EXEC SQLIMS and an END-EXEC.

Example

Use EXEC SQLIMS and END-EXEC. to delimit an SQL statement in a COBOL program:

```
EXEC SQLIMS
  an SQL statement
END-EXEC.
```

SQL aggregate functions supported for COBOL

SQL aggregate functions are supported for COBOL and .NET applications (by using IMS Enterprise Suite IMS Data Provider for Microsoft .NET).

- AVG
- COUNT
- MAX
- MIN
- SUM

These functions work with the ORDER BY clause, as well as GROUP BY:

- ORDER BY
 - ASC
 - DESC
- GROUP BY

Restriction: The supported SQL aggregate functions accept only a single field name in a segment as the argument (the DISTINCT keyword is not allowed).

The following table shows the data types of the fields that are accepted by the aggregate functions, along with the resulting data type in the ResultSet.

Table 95. Supported SQL aggregate functions and their supported data types		
Function	Argument types	Results and result types
AVG	All supported numeric data types for SQL statements in COBOL programs are supported, including TINYINT, SMALLINT, INTEGER, BIGINT, Zoned Decimal, and Packed Decimal	• The function is applied to the set of values derived from the argument values by excluding null values. • For INTEGER, the fractional part of the average is discarded. The result can be null. • For non-DECIMAL data types, the result data type is always LONG. • If the data type of the argument value is DECIMAL, the result is packed decimal. The scale of the result is the same as the scale of the argument value, and the precision of the result is 31. • If the function is applied to an empty set, the result is the null value. • The averaged value must be within the range of the data type of the result.
COUNT	Any supported data type	• NULL values are not counted when the total number of values in a given column is counted (COUNT(column)). • NULL values are counted when the number of rows in a table is counted (COUNT(*)). • COUNT(*) of an empty table returns one row with a value of 0.

Table 95. Supported SQL aggregate functions and their supported data types (continued)

Function	Argument types	Results and result types
MAX	TINYINT, INTEGER, BIGINT, Zoned Decimal, Packed Decimal, CHAR, BINARY, DATE, TIME, and TIMESTAMP	- Character string arguments and binary string arguments cannot have a length attribute greater than 32704. - The data type of the result and its other attributes (for example, the length and CCSID of a string or a datetime value) are the same as the data type and attributes of the argument values. - The result can be null. - The function is applied to the set of values derived from the argument values by the elimination of null values. - If the function is applied to an empty set, the result is the null value.
MIN	TINYINT, INTEGER, BIGINT, Zoned Decimal, Packed Decimal, CHAR, BINARY, DATE, TIME, and TIMESTAMP	- The data type of the result and its other attributes (for example, the length and CCSID of a string or a datetime value) are the same as the data type and attributes of the argument values. - The result can be null. - The function is applied to the set of values derived from the argument values by the elimination of null values. - If the function is applied to an empty set, the result is the null value.
SUM	All supported numeric data types for SQL statements in COBOL programs are supported, including TINYINT, SMALLINT, INTEGER, BIGINT, Zoned Decimal and Packed Decimal	- The sum must be within the range of the data type of the result. - The function is applied to the set of values from the argument values by eliminating null values. - If the function is applied to an empty set, the result is the null value. - The order in which the summation is performed is undefined, but every intermediate result must be within the range of the result data type. - If the data type of the argument value is DECIMAL, the result is packed decimal. The scale of the result is the same as the scale of the argument value, and the precision of the result is 31.

Column names generated by aggregate functions

The generated column name from an aggregate function is a combination of the aggregate function name and the field name, separated by an underscore character (_). For example, the statement `SELECT MAX(age)` results in a column name `MAX_age`.

If the aggregate function argument field is table-qualified, the generated column name is the combination of the aggregate function name, the table name, and the column name, separated by underscore characters (_). For example, `SELECT MAX(Employee.age)` results in a column name `MAX_Employee_age`.

The aggregate function is executed first, and then the required number of rows of result are fetched from the result set.

Using the ORDER BY and GROUP BY clauses

The field names that are specified in a GROUP BY or ORDER BY clause must match exactly the field name that is specified in the SELECT statement. In order to GROUP BY properly, the fields specified in SELECT list must also be specified in the GROUP BY list.

```
SELECT HOSPNAME, COUNT(PATNAME) FROM PCB01.HOSPITAL, PATIENT GROUP BY HOSPNAME
ORDER BY HOSPNAME
```

Restriction: Aggregate functions cannot be used in the GROUP BY or ORDER BY statements. For example, GROUP BY COUNT(PATNAME) or ORDER BY AVG(COST) is not supported.

Related reference

[SELECT \(Application Programming APIs\)](#)

Adding and modifying data

Your application program can query, modify, or delete data in any IMS segment for which you have the appropriate level of access.

Inserting rows

You can insert data into segments using the SQL INSERT statement.

Use an INSERT statement to add new rows to a segment or view. Using an INSERT statement, you can specify the field values to insert a single row. You can specify constants or parameter markers, by using the VALUES clause.

For every row that you insert, you must provide a value for every key field. If you do not specify a value in the INSERT call, IMS sets a value of 0.

Inserting a single row:

You can use the VALUES clause of the INSERT statement to insert a single row of field values into a segment. You can either name all of the fields for which you are providing values, or you can omit the list of field names. If you omit the field name list, you must specify values for **all** of the fields. The fields are ordered first by their field position within the IMS catalog and then by their length.

When inserting a record in a table at a non-root level, you must specify values for all the foreign key fields of the table. Foreign key fields properly position the new record (or segment instance) to be inserted in the hierarchic path using standard SQL processing, similar to foreign keys in a relational database.

Recommendation: For INSERT statements, name all of the fields for which you are providing values for the following reasons:

- Your INSERT statement is independent of the segment format. (For example, you do not need to change the statement when a field is added to the segment.)
- You can verify that you are specifying the values in order.
- Your source statements are more self-descriptive.

When you list the field names, you must specify their corresponding values in the same order as in the list of field names.

Example: The following statement inserts information about a new hospital into the HOSPITAL segment.

```
INSERT INTO PCB01.HOSPITAL (HOSPCODE, HOSPNAME)
VALUES ('R1210050000A', 'MALLEY CLINIC')
```

After inserting a new hospital into your HOSPITAL segment, you can use a SELECT statement to see what you have loaded into the segment. The following SQL statement shows you all of the new department rows that you have inserted:

```
SELECT HOSPCODE, HOSPNAME
FROM PCB01.HOSPITAL
```

The result segment looks similar to the following output:

HOSPCODE	HOSPNAME
R1210010000A	ALEXANDRIA
R1210020000A	SANTA TERESA
R1210030000A	SANTA CLARA
R1210040000A	NEW ENGLAND
R1210050000A	MALLEY CLINIC

Example: The following statement inserts a new WARD record under a specific HOSPITAL table. In this example, the WARD table has the foreign key HOSPITAL_HOSPCODE. The new record will be inserted if and only if there is a HOSPCODE in the HOSPITAL table with the value of 'R1210050000A'

```
INSERT INTO PCB01.WARD
WARDNO, HOSPITAL_HOSPCODE, WARDNAME)
VALUES ('0001', 'R1210050000A', 'EMGY')
```

Example: The following statement also inserts a row into the HOSPITAL segment without specifying the column names. All the columns values must be specified in the VALUES clause.

```
INSERT INTO PCB01.HOSPITAL
VALUES (900, 'R1210050000A', 'MALLEY CLINIC');
```

Updating segment data

You can change a field value to another value or remove the field value altogether.

To change the data in a segment, use the UPDATE statement. You can also use the UPDATE statement to remove a value from a field (without removing the row) by changing the field value to an empty string.

Example: The following statement updates the hospital name to 'MISSION CREEK' for hospital 'H001007'.

```
UPDATE HOSPITAL SET HOSPNAME = 'MISSION CREEK'
WHERE HOSPITAL.HOSPCODE = 'H001007'
```

The SET clause names the fields that you want to update and provides the values that you want to assign to those fields. You can replace a field value in the SET clause with any of the following items:

- An expression, which can be any of the following items:
 - A constant
 - A parameter marker

Next, identify the rows to update:

- To update a single row, use a WHERE clause that locates one, and only one, row.
- To update several rows, use a WHERE clause that locates only the rows that you want to update.

If you omit the WHERE clause, IMS updates every row in the segment with the values that you supply. If IMS cannot find the row you want to identify, a SQLIMSCODE of 100 will be returned to the application.

If the UPDATE is successful, SQLIMSERRD(3) in the SQLIMSCA contains the number of updated rows. This number includes only the number of updated rows in the segment that is specified in the UPDATE statement.

Update rules:

Update values must satisfy the following rules. If they do not, or if other errors occur during the execution of the UPDATE statement, no rows are updated and the position of the cursors are not changed.

- Assignment. Update values are assigned to columns using the assignment rules described in Language elements.
- When updating a record in a table at a non-root level, you must specify values for all the foreign key fields of the table to identify the exact record (or segment instance) to update.
- Making an UPDATE on a foreign key field is invalid.

If IMS finds an error while executing your UPDATE statement (for example, an update value that is too large for the field), it returns an error. Upon receiving the error, the application will have to decide on how to manage the rows that were already changed. It can either commit or rollback the changes.

The following statement updates a WARD record under a specific HOSPITAL. In this example, the WARD table has the virtual foreign key HOSPITAL_HOSPCODE. The record will be updated if and only if there is a HOSPCODE in the HOSPITAL table with the value of 'H5140070000H'.

```
UPDATE WARD SET WARDNAME = 'EMGY',  
DOCCOUNT = '2', NURCOUNT = '4'  
WHERE HOSPITAL_HOSPCODE = 'H5140070000H'  
AND WARDNO = '01'
```

Deleting data from segments

You can delete data from a segment by deleting one or more rows from the segment, by deleting all rows from the segment, or by dropping fields from the segment.

To delete one or more rows in a segment:

- Use the DELETE statement with a WHERE clause to specify a search condition.

The DELETE statement removes zero or more rows of a segment, depending on how many rows satisfy the search condition that you specify in the WHERE clause.

The following DELETE statement deletes each row in the HOSPITAL segment that has Hospital named ALEXANDRIA and SANTA TERESA.

```
DELETE FROM PCB01.HOSPITAL WHERE HOSPNAME = 'ALEXANDRIA' OR HOSPNAME = 'SANTA TERESA';
```

When this statement executes, IMS deletes any row from the HOSPITAL segment that meets the search condition.

If IMS finds an error while executing your DELETE statement, the application commits and rolls back the changes and returns error codes in the SQLIMSCODE and SQLIMSSTATE variables in the SQLIMSCA. (The data in the segment does not change.)

If the DELETE is successful, SQLIMSERRD(3) in the SQLIMSCA contains the number of deleted rows. This number includes only the number of deleted rows in the segment that is specified in the DELETE statement.

To delete every row in a segment:

- Use the DELETE statement without specifying a WHERE clause.

The following DELETE statement deletes every row in the HOSPITAL segment:

```
DELETE FROM HOSPITAL;
```


If the statement executes, the segment continues to exist (that is, you can insert rows into it), but it is empty.

Related concepts

[SQL communication area \(SQLIMSCA\) \(Application Programming APIs\)](#)

Accessing data

Your program can read data from any IMS segments for which you have read access using SQL SELECT statements.

Related concepts

[Writing SQL queries to access an IMS database with the IMS Universal JDBC driver \(Application Programming\)](#)

Retrieving data by using the SELECT statement

The simplest way to retrieve data is to use the SQL SELECT statement to specify a result segment. You can specify the fields and rows that you want to retrieve.

You do not need to know the field names to select IMS data. Use an asterisk (*) in the SELECT clause to indicate that you want to retrieve all fields of each selected row of the named segment. To view the values of these fields, you must specify the field name.

The fields in a SELECT * statement are ordered first by their field position within the IMS catalog and then by their length.

Example: SELECT *: The following statement retrieves all fields for the PATIENT segment:

```
SELECT *  
FROM PCB01.HOSPITAL;
```

The result segment looks similar to the following output:

HOSPCODE	HOSPNAME
R1210010000A	ALEXANDRIA
R1210020000A	SANTA TERESA
R1210030000A	SANTA CLARA
R1210040000A	NEW ENGLAND

Because the example does not specify a WHERE clause, the statement retrieves data from all rows.

SELECT * is not recommended when fetching into a static host structure because of host variable compatibility and performance reasons.

Note:

- Suppose that you add a field to the segment to which SELECT * refers. If you have not defined a receiving host variable for that field, an error occurs.
- If you list the field names in an SELECT statement instead of using an asterisk, you can avoid the problem that sometimes occurs with SELECT *. You can also see the relationship between the receiving host variables and the fields in the result segment.

Selecting some fields: SELECT field-name:

Select the field or fields you want to retrieve by naming each field. All fields appear in the order you specify, not in their order in the segment.

Example: SELECT field-name: The following statement retrieves the ward names and patient names from the WARD and PATIENT tables, respectively:

```
SELECT HOSPNAME FROM PCB01.HOSPITAL
```

The result segment looks similar to the following output:

```
+-----+
|HOSPNAME|
+-----+
|ALEXANDRIA|
|SANTA TERESA|
|SANTA CLARA|
|NEW ENGLAND|
+-----+
```

With a single SELECT statement, you can select data from one field or as many as 750 fields.

Selecting rows using search conditions: WHERE:

Use a WHERE clause to select the rows that meet certain conditions. A WHERE clause specifies a search condition. A *search condition* consists of one or more predicates. A *predicate* specifies a test that you want IMS to apply to each segment row.

IMS evaluates a predicate for each row as true, false, or unknown. Results are unknown only if an operand is null.

The following segment lists the type of comparison, the comparison operators, and an example of each type of comparison that you can use in a predicate in a WHERE clause.

Table 96. Comparison operators used in conditions

Type of comparison	Comparison operator	Example
Equal to	=	HOSPCODE = 'R1210010000A'
Not equal to	<>	HOSPCODE <> 'R1210020000A'
Less than	<	SALARY < 30000
Less than or equal to	<=	AGE <= 25
Not less than	>=	AGE >= 21
Greater than	>	WARDNO > '0001'
Greater than or equal to	>=	WARDNO >= '0003'
Not greater than	<=	PATNUM <= '0010'
At least one of two conditions	OR	HOSPCODE >= 'R1210010000A' OR HOSPCODE < 'R1210050000A'
Both of two conditions	AND	HOSPCODE = 'R1210050000A' AND HOSPNAME = 'SANTA TERESA'

Both of these forms of the predicate create an expression for which one value is equal to another value or both values are equal to null.

Related concepts

[“Host variables” on page 510](#)

Use host variables to pass a single data item between IMS and your application.

[Predicates \(Application Programming APIs\)](#)

Formatting the result segment

An SQL statement returns data in a segment called a result segment. You can specify certain attributes of the result segment, such as the field names, how the rows are ordered, and whether the rows are numbered.

Result segments

The data that is retrieved by an SQL statement is always in the form of a segment, which is called a *result segment*. Like the segments from which you retrieve the data, a result segment has rows and fields. A program fetches this data one row at a time.

Example result segment: Assume that you issue the following SELECT statement, which retrieves the hospital code and name from the HOSPITAL segment and order the result with the name of the hospital in ascending order:

```
SELECT HOSPCODE, HOSPNAME
FROM PCB01.HOSPITAL
ORDER BY HOSPNAME
```

The result segment looks similar to the following output:

HOSPCODE	HOSPNAME
R1210010000A	ALEXANDRIA
R1210040000A	NEW ENGLAND
R1210030000A	SANTA CLARA
R1210020000A	SANTA TERESA

Ordering the result segment rows

If you want to guarantee that the rows in your result segment are ordered in a particular way, you must specify the order in the SELECT statement. Otherwise, IMS can return the rows in any order.

To retrieve rows in a specific order, use the ORDER BY clause. Using ORDER BY is the only way to guarantee that your rows are ordered as you want them. The following topics show you how to use the ORDER BY clause.

Specifying the sort key in the ORDER BY clause:

The order of the selected rows depends on the sort keys that you identify in the ORDER BY clause. A *sort key* is a field name of the segment. IMS orders the rows by the first sort key, followed by the second sort key, and so on.

You can list the rows in ascending or descending order. Null values appear last in an ascending sort and first in a descending sort.

Example: ORDER BY clause with a field name as the sort key: The following statement retrieves all hospital names, sorted in alphabetical order:

```
SELECT HOSPITAL.HOSPNAME
FROM PCB01.HOSPITAL
ORDER BY HOSPITAL.HOSPNAME ASC
```

Optimizing retrieval for a small set of rows

When you need only a few of the thousands of rows that satisfy a query, you can tell IMS to return only a specified number of rows.

Question: How can I tell IMS that I want only a few of the thousands of rows that satisfy a query?

Answer: Use or FETCH FIRST *n* ROWS ONLY.

If you want to retrieve only the first few rows. For example, to retrieve the first 50 rows, code:

```
SELECT * FROM PCB01.HOSPITAL
FETCH FIRST 50 ROWS ONLY
```

Use `FETCH FIRST n ROWS ONLY` to limit the number of rows in the result segment to *n* rows. `FETCH FIRST n ROWS ONLY` has the following benefits:

- When you use `FETCH` statements to retrieve data from a result segment, `FETCH FIRST n ROWS ONLY` causes IMS to retrieve only the number of rows that you need. This can have performance benefits, especially in distributed applications. If you try to execute a `FETCH` statement to retrieve the *n*+1st row, IMS returns a +100 SQLCODE.

Implications of using `SELECT *`

Generally, you should use `SELECT *` only when you want to select all fields. Otherwise, specify the specific fields that you want to view.

Question: What are the implications of using `SELECT *` ?

Answer: Generally, you should select only the fields you need because IMS is sensitive to the number of fields selected. Use `SELECT *` only when you are sure you want to select all fields, except hidden fields.

Support for variable-length database segments

SQL language conventions assume that the target database is relational. Relational database managers do not use the concept of "variable-length" data structures that are managed by an external application program.

Because IMS is a hierarchical database, IMS translates SQL statements into DL/I calls that can be interpreted by the IMS Database Manager. When using the IMS Database Manager, applications must manage the length of individual variable-length segment instances with the LL field for the segment. Applications that use the IMS Universal DL/I driver or SQL support for COBOL are responsible for managing the LL field.

For application programs that use the IMS Universal database resource adapter or the IMS Universal JDBC driver driver treat IMS databases as standard JDBC data sources. The IMS Universal Database resource adapter and IMS Universal JDBC driver internally manage the LL field on behalf of the application, so that the application program does not need to manage the segment length or the size of the I/O area. For read operations, the IMS Universal Database resource adapter and IMS Universal JDBC driver handle the offsets and lengths of all the segments and fields returned. By default, the SQL result set contains the LL field information. For update or insert operations, each instance of a variable-length segment is automatically expanded to contain the largest field (determined by the field length and offset) in the segment instance.

For SQL support for COBOL, COBOL applications are responsible for managing the LL field.

In a variable-length segment, some fields might be nullable. In IMS, a nullable field is a field that has a starting offset or combined offset and length larger than the minimum length of the segment. You can determine if a nullable field exists for a particular segment instance by comparing the LL value for the instance to the combined offset and length for the nullable field. If the LL value is less than the combined offset and length of the field, the field is null. For example, if a segment definition includes a field that starts at offset 50 and is length 5, it is nullable if the minimum length of the segment is less than 55. It is null for a particular segment instance if the LL value for that instance is less than 55.

You can also use the null indicator variable for COBOL or `java.sql.ResultSet.wasNull` method for IMS Universal Database resource adapter and IMS Universal JDBC driver to determine whether a nullable field exists in an instance of a variable-length segment without examining the LL data.

Using the LL field with the IMS Universal Database resource adapter and IMS Universal JDBC driver

By default, the LL field for a variable length segment is returned as a visible column for SQL queries, and the IMS Universal database resource adapter or IMS Universal JDBC driver manages the LL field on behalf of your application program. If you use this option, you can make LL field data inaccessible to any query type (including `SELECT *`) by setting the connection property to **false**.

Alternatively, you can use your application program to set and control the LL field. To do so, set the **llField** property to **true** through the standard properties list of either of the following interfaces:

```
java.sql.DriverManager.getConnection(String url, Properties properties)
```

```
com.ibm.ims.jdbc.IMSDataSource.setProperties(Properties properties)
```

When the `llField=true` property is set, the LL field is exposed as a normal column in the standard SQL result set for all operations. You can read, insert, or update the LL field data directly. Deleting the LL field data also deletes the rest of the associated database record. To set a field to the null state, set the length of the segment (the value of the LL field column) to be smaller than the offset of the field within the segment.

The LL field is 2 bytes long and must be handled as BINARY, SHORT, or USHORT data.

Using the LL field with COBOL

The LL field is treated as a normal column in the standard SQL result set for all operations. You can read, insert, or update the LL field data directly. Deleting the LL field data also deletes the rest of the associated database record. To set a field to the null state, set the length of the segment (the value of the LL field column) to be smaller than the offset of the field within the segment.

The LL field is 2 bytes long and must be handled as BINARY, SHORT, or USHORT data.

Checking for null field instances with the IMS Universal DL/I driver

Applications that use the IMS Universal DL/I driver always receive the LL field data for a variable-length segment. You can determine if a field is null in a segment instance in one of two ways: either compare the LL field data to the offset of the field, or use the `com.ibm.ims.dli.Path.wasNull()` method.

The `com.ibm.ims.dli.Path.wasNull()` method returns a boolean value for the null state of the last field that was read. The returned value is `true` if the field is null. You must attempt to read a field before calling the `wasNull()` method to determine whether the field is null.

Checking for null field instances in COBOL

Applications that use the SQL for COBOL always receive the LL field data for a variable-length segment. You can determine if a field is null in a segment instance in one of two ways: either compare the LL field data to the offset of the field, or use the null indicator variable.

Related concepts

[Variable-length segments \(Database Administration\)](#)

Related tasks

[“Declaring host variables and indicator variables” on page 510](#)

You can use host variables in SQL statements in your program to pass data between IMS and your application.

[“Declaring host variables and indicator variables in COBOL” on page 525](#)

You can use host variables and host structures in SQL statements in your program to pass data between IMS and your application.

Retrieving a set of rows by using a cursor

In an application program, you can retrieve a set of rows from IMS.

Use the following type of cursors to retrieve rows from a result segment:

- A row-positioned cursor retrieves at most a single row at a time from the result segment into host variables. At any point in time, the cursor is positioned on at most a single row. For information about how to use a row-positioned cursor, see [“Accessing data by using a row-positioned cursor” on page 546](#).

Cursors

A *cursor* is a mechanism that points to one row in a set of rows from the segment. Your application program can use a cursor to retrieve rows from a segment.

Types of cursors

You can declare row-positioned non-scrollable cursor to retrieve data from the result table.

Using a non-scrollable cursor:

The simplest type of cursor is a non-scrollable cursor. A row-positioned non-scrollable cursor moves forward through its result segment one row at a time.

Accessing data by using a row-positioned cursor

A row-positioned cursor points to a single row and retrieves at most a single row at a time from the result segment. You can specify a fetch request to specify which rows to retrieve, relative to the current cursor position.

To access data by using a row-positioned cursor:

1. Execute a DECLARE CURSOR statement to define the result segment on which the cursor operates. See [“Declaring a row cursor” on page 546](#).
2. Execute an OPEN CURSOR to make the cursor available to the application. See [“Opening a row cursor” on page 546](#).
3. Specify what the program is to do when all rows have been retrieved. See [“Specifying the action that the row cursor is to take when it reaches the end of the data” on page 547](#).
4. Execute the SQL statements to retrieve data from the segment. See [“Executing SQL statements by using a row cursor” on page 547](#).
5. Execute a CLOSE CURSOR statement to make the cursor unavailable to the application. See [“Closing a row cursor” on page 548](#).

Your program can have several cursors, each of which performs the previous steps.

Declaring a row cursor

Before you can use a row-positioned cursor to retrieve rows, you must declare the cursor. When you declare a cursor, you identify a set of rows that are to be accessed with the cursor.

To declare a row cursor, issue a DECLARE CURSOR statement. The DECLARE CURSOR statement names a cursor and specifies a prepared SELECT statement. The SELECT statement defines the criteria for the rows that are to make up the result segment.

The following example shows a cursor named C1 that is declared using a simple form of the DECLARE CURSOR statement for a prepared statement STMT.

```
EXEC SQLIMS
  DECLARE C1 CURSOR FOR STMT
END-EXEC.
```

Related reference

[DECLARE CURSOR \(Application Programming APIs\)](#)

[SELECT \(Application Programming APIs\)](#)

Opening a row cursor

After you declare a row cursor, you must tell IMS that you are ready to process the first row of the result segment. This action is called opening the cursor.

To open a row cursor, execute the OPEN statement in your program. IMS then uses the SELECT statement within DECLARE CURSOR to identify a set of rows. If you use parameter markers in the search condition of that SELECT statement, you must specify the values for the parameter makers with the USING clause. IMS uses the **current value** of the variables to select the rows. The result segment that satisfies the search condition might contain zero, one, or many rows.

An example of an OPEN statement is:

```
EXEC SQLIMS  
  OPEN C1  
END-EXEC.
```

An example of an OPEN statement when the prepared SELECT statement has parameter markers. Assume that the prepared SELECT statement has a parameter marker in the WHERE clause:

```
SELECT HOSPCODE, HOSPAME FROM PCB01.HOSPITAL  
WHERE HOSPNAME = ?
```

Use the OPEN statement with the USING clause to set the value of the parameter marker from host variable PARM1:

```
EXEC SQLIMS  
  OPEN C1 USING :PARM1  
END-EXEC.
```

Specifying the action that the row cursor is to take when it reaches the end of the data

Your program must be coded to recognize and handle an end-of-data condition whenever you use a row cursor to fetch a row.

To determine whether the program has retrieved the last row of data, test the SQLIMSCODE field for a value of 100 or the SQLIMSSTATE field for a value of '02000'. These codes occur when a FETCH statement has retrieved the last row in the result segment and your program issues a subsequent FETCH. For example:

```
IF SQLIMSCODE = 100 GO TO DATA-NOT-FOUND.
```

An alternative to this technique is to code the WHENEVER NOT FOUND statement. The WHENEVER NOT FOUND statement causes your program to branch to another part that then issues a CLOSE statement. For example, to branch to label DATA-NOT-FOUND when the FETCH statement does not return a row, use this statement:

```
EXEC SQLIMS  
  WHENEVER NOT FOUND GO TO DATA-NOT-FOUND  
END-EXEC.
```

For more information about the WHENEVER NOT FOUND statement, see [“Checking the execution of SQL statements” on page 521](#).

Executing SQL statements by using a row cursor

You can use row cursors to execute FETCH statements.

Execute a FETCH statement to copy data from a row of the result segment into one or more host variables.

The following example shows a FETCH statement that retrieves selected fields from the hospital segment:

```
EXEC SQLIMS  
  FETCH C1 INTO :HOSPCODE, :HOSPNAME, :WARDNAME, :PATNAME  
END-EXEC.
```

When your program executes the FETCH statement, IMS positions the cursor on a row in the result segment. That row is called the *current row*. IMS then copies the current row contents into the program host variables that you specify on the INTO clause of FETCH. This sequence repeats each time you issue FETCH, until you process all rows in the result segment.

Closing a row cursor

Close a row cursor when it finishes processing rows if you want to free the resources or if you want to use the cursor again. Otherwise, you can let IMS automatically close the cursor when the current transaction terminates or when your program terminates.

To free the resources that are held by the cursor, close the cursor explicitly by issuing the CLOSE statement.

If you want to use the rowset cursor again, reopen it.

Issue a CLOSE statement.

An example of a CLOSE statement looks like this:

```
EXEC SQLIMS  
  CLOSE C1  
END-EXEC.
```

Commit or roll back data

After your application issues SQL statements to modify data in your IMS database, your application might want to commit or roll back database changes. To commit or roll back IMS database changes, use IMS DB system services DL/I calls.

For example, you might issue a ROLB to roll back the changes and issue a CHKP to commit the changes.

The SQL keyword COMMIT and ROLLBACK are currently not supported.

Related reference

[“Recovering databases and maintaining database integrity” on page 263](#)

You can issue checkpoints, restart programs, and maintain database integrity in your application programs.

Preparing an application to run on IMS

To prepare and run applications that contain SQL statements, you must coprocess, compile, and link-edit them.

Tip: To avoid rework, first test your SQL statements using IMS Enterprise Suite Explorer for Development. Then compile your program with SQL statements, and resolve all compiler errors. Finally, proceed with the deployment and compile your COBOL program using the IMS coprocessor to translate the SQL statements.

Processing SQL statements

The first step in preparing an SQL application to run is to process the SQL statements in the program. To process the statements, use the IMS coprocessor. During this step, the SQL statements are replaced with calls to IMS language interface module (DFSLI000).

For COBOL applications, you can use one of the following techniques to process SQL statements:

- Invoke the IMS coprocessor for the host language that you are using as you compile your program. You can use the IMS coprocessor with COBOL host compilers. To invoke the IMS coprocessor, specify the SQLIMS compiler option followed by its suboptions.
 - For COBOL, you need Enterprise COBOL for z/OS Version 5 Release 1 or later to use this technique. For more information about the COBOL IMS coprocessor, see *Enterprise COBOL for z/OS Programming Guide*.

The IMS coprocessor performs precompiler functions at compile time. When you use the IMS coprocessor, the compiler scans the program and returns the modified source code.

Processing SQL statements by using the IMS coprocessor

The IMS coprocessor processes SQL statements at compile time.

To process SQL statements by using the IMS coprocessor, perform the following action:

- Submit a JCL job to process the IMS application that contains SQL statement. Include the following information:

- Specify the SQLIMS compiler option when you compile your program:

The SQLIMS compiler option indicates that you want the compiler to invoke the IMS coprocessor. To use IMS co-processor and use the default options, specify just SQLIMS.

For example:

```
//COBOL1 EXEC PGM=IGYCRCTL,  
// PARM='LIST,XREF,CP(37),SQLIMS'
```

The following default option values will be used by the IMS coprocessor:

PERIOD

Period (.) will be recognized as the decimal point indicator in decimal or floating point literals within the SQL statement.

APOSTSQL

Apostrophe (') as the string delimiter and the double quotation mark (") as the SQL escape character within SQL statements.

In addition, the COBOL CODEPAGE option will be used to determine what CCSID the source program is written. Currently, only EBCDIC CCSID 37 and 1140 are supported.

- Include DD statements for the following data sets in the JCL for your compile step:

- IMS load library (IGYV5R10.SQGYCOMP)

The IMS coprocessor calls IMS modules to process the SQL statements. You therefore need to include the name of the IMS load library data set in the STEPLIB concatenation for the compile step.

- Library for SQL INCLUDE statements

If your program contains SQL INCLUDE *member-name* statements that specify secondary input to the source program, you need to also specify the data set for *member-name*. Include the name of the data set that contains *member-name* in the SYSLIB concatenation for the compile step.

Part 6. Java application development for IMS

IMS provides support for developing applications using the Java programming language.

Chapter 38. IMS solutions for Java development overview

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

The IMS solutions for Java development include the IMS Universal drivers, the IMS Java dependent region resource adapter, and the IMS Transaction Manager Resource Adapter (IMS TM Resource Adapter). These solutions are provided through the Java On Demand feature FMID. For solutions that can also run on distributed platforms, you can download them from the IMS download site at [IMS Transaction Manager \(TM\) Resource Adapter Version 14](#).

For developing Java applications that run in IMS, see the [Java in IMS solution adoption kit \(System Definition\)](#) for end-to-end guidance, including tasks for system programmers to set up the environment.

IMS Universal drivers

The IMS Universal drivers are a set of Java drivers and resource adapters that enable access to IMS from z/OS and distributed (non-z/OS) platforms. The IMS Universal drivers are built on industry standards and open specifications. Two types of connectivity are supported by the IMS Universal drivers: local connectivity to IMS databases on the same LPAR (*type-2 connectivity*) and distributed connectivity through TCP/IP (*type-4 connectivity*). Java applications that use the type-2 IMS Universal drivers must reside on the same logical partition (LPAR) as the IMS subsystem. Java applications that use the type-4 IMS Universal drivers can reside on the same logical partition (LPAR) or on a different LPAR from the IMS subsystem.

The IMS Universal drivers enable access to IMS from multiple environments, including:

- WebSphere Application Server for z/OS
- CICS Transaction Server for z/OS
- IMS on the host in Java batch processing (JBP) and Java message processing (JMP) regions

The IMS Universal drivers include:

- IMS Universal Database resource adapter: A Java EE Connector Architecture (JCA) 1.6-compliant resource adapter
- IMS Universal JDBC driver: A Java Database Connectivity (JDBC) driver that implements the JDBC 4.0 API
- IMS Universal DL/I driver: A Java API for making calls with traditional DL/I programming semantics

IMS Java dependent region resource adapter

The IMS Java dependent region resource adapter is a set of Java classes and interfaces that support IMS database access and IMS message queue processing within JBP and JMP regions. The IMS Java dependent region resource adapter provides Java application programs running in JMP or JBP regions with similar DL/I functionality to that provided in message processing program (MPP) and non-message-driven BMP regions, such as:

- Accessing IMS message queues to read and write messages
- Performing program switches
- Commit and rollback processing
- Accessing IMS databases in an IMS DB/TM environment
- Accessing GSAM databases in IMS DB/TM and DCCTL environments
- Database recovery (CHKP/XRST)

IMS TM Resource Adapter

Generic Java EE 1.7 or later certified application servers, including WebSphere Application Server Liberty Profile, are supported if the IMS TM resource adapter installation verification program (IVP) can run successfully. Only a subset of IMS TM resource adapter functions are supported. See the restrictions topic for usage restrictions. For generic Java EE 1.7 or later certified application servers, IBM support is provided for specific functions if:

- The problem can be recreated in the supported versions of WebSphere Application Server or WebSphere Application Server Liberty Profile.
- The problem is confirmed to be caused by the IMS TM resource adapter through diagnostic traces.

While the IMS Universal drivers and the IMS Java dependent region resource adapter provide the interfaces and classes for accessing IMS data through a Java or Java EE application, the IMS TM resource adapter, as its name suggests, is primarily for interacting with IMS transactions through a Java EE application.

Related concepts

[“Programming with the IMS Universal drivers” on page 559](#)

Use these topics to design, write, and maintain application programs for IMS 15.3 using the IMS Universal drivers.

[“Programming Java dependent regions” on page 673](#)

Use these topics to design, write, and maintain application programs for running in the Java dependent regions.

Related tasks

[Configuring external Java environment connections \(Communications and Connections\)](#)

Chapter 39. Comparison of hierarchical and relational databases

The following information describes the differences between the hierarchical model for IMS databases and the standard relational database model.

A database segment definition defines the fields for a set of segment instances similar to the way a relational table defines columns for a set of rows in a table. In this way, segments relate to relational tables, and fields in a segment relate to columns in a relational table.

The name of an IMS segment becomes the table name in an SQL query, and the name of a field becomes the column name in the SQL query.

A fundamental difference between segments in a hierarchical database and tables in a relational database is that, in a hierarchical database, segments are implicitly joined with each other. In a relational database, you must explicitly join two tables. A segment instance in a hierarchical database is already joined with its parent segment and its child segments, which are all along the same hierarchical path. In a relational database, this relationship between tables is captured by foreign keys and primary keys.

This section compares the Dealership sample database, to a relational representation of the database. The Dealership sample DBDs are available with the IMS Enterprise Suite Explorer for Development, in the <installation location>\IMS Explorer samples directory.

Important: This information provides only a comparison between relational and hierarchical databases.

The Dealership sample database contains five segment types, which are shown in the following figure. The root segment is the Dealer segment. Under the Dealer segment is its child segment, the Model segment. Under the Model segment are its children: the segments Order, Sales, and Stock.

The following figure shows the structure and each segment of the Dealership sample database.




Figure 102. Segments of the Dealership sample database

The Dealer segment identifies a dealer that sells cars. The segment contains a dealer name in the field DLRNAME, and a unique dealer number in the field DLRNO.

Dealers carry car types, each of which has a corresponding Model segment. A Model segment contains a type code in the field MODTYPE.

Each car that is ordered for the dealership has an Order segment. A Stock segment is created for each car that is available for sale in the dealer's inventory. When the car is sold, a Sales segment is created.

The following figure shows a relational representation of the IMS database record shown in [Figure 102](#) on page 555.




Figure 103. Relational representation of the Dealership sample database

If a segment does not have a unique key, which is similar to a primary key in relational databases, view the corresponding relational table as having a generated primary key added to its column (field) list. An example of a generated primary key is in the Model table (segment) of the figure above. Similar to referential integrity in relational databases, you cannot insert, for example, an Order (child) segment to the database without it being a child of a specific Model (parent) segment.

Also note that the field (column) names have been renamed. You can rename segments and fields to more meaningful names by using the IMS Explorer for Development.

An occurrence of a segment in a hierarchical database corresponds to a row (or tuple) of a table in a relational database.

The following figure shows three Dealership database records.




Figure 104. Segment occurrences in the Dealership sample database

The Dealer segment occurrences have dependent Model segment occurrences.

The following figure shows the relational representation of the dependent model segment occurrences.




Figure 105. Relational representation of segment occurrences in the Dealership database

In the following example that shows the SELECT statement of an SQL call, Model is a segment name that is used as a table name in the query:

```
SELECT * FROM Model
```

The following example, ModelTypeCode is the name of a field that is contained in the Model segment and it is used in the SQL query as a column name:

```
SELECT * FROM Model WHERE ModelTypeCode = '062579'
```

In the two preceding examples, Model and ModelTypeCode are alias names that are assigned with the EXTERNALNAME parameter of the SEGM and FIELD statements in the DBD, respectively. The EXTERNALNAME parameter is an optional parameter that specifies an external alias name for client applications to use when referencing the field or segment, and does not need to conform to the

8-character limit for host resource names. External alias names are only used when the IMS catalog is active. If the IMS catalog is active but no alias name is specified on the EXTERNALNAME parameter for a segment or field, use the 8-character IMS name for the resource instead.

Chapter 40. Programming with the IMS Universal drivers

Use these topics to design, write, and maintain application programs for IMS 15.3 using the IMS Universal drivers.

Related concepts

[“IMS solutions for Java development overview” on page 553](#)

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

[“Overview of the IMS Java dependent regions” on page 673](#)

The IMS Java dependent regions are two types of IMS dependent regions that provide a Java Virtual Machine (JVM) environment for Java applications: Java message processing (JMP) regions and Java batch processing (JBP) regions.

Related reference

[Software requirements for Java application programs that use the IMS Universal drivers or JDR resource adapter \(Release Planning\)](#)

IMS Universal drivers overview

The IMS Universal drivers are software components that provide Java applications with connectivity and access to IMS databases from z/OS and distributed environments through TCP/IP. Java applications that use the type-2 IMS Universal drivers must reside on the same logical partition (LPAR) as the IMS subsystem. Java applications that use the type-4 IMS Universal drivers can reside on the same logical partition (LPAR) or on a different LPAR from the IMS subsystem.

Programming approaches

The IMS Universal drivers provide an application programming framework that offers multiple options for access to IMS data. These programming options include:

IMS Universal Database resource adapter

Provides connectivity to IMS databases from a Java Platform, Enterprise Edition (Java EE) environment, and access to IMS data using the Common Client Interface (CCI) and Java Database Connectivity (JDBC) interfaces.

IMS Universal JDBC driver

Provides a stand-alone JDBC 4.0 driver for making SQL-based database calls to IMS databases.

IMS Universal DL/I driver

Provides a stand-alone Java API for writing granular queries to IMS databases using programming semantics similar to traditional DL/I calls.

Open standards

The IMS Universal drivers are built on the following industry open standards and interfaces:

Java EE Connector Architecture (JCA)

JCA is the Java standard for connecting Enterprise Information Systems (EISs) such as IMS into the Java EE framework. Using JCA, you can simplify application development and take advantage of the services that can be provided by a Java EE application server, such as connection management, transaction management, and security management. The Common Client Interface (CCI) is the interface in JCA that provides access from Java EE clients, such as Enterprise JavaBeans (EJB) applications, JavaServer Pages (JSP), and Java servlets, to backend IMS subsystems.

Java Database Connectivity (JDBC)

JDBC is the SQL-based standard interface for database access. It is the industry standard for database-independent connectivity between the Java programming language and any database that has implemented the JDBC interface.

Distributed Relational Database Architecture (DRDA) specification

DRDA is an open architecture that enables communication between applications and database systems on disparate platforms. These applications and database systems can be provided by different vendors and the platforms can be different hardware and software architectures. DRDA provides distributed database access with built-in support for distributed, two-phase commit transactions.

Related reference

[DRDA DDM command architecture reference \(Application Programming APIs\)](#)

[Software requirements for Java application programs that use the IMS Universal drivers or JDR resource adapter \(Release Planning\)](#)

Distributed and local connectivity with the IMS Universal drivers

The IMS Universal drivers support distributed (*type-4*) and local (*type-2*) connectivity to IMS databases.

Distributed connectivity with the type-4 IMS Universal drivers

With type-4 connectivity, the IMS Universal drivers can run on any platform that supports TCP/IP and a Java Virtual Machine (JVM), including z/OS. To access IMS databases, the type-4 IMS Universal drivers first establish a TCP/IP-based socket connection to IMS Connect. IMS Connect is responsible for routing the request to the IMS databases by using the Open Database Manager (ODBM), and sending the response back to the client application. The DRDA protocol is used internally in the implementation of the type-4 IMS Universal drivers. You do not need to know DRDA to use the type-4 IMS Universal drivers.

The type-4 IMS Universal drivers support two-phase commit (XA) transactions. IMS Connect builds the necessary z/OS Resource Recovery Services (RRS) structure to support the two-phase commit protocol. If two-phase commit transactions are not used, RRS is not required.

To establish a connection to IMS, the **driverType** connection property must be set to indicate distributed (*type-4*) connectivity to IMS

After successful authentication, the IMS Universal drivers send other socket connection information, such as program specification block (PSB) name and IMS database subsystem, to IMS Connect and ODBM to allocate the PSB to connect to the database.

A connection to an IMS database is established only when a program specification block (PSB) is allocated. Authorization for a particular PSB is done by the ODBM component during the allocation of a PSB.

The type-4 IMS Universal drivers support connection pooling, which limits the time that is needed for allocation and deallocation of TCP/IP socket connections. To maximize connection reuse, only the socket attributes of a connection are pooled. These attributes include the IP address and port number that the host IMS Connect is listening on. As a result, the physical socket connection can be reused and additional attributes can be sent on this socket to connect to an IMS database. When a client application of the type-4 IMS Universal drivers makes a connection to IMS, this means:

- A one-to-one relationship is established between a client socket and an allocated PSB that contains one or more IMS databases.
- A one-to-many relationship is established between IMS Connect and the possible number of database connections it can handle at one time.
- IMS Connect does the user authentication.
- ODBM ensures that the authenticated user is authorized to access the given PSB.

The following figure shows how the type-4 IMS Universal drivers route communications between your Java client applications that are running in a distributed environment and an IMS subsystem.




Figure 106. Distributed connectivity with the type-4 IMS Universal drivers

You can also use the type-4 IMS Universal drivers if your Java clients are running in a z/OS environment but are on a separate logical partition from the IMS subsystem. Use type-4 connectivity from a z/OS environment if you want to isolate the application runtime environment from the IMS subsystem environment.

Local connectivity with the type-2 IMS Universal drivers

Local connectivity with the type-2 IMS Universal drivers is targeted for the z/OS platform and runtime environments. Use type-2 connectivity to connect to IMS subsystems in the same logical partition (LPAR).

The following table shows the z/OS runtime environments that support client applications of the type-2 IMS Universal drivers.

Table 97. z/OS runtime environment support for the type-2 IMS Universal drivers	
z/OS runtime environment	Type-2 IMS Universal drivers supported
WebSphere Application Server for z/OS or WebSphere Application Server Liberty	IMS Universal Database resource adapter
IMS Java dependent regions (BMP, JMP, JBP, IFP, and MPR regions); CICS	- IMS Universal DL/I driver - IMS Universal JDBC driver
CICS OSGi JVM server	- IMS Universal DL/I driver - IMS Universal JDBC driver

Table 97. z/OS runtime environment support for the type-2 IMS Universal drivers (continued)

z/OS runtime environment	Type-2 IMS Universal drivers supported
CICS Liberty JVM server	• Not supported
Db2 Java Stored Procedures	• Not supported
z/OS batch	• Not supported

Because it runs on the same LPAR as the IMS subsystem, during connection time, a client application of the type-2 IMS Universal drivers does not need to supply an IP address, port number, user ID, or password. The **driverType** property must be set to indicate local (*type-2*) connectivity to IMS.

The following figure shows how the type-2 IMS Universal drivers route communications between your Java client applications that are running in an LPAR inside a z/OS mainframe environment and an IMS subsystem that is located in the same LPAR.




Figure 107. Local connectivity with the type-2 IMS Universal drivers

RRSLocalOption connectivity type

In addition to type-4 and type-2 connectivity, the RRSLocalOption connectivity type is supported by the IMS Universal Database resource adapter running on WebSphere Application Server for z/OS. With RRSLocalOption connectivity, applications using the IMS Universal Database resource adapter do not issue commit or rollback calls. Instead, transaction processing is managed by WebSphere Application Server for z/OS. Two-phase commit (XA) transaction processing is not supported with RRSLocalOption connectivity type.

Related concepts

[IMS Connect support for access to IMS DB \(Communications and Connections\)](#)

[CSL ODBM administration \(System Administration\)](#)

Configuring JAXB with JDK 9 or higher

The IMS Universal drivers leverage JAXB, part of the `java.se.ee.module` for XML parsing. If you are using JDK 9 or higher, add JAXB to the classpath to enable full functionality of IMS Universal drivers.

JAXB is an XML to Java binding technology that supports transformation between schema and Java objects and between XML instance documents and Java object instances. JAXB consists of a runtime application programming interface (API) and accompanying tools that simplify access to XML documents. JAXB also helps to build XML documents that both conform and validate to the XML schema. Java API for XML-Based Web Services (JAX-WS) leverages the JAXB API and tools as the binding technology for mappings between Java objects and XML documents. JAX-WS tooling relies on JAXB tooling for default data binding for two-way mappings between Java objects and XML documents.

JAXB is required by the IMS Universal drivers for processing the XML metadata obtained locally or from IMS Catalog.



Attention: If you are using JDK 9 or higher, add JAXB to the classpath to enable the full functionality of IMS Universal drivers.

If you are configuring JAXB using JDK 9 or later, please note that the Java EE and CORBA modules are deprecated in JDK 9 and 10, and are removed completely in JDK 11 and later. Enabling any deprecated modules within the JDK itself is not recommended.

Adding JAXB in a Java Project in Eclipse

All applications leveraging IMS Universal drivers need the following JARS added to their classpath:

- <https://mvnrepository.com/artifact/javax.activation/activation/1.1.1>
- <https://mvnrepository.com/artifact/javax.xml.bind/jaxb-api/2.3.1>
- <https://mvnrepository.com/artifact/com.sun.xml.bind/jaxb-osgi/2.3.1>



Attention: A more recent version of the JAXB libraries mentioned above might be available. Check for the newer versions of JAXB 2.x.x.

Adding JAXB in a Maven project

The following dependencies are required if you are using a Maven project:

```
<dependency>
  <groupId>org.glassfish.jaxb</groupId>
  <artifactId>jaxb-runtime</artifactId>
  <version>2.3.1</version>
</dependency>
<dependency>
  <groupId>javax.xml.bind</groupId>
  <artifactId>jaxb-api</artifactId>
  <version>2.3.1</version>
</dependency>
<dependency>
  <groupId>javax.activation</groupId>
  <artifactId>activation</artifactId>
  <version>1.1.1</version>
</dependency>
```

The following JVM argument is required both in the case of Eclipse and Maven:

-Djavax.xml.bind.context.factory=com.sun.xml.bind.v2.ContextFactory

Comparison of IMS Universal drivers programming approaches for accessing IMS

Depending on your IT infrastructure, solution architecture, and application design, choose the IMS Universal drivers programming approach that is best for your development scenario.

The following table lists the recommended IMS Universal drivers programming approach to use, based on the application programmer's choice of application platform, data access method, and transaction processing option.

Table 98. Comparison of programming approaches for accessing IMS

Application platform	Data access method	Transaction processing required	Recommended approach
WebSphere Application Server for distributed platforms or WebSphere Application Server for z/OS or WebSphere Application Server Liberty	CCI programming interface to perform SQL or DL/I data operations.	Local transaction processing only.	Use the IMS Universal Database resource adapter with local transaction support (imsudbLocal.rar), and make SQL calls with the SQLInteractionSpec class or DL/I calls with the DLIIInteractionSpec class.
	CCI programming interface to perform SQL or DL/I data operations.	Two-phase (XA) commit processing ¹ or local transaction processing.	Use the IMS Universal Database resource adapter with XA transaction support (imsudbXA.rar), and make SQL calls with the SQLInteractionSpec class or DL/I calls with the DLIIInteractionSpec class.
	JDBC programming interface to perform SQL data operations.	Local transaction processing only.	Use the IMS Universal JCA/JDBC driver version of the IMS Universal Database resource adapter with local transaction support (imsudbJLocal.rar), and make SQL calls with the JDBC API.
	JDBC programming interface to perform SQL data operations.	Two-phase (XA) commit processing ¹ or local transaction processing.	Use the IMS Universal JCA/JDBC driver version of the IMS Universal Database resource adapter with XA transaction support (imsudbJXA.rar), and make SQL calls with the JDBC API.

Table 98. Comparison of programming approaches for accessing IMS (continued)

Application platform	Data access method	Transaction processing required	Recommended approach
Standalone Java application (outside a Java EE application server) that resides on a distributed platform or a z/OS platform	JDBC programming interface to perform SQL data operations.	Two-phase (XA) commit processing ² or local transaction processing.	Use the IMS Universal JDBC driver (imsudb.jar ³), and make SQL calls with the JDBC API.
	Traditional DL/I programming semantics to perform data operations.	Two-phase (XA) commit processing ² or local transaction processing.	Use the IMS Universal DL/I driver (imsudb.jar ³), and make DL/I calls with the PCB class.
Standalone non-Java application that resides on a distributed platform or a z/OS platform	Data access using DRDA protocol.	Two-phase (XA) commit processing or local transaction processing.	Use a programming language of your choice to issue DDM commands to IMS Connect. The application programmer is responsible for implementing the two-phase commit mechanism.

Note:

1. XA transaction support is available only with type-4 connectivity.
2. The driver is enabled for local and XA transactions, but the application programmer is responsible for implementing the two-phase commit mechanism. XA transaction support is available only with type-4 connectivity.
3. The JDBC driver and the DL/I driver are provided in the same .jar file because the SQL calls made through the JDBC driver are converted into DL/I calls by using the methods and classes in the DL/I driver.

Support for variable-length database segments with the IMS Universal drivers

The IMS Universal database resource adapter and the IMS Universal JDBC driver manage variable-length segments on behalf of client application programs by default. Application programs that use the IMS Universal DL/I driver must manage the LL field data for variable-length segments.

Using the LL field with the IMS Universal database resource adapter or IMS Universal JDBC driver

By default, the LL field for a variable length segment is returned as a visible column for SQL queries, and the IMS Universal database resource adapter or IMS Universal JDBC driver manages the LL field on behalf of your application program. If you use this option, you can make LL field data inaccessible to any query type (including SELECT *) by setting the connection property to **false**.

Alternatively, you can use your application program to set and control the LL field. To do so, set the **llField** property to **true** through the standard properties list of either of the following interfaces:

```
java.sql.DriverManager.getConnection(String url, Properties properties)
```

```
com.ibm.ims.jdbc.IMSDataSource.setProperties(Properties properties)
```

When the `llField=true` property is set, the LL field is exposed as a normal column in the standard SQL result set for all operations. You can read, insert, or update the LL field data directly. Deleting the LL field data also deletes the rest of the associated database record. To set a field to the null state, set the length of the segment (the value of the LL field column) to be smaller than the offset of the field within the segment.

The LL field is 2 bytes long and must be handled as BINARY, SHORT, or USHORT data.

You can also use the `java.sql.ResultSet.isNull` method to determine whether a nullable field exists in an instance of a variable-length segment without examining the LL data.

Checking for null field instances with the IMS Universal DL/I driver

Applications that use the IMS Universal DL/I driver always receive the LL field data for a variable-length segment. You can determine if a field is null in a segment instance in one of two ways: either compare the LL field data to the offset of the field, or use the `com.ibm.ims.dli.Path.isNull()` method.

The `com.ibm.ims.dli.Path.isNull()` method returns a boolean value for the null state of the last field that was read. The returned value is `true` if the field is null. You must attempt to read a field before calling the `isNull()` method to determine whether the field is null.

Related concepts

[Variable-length segments \(Database Administration\)](#)

Related tasks

[How to specify variable-length segments \(Database Administration\)](#)

Support for flattening complex structures

The `flattenTables` connection property produces a flattened view of a database table. Although the copybook structure in the IMS catalog is unchanged, the information about the structure of the table is altered for that particular connection. Enabling the `flattenTables` connection property can simplify the process of querying a database table.

Complex structures that are displayed according to the `flattenTables` connection property are formatted according to these naming conventions:

- Static arrays are referenced by the name of the array, index of the array, and name of the field.

An example of this convention is `CLASSES_2_INST`, where `CLASSES` is the name of the array, `2` is the index of the array, and `INST` is the name of the field.

- Only the sub-elements of structures are displayed in the flattened structure view (not the names of the structures themselves).

Restriction: The `flattenTables` connection property supports static arrays and structures only. Dynamic arrays are not altered.

The following view displays a copybook structure with the `flattenTables` connection property disabled.

```
01 SEGM.
  05 CLASSES OCCURS 4 TIMES.
    10 INST PIC X(15).
    10 GRADE PIC X(1).
    10 BOOKS OCCURS 2 TIMES.
      15 AUTHOR.
        20 FIRSTNAME PIC X(10).
        20 LASTNAME PIC X(15).
      15 TITLE PIC X(20).
```

The following view displays the same copybook structure as in the previous example, but with the `flattenTables` connection property enabled.

```
01 SEGM.
  05 CLASSES_1_INST PIC X(15).
  05 CLASSES_1_GRADE PIC X(1).
```

```

05 CLASSES_1_BOOKS_1_FIRSTNAME PIC X(10).
05 CLASSES_1_BOOKS_1_LASTNAME PIC X(15).
05 CLASSES_1_BOOKS_1_TITLE PIC X(20).
05 CLASSES_1_BOOKS_2_FIRSTNAME PIC X(10).
05 CLASSES_1_BOOKS_2_LASTNAME PIC X(15).
05 CLASSES_1_BOOKS_2_TITLE PIC X(20).
05 CLASSES_2_INST PIC X(15).
05 CLASSES_2_GRADE PIC X(1).
05 CLASSES_2_BOOKS_1_FIRSTNAME PIC X(10).
05 CLASSES_2_BOOKS_1_LASTNAME PIC X(15).
05 CLASSES_2_BOOKS_1_TITLE PIC X(20).
05 CLASSES_2_BOOKS_2_FIRSTNAME PIC X(10).
05 CLASSES_2_BOOKS_2_LASTNAME PIC X(15).
05 CLASSES_2_BOOKS_2_TITLE PIC X(20).
05 CLASSES_3_INST PIC X(15).
05 CLASSES_3_GRADE PIC X(1).
05 CLASSES_3_BOOKS_1_FIRSTNAME PIC X(10).
05 CLASSES_3_BOOKS_1_LASTNAME PIC X(15).
05 CLASSES_3_BOOKS_1_TITLE PIC X(20).
05 CLASSES_3_BOOKS_2_FIRSTNAME PIC X(10).
05 CLASSES_3_BOOKS_2_LASTNAME PIC X(15).
05 CLASSES_3_BOOKS_2_TITLE PIC X(20).
05 CLASSES_4_INST PIC X(15).
05 CLASSES_4_GRADE PIC X(1).
05 CLASSES_4_BOOKS_1_FIRSTNAME PIC X(10).
05 CLASSES_4_BOOKS_1_LASTNAME PIC X(15).
05 CLASSES_4_BOOKS_1_TITLE PIC X(20).
05 CLASSES_4_BOOKS_2_FIRSTNAME PIC X(10).
05 CLASSES_4_BOOKS_2_LASTNAME PIC X(15).
05 CLASSES_4_BOOKS_2_TITLE PIC X(20).

```

Related tasks

[“Connecting using the IMS Universal Database resource adapter in a managed environment” on page 573](#)

In a managed (or three-tier) environment, your Java EE application interacts with a Java EE application server, such as WebSphere Application Server, and the IMS Universal Database resource adapter to communicate with an IMS database.

[“Connecting using the IMS Universal JCA/JDBC driver in a managed environment” on page 579](#)

To access IMS databases using a JDBC programming interface in a managed (or three-tier) environment, you need to deploy the IMS Universal JCA/JDBC driver on your Java EE application server and configure the connection properties.

[“Connecting to an IMS database using the JDBC DataSource interface” on page 595](#)

Using the DataSource interface is the preferred way to connect to IMS from your IMS Universal JDBC driver application.

[“Connecting to an IMS database by using the JDBC DriverManager interface” on page 601](#)

A JDBC application can establish a connection to a data source using the JDBC DriverManager interface, which is part of the `java.sql` package.

Generating the runtime Java metadata class

To connect to an IMS database using the IMS Universal drivers, you need to include on your Java classpath the Java metadata class that provides the database view.

Note: If you are using the IMS catalog, an IMS Universal drivers application program can obtain the necessary metadata directly from the catalog database without a Java metadata class file.

The Java metadata class represents the application view information specified by a program specification block (PSB) and its related Program Control Blocks (PCBs). The Java metadata class provides a one-to-one mapping to the segments and fields defined in the PSB.

The Java metadata class must be compiled and made available through the classpath for any Java application attempting to access IMS data using that PSB.

During database connection setup, pass the name of this metadata class to the resource adapter or JDBC driver. The Java metadata class is used at runtime by the IMS Universal drivers to process both SQL and Java-based DL/I calls.

The default segment encoding of the database metadata class produced by the IMS Explorer for Development is cp1047. To change the segment encoding, use the `com.ibm.ims.base.DLIBaseSegment.setDefaultEncoding` method.

Related concepts

[IMS Explorer for Development overview](#)

Hospital database example

The code examples for the IMS Universal drivers application programming topics use the Hospital database.

The following figure shows the hierarchical structure of the segments in the Hospital database.




Figure 108. Segments of the Hospital database

Each node in the figure represents a segment:

- The HOSPITAL segment is the root segment in the database.
- PAYMENTS and WARD are child segments of the HOSPITAL segment.
- WARD has a direct descendent segment named PATIENT.
- ILLNESS and BILLING are the child segments of the PATIENT.
- ILLNESS has a child segment named TREATMENT that stores details about patient treatment.
- The child segment of ILLNESS, DOCTOR, is the lowest level segment in the database hierarchy.

The tables that follow show the layouts of each segment in the Hospital database.

HOSPITAL segment

The following table shows the HOSPITAL segment, which has two fields:

- The hospital code (HOSPCODE)
- The hospital name (HOSPNAME)

HOSPCODE is a unique key field.

Field name	Field length (in bytes)
HOSPCODE	12
HOSPNAME	17

PAYMENTS segment

The following table shows the PAYMENTS segment, which has two fields:

- The patient number (PATNUM)
- The payment amount (AMOUNT)

Field name	Field length (in bytes)
PATNUM	4
AMOUNT	8

WARD segment

The following table shows the WARD segment, which has five fields:

- The ward number (WARDNO)
- The ward name (WARDNAME)
- The patient count (PATCOUNT)
- The nurse count (NURCOUNT)
- The doctor count (DOCCOUNT)

WARDNO is a unique key field.

Field name	Field length (in bytes)
WARDNO	2
WARDNAME	4
PATCOUNT	8
NURCOUNT	4
DOCCOUNT	2

PATIENT segment

The following table shows the PATIENT segment, which has two fields:

- The patient number (PATNUM)
- The patient name (PATNAME)

PATNUM is a unique key field.

Field name	Field length (in bytes)
PATNUM	12
PATNAME	17

ILLNESS segment

The following table shows the ILLNESS segment, which has one field:

- The illness name (ILLNAME)

Field name	Field length (in bytes)
ILLNAME	15

TREATMNT segment

The following table shows the TREATMNT segment, which has three fields:

- The day of treatment (TREATDAY)
- The type of treatment (TREATMNT)
- The treatment comments (COMMENTS)

Field name	Field length (in bytes)
TREATDAY	8
TREATMNT	15
COMMENTS	10

DOCTOR segment

The following table shows the DOCTOR segment, which has two fields:

- The doctor number (DOCTNO)
- The doctor name (DOCNAME)

Field name	Field length (in bytes)
DOCTNO	4
DOCNAME	20

BILLING segment

The following table shows the BILLING segment, which has two fields:

- The bill amount (AMOUNT)
- The bill comments (COMMENTS)

Field name	Field length (in bytes)
AMOUNT	8
COMMENTS	20

Related concepts

[“Specifying segment search arguments using the SSAList interface” on page 646](#)

The SSAList interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Programming using the IMS Universal Database resource adapter

This information describes how to write programs with the IMS Universal Database resource adapter to access IMS databases.

Overview of the IMS Universal Database resource adapter

The IMS Universal Database resource adapter is based on the Java Platform, Enterprise Edition (Java EE) Connector Architecture (JCA) 1.6 standard. The purpose of the JCA is to connect Enterprise Information Systems (EISs), such as IMS, into the Java EE platform. JCA provides a number of services that are managed by a Java EE application server. These services include security credential management, connection pooling, and transaction management.

These services are provided by means of system level contracts between the IMS Universal Database resource adapter and the Java EE application server, without the need for additional coding by the application programmer.

The JCA specification defines a programming interface called the Common Client Interface (CCI). This interface is used to communicate with any EIS. The IMS Universal Database resource adapter implements the CCI for interactions with IMS databases. The CCI interfaces for the IMS Universal Database resource adapter are in the `com.ibm.ims.db.cci` package. The CCI implementation provided by IMS allows applications to make either SQL or DL/I calls to access the IMS database.

In addition to the CCI interface provided by the IMS Universal Database resource adapter, you can also write JDBC applications to access your IMS data from a managed environment, while leveraging the Java EE services provided by the application server. This capability is provided by the IMS Universal JCA/JDBC driver version of the IMS Universal Database resource adapter. The IMS Universal JCA/JDBC driver is based on the Java Platform, Enterprise Edition (Java EE) Connector Architecture (JCA) 1.6 and Java Database Connectivity (JDBC) 4.0 standard.

The IMS Universal Database resource adapter communicates with IMS Connect as the TCP/IP endpoint to access IMS.

Preparing to write a Java application with the IMS Universal drivers

Java application programs that use the IMS Universal drivers require the Java Development Kit (JDK) 7.0. Java programs that run in JMP and JBP regions require JDK 7.0 or later. Java application programs that use the IMS Universal drivers must have access to database metadata in order to interact with IMS databases. This metadata can either be accessed directly in the IMS catalog database or it can be generated as a Java metadata class with the IMS Enterprise Suite Explorer for Development.

Transaction types and programming interfaces supported by the IMS Universal Database resource adapter

The IMS Universal Database resource adapter provides four types of support for optimized transaction management and performance.

The types of transaction support provided by the IMS Universal Database resource adapter are:

IMS Universal Database resource adapter with local transaction support (imsudbLocal.rar)

This resource adapter provides a CCI programming interface and LocalTransaction support when deployed on any supported Java EE application server.

IMS Universal Database resource adapter with XA transaction support (imsudbXA.rar)

This resource adapter provides a CCI programming interface and both XATransaction and LocalTransaction support when deployed on any supported Java EE application server.

IMS Universal JCA/JDBC driver with local transaction support (imsudbJLocal.rar)

This resource adapter provides a JDBC programming interface and LocalTransaction support when deployed on any supported Java EE application server.

IMS Universal JCA/JDBC driver with XA transaction support (imsudbJXA.rar)

This resource adapter provides a JDBC programming interface and both XATransaction and LocalTransaction support when deployed on any supported Java EE application server.

Restriction: XA transaction support is available only with type-4 connectivity.

For global or two-phase commit transaction processing, use the IMS Universal Database resource adapters with XA transaction support. For single-phase commit functionality, use either the IMS Universal Database resource adapters with XA transaction support or with local transaction support.

In order to provide for different transactional qualities of service for Java EE applications, it is possible to deploy two or more separate types of IMS Universal Database resource adapters into the same Java EE application server.

When carrying out multiple interactions with IMS databases using the IMS Universal Database resource adapter, you might want to group all actions together to ensure that they either all succeed or all fail. This can be done using container-managed or bean-managed transaction demarcation.

In *container-managed transactions*, all work performed in an EJB method invocation is part of one unit of work, and no explicit demarcation by the application is required. Transactional integrity is managed by the Java EE application server.

In *bean-managed transactions*, you must use the `javax.resource.cci.LocalTransaction` or `javax.transaction.UserTransaction` interface to programmatically demarcate units of work explicitly. Bean-managed transactions that use the `LocalTransaction` interface can group work performed only through the resource adapter; the `UserTransaction` interface allows all transactional resources within the application to be grouped. Use a bean-managed EJB if you need to have multiple units of works within the same EJB method invocation.

When using the type-2 IMS Universal Database resource adapter, if you specify a `driverType` connection property of 2, you can use `javax.resource.cci.LocalTransaction` for bean-managed transactions or the JDBC Connection interface. If you specify a `driverType` connection property of 2_CTX, you can use the `javax.transaction.UserTransaction` for application programs that issue explicit commit and rollback calls.

Software configurations supported by the IMS Universal Database resource adapter

The IMS Universal Database resource adapter has some requirements and restrictions on the supported software configurations.

The following table lists the supported software configurations.

Table 99. Supported software configurations for IMS Universal Database resource adapter

Version of the IMS Universal Database resource adapter	Supported version of IBM WebSphere Application Server	Non-IBM Java EE application servers
Version 14	- Version 8.5 or later - Version 8.5 or later for IBM WebSphere Liberty servers	Any generic Java EE 1.4 or later certified application servers ^{“1”} on page 573 ^{“2”} on page 573
Version 13	Version 7.0.1 or later	Any generic Java EE 1.4 or later certified application servers ^{“1”} on page 573 ^{“2”} on page 573
Version 12	Version 6.1 or later	Any generic Java EE 1.4 or later certified application servers ^{“1”} on page 573 ^{“2”} on page 573

1. For generic Java EE application servers, IBM support for the IMS Universal Database resource adapter is limited to specific features that meet one of the following criteria:
 - The problem can be recreated in the supported versions of WebSphere Application Server or WebSphere Application Server Liberty Profile.
 - The problem is confirmed to be caused by the IMS Universal Database resource adapter through diagnostic traces.
2. For generic Java EE application servers, the following features are not supported:
 - Two-phase (XA) commit processing.
 - Generic Java EE application servers on z/OS.
 - Other database resource adapters other than the IMS Universal Database resource adapter for local transactions with the JDBC programming interface (imsudbJLocal.rar).

Connecting to IMS with the IMS Universal Database resource adapter

The IMS Universal Database resource adapter provides connectivity to IMS databases from a Java EE-managed environment.

The Common Client Interface (CCI) provides `ConnectionFactory` and `Connection` interfaces to establish a connection with an Enterprise Information System (EIS). When using the CCI programming interface with the IMS Universal Database resource adapter, your Java application component looks up a `ConnectionFactory` instance using the Java Naming and Directory Interface (JNDI) and uses the `ConnectionFactory` instance to get a connection to an IMS database.

Similarly, when using the JDBC programming interface with the IMS Universal JCA/JDBC driver, your Java application component looks up a `DataSource` instance using JNDI, and uses the `DataSource` instance to obtain a `Connection` object

RRSLocalOption connectivity type

In addition to type-4 and type-2 connectivity, the `RRSLocalOption` connectivity type is supported by the IMS Universal Database resource adapter running on WebSphere Application Server for z/OS. With `RRSLocalOption` connectivity, applications using the IMS Universal Database resource adapter do not issue commit or rollback calls. Instead, transaction processing is managed by WebSphere Application Server for z/OS. Two-phase commit (XA) transaction processing is not supported with `RRSLocalOption` connectivity type.

Connecting using the IMS Universal Database resource adapter in a managed environment

In a managed (or three-tier) environment, your Java EE application interacts with a Java EE application server, such as WebSphere Application Server, and the IMS Universal Database resource adapter to communicate with an IMS database.

To configure and use a CCI `Connection` object in WebSphere Application Server to access an IMS database:

1. Deploy the IMS Universal Database resource adapter in WebSphere Application Server using the administrative console.
2. Create a connection factory for use with the IMS Universal Database resource adapter in WebSphere Application Server through the administrative console.
 - a) Specify a name for the connection factory and a Java Naming and Directory Interface (JNDI) name.
 - b) Set the following custom connection properties for the IMS Universal Database resource adapter:

DatastoreName

The name of the IMS data store to access.

When using type-4 connectivity, the **DatastoreName** property must match either the name of the data store defined to ODBM or be blank. The data store name is defined in the ODBM CSLDCxxx PROCLIB member using either the DATASTORE(NAME=*name*) or DATASTORE(NAME=*name*, ALIAS(NAME=*aliasname*)) parameter. If an alias is specified, you must specify the *aliasname* as the value of the **datastoreName** property. If the **DatastoreName** value is left blank (or not supplied), IMS Connect connects to any available instance of ODBM as it is assumed that data sharing is enabled between all datastores defined to ODBM.

When using type-2 connectivity, set the **DatastoreName** property to the IMS subsystem alias. This is not required to be set for the Java Dependent Region run time.

DatabaseName

The location of the database metadata representing the target IMS database.

The **DatabaseName** property can be specified in one of two ways, depending on whether the metadata is stored in the IMS catalog or as a static metadata class generated by the IMS Enterprise Suite Explorer for Development.

If your IMS system uses the IMS catalog, the **DatabaseName** property is the name of the PSB that your application uses to access the target IMS database.

If you are using the IMS Explorer for Development, the **databaseName** property is the fully qualified name of the Java metadata class generated by the IMS Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **DatabaseName** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

MetadataURL

The location of the database metadata representing the target IMS database.

This property is deprecated. Use **DatabaseName** instead.

The **MetadataURL** property is the fully qualified name of the Java metadata class generated by the IMS Enterprise Suite Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **MetadataURL** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

PortNumber

The TCP/IP server port number to be used to communicate with IMS Connect. The port number is defined using the DRDAPORT parameter on the ODACCESS statement in the IMS Connect configuration PROCLIB member. The default port number is 8888. Do not set this property when using type-2 connectivity.

DatastoreServer

The name or IP address of the data store server (IMS Connect). You can provide either the host name (for example, `dev123.svl.ibm.com`) or the IP address (for example, `192.166.0.2`). Do not set this property when using type-2 connectivity.

DriverType

The type of driver connectivity to use. The **DriverType** value must be "4" for type-4 connectivity or "2" for type-2 connectivity. If the driver is running on WebSphere Application Server for z/OS, you can also set the **DriverType** value to "2_CTX" for RRSLocalOption connectivity.

user

The user name for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

password

The password for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

allMetadata

Optional. When this property is set to `true`, the DatabaseMetadata interface returns information for all resources in the IMS catalog. When the property is set to `false`, the DatabaseMetadata interface returns information for the allocated PSB. The default value for this property is `false`.

signedCompare

Optional. When this property is set to `"true"`, special SSAs are generated to support ranged queries over signed data types. If the property is set to `"false"`, standard binary comparisons are performed based on the binary representation of the data type value. Setting the value to `"false"` can increase performance but might result in incorrect results. The default value for this property is `"true"`.

flattenTables

Optional. Indicates whether to produce a flattened view of the database tables. A value of `true` exposes the sub-elements of a STRUCT or an ARRAY as additional columns of the table. The default value is `false`.

IMS Explorer flattens the copybook structures when you import the copybook. Although the copybook itself remains unchanged in the IMS catalog, the information about the structure of each table is altered for that particular connection.

The the `flattenTables` property allows you to query the fields in complex structures directly. For more information about support for flattening complex structures, see [“Support for flattening complex structures” on page 566](#).

Restriction: The `flattenTables` connection property supports static arrays and structures only. Dynamic arrays are not altered.

initStatusGroup

Optional. When a connection is made and the PSB is allocated, this property will indicate that the driver should automatically issue an `INIT STATUS GROUPA` or `INIT STATUS GROUPB` if a value of 'A' or 'B' is provided. The default will not issue an `INIT STATUS GROUP` call.

sslKeyStoreType

Optional. Specifies the format of the file that contains cryptographic objects needed to establish a secure socket connection. The valid values are `"JKS"` and `"PKCS12"`. This value is only used when **`sslConnection`** is set to `"true"` and **`sslKeyStoreType`** is not specified. The **`sslKeyStoreType`** parameter defaults to `"JKS"`.

sslSecureSocketProtocol

Optional. Specifies the cryptographic communication protocol for the new connection. Specify a protocol that is supported by the server and provides the highest level of security. The valid values are `"SSL"`, `"SSLv3"`, `"TLSv1.1"`, and `"TLSv1.2"`. This value is only used when **`sslConnection`** is set to `"true"`. If **`sslConnection`** is set to `"true"` and **`sslSecureSocketProtocol`** is not specified, a default protocol will be determined at runtime by the JRE and the server.

t2OutputBufferSize

Optional. The size of the output buffer in bytes for the results from a SELECT operation for a type-2 connection.

The minimum value for `t2OutputBufferSize` is 500000. If any value less than 500000 is set, this property value will be adjusted to 500000. There is no maximum bound. The default value is 1280000.

treatInvalidDecimalAsNull

Optional. Indicates whether to interpret certain Decimal values that appear invalid in Java applications (such as PACKEDDECIMAL and ZONEDDECIMAL with invalid sign bits) as null. By default, this property is "false", and a conversion exception is thrown when the Java applications are processing invalid values.

removeInvalidCaseFields

Optional. If you use SELECT * to access a map case, the JDBC Universal Driver returns the set of all map case segment fields. When this property is set to true, the results contain only map case fields that satisfy a WHERE clause with a DEPENDINGON field condition. To learn more, see [“Using the removeInvalidCaseFields property” on page 610](#).

currentSchema

Optional. Specifies the default schema name that is used to qualify unqualified database objects in dynamically prepared SQL statements.

dbViewLocation

Optional. Specifies the fully qualified path to a databaseView metadata class. You can use this property to include a metadata class that is not located in your project path.

dpsbOnCommit

Optional. Set this property to **true** to deallocate the PSB when a commit occurs.

Recommendation: Do not set this property to **true** except in a managed environment with integrated connection pooling.

fetchSize

Optional. Gives the client a hint about the number of rows to get from the database when more rows are needed. The number specified for this property only affects data retrieved with the current connection. If the value specified is 0, all of the applicable rows are returned.

The default value for this property is 0 for both managed and unmanaged connections.

llField

Optional. The default value for this property is **true**. You can modify the LL field value to change the length of a variable length segment instance. Set this property to **false** to hide the LL field data from the column list in the result set.

maxRows

Optional. Specifies the maximum number of rows to return in a query result set. The default value is 0, which returns all of the applicable rows in the result set.

expandArrayResultSet

Optional. This property enables IMS Universal JDBC driver to create an **ArrayResultSet** with the array elements as the actual fields in the ResultSet instead of a Struct object containing those elements as the only entry into that ResultSet. To learn more, see [“Using the expandArrayResultSet property” on page 611](#).

ssaOptimization

Optional. When set to true, the SSA is optimized to a fully qualified key if subfields of that key are provided. To learn more, see [Optimizing SQL statements to promote partial key subfields in the WHERE Clause \(Application Programming\)](#).

traceFile

Optional. Specifies the name of the trace file for the connection.

traceFileAppend

Optional. If the specified trace file exists, setting this property to true specifies that the trace data for the new connection must be appended to the existing trace file instead of overwriting it.

This property is ignored if no value is specified for **traceFile**.

traceDirectory

Optional. Specifies the file system directory where the trace file is located. By default, this path is the directory where the application is executed.

This property is ignored if no value is specified for **traceFile**.

traceLevel

Optional. Specifies which traces are enabled for the connection. The valid values for this property are defined in the Java API documentation for the IMSDataSource class.

By default, all traces are disabled.

This property is ignored if no value is specified for **traceFile**.

Trace level	traceLevel package value	traceLevel constant field in IMSDataSource	traceLevel decimal value
All	com.ibm.ims.db.opendb.*	TRACE_ALL	-1
DL/I	com.ibm.ims.db.opendb.dli.*	TRACE_DLI	28
DRDA	com.ibm.ims.db.opendb.drda.*	TRACE_DRDA	1
JDBC	com.ibm.ims.db.opendb.jdbc.*	TRACE_JDBC	32
Java EE	com.ibm.ims.opendb.spi.* com.ibm.ims.db.opendb.cci.*	TRACE_JEE	192

3. To use the connection factory that is created in WebSphere Application Server in the previous step, use one of the following approaches:

- For Java EE 6 or later applications, you can use resource injection to get a handle to the connection factory object. Below is an example:

```
// When using imsudbLocal.rar or imsudbXA.rar the connection
// information is returned as a javax.resource.cci.ConnectionFactory object
@Resource(name="imsdblocal")
private javax.resource.cci.ConnectionFactory cf;

// When using imsudbJLocal.rar or imsudbJXA.rar the connection
// information is returned as a javax.sql.DataSource object
@Resource(name="imsdblocal")
private javax.sql.DataSource ds;
```

- For earlier Java EE applications or other applications, you can use a lookup facility to look up the connection factory object programmatically. Below is an example:

```
// When using imsudbLocal.rar or imsudbXA.rar the connection
// information is returned as a javax.resource.cci.ConnectionFactory object
Context initctx = new InitialContext();
javax.resource.cci.ConnectionFactory cf =
    (javax.resource.cci.ConnectionFactory)initctx.lookup
    ("java:comp/env/imsdblocal");

// When using imsudbJLocal.rar or imsudbJXA.rar the connection
// information is returned as a javax.sql.DataSource object
Context initctx = new InitialContext();
javax.resource.cci.ConnectionFactory cf =
    (javax.sql.DataSource)initctx.lookup
    ("java:comp/env/imsdblocal");
```

4. Invoke the `getConnection()` method to get the connection to an IMS database from the connection factory (cf) or data source (ds) object.

Follow one of the following methods:

- If the application does not need to override any connection properties that are set in the connection factory deployment, call the `getConnection()` method:

```
// When using imsudbLocal.rar or imsudbXA.rar
javax.resource.cci.Connection conn = cf.getConnection();

// When using imsudbJLocal.rar or imsudbJXA.rar
javax.sql.Connection conn = ds.getConnection();
```

- If the application needs to override certain connection properties such as *user*, *password*, *databaseName*, or *metadataURL*, provide an additional `IMSConnectionSpec` object with the override values. Take the new `IMSConnectionSpec` object as the argument when you call the `getConnection()` method. Here is an example of overriding user and the password:

```
// specify connection properties to override user and password
IMSConnectionSpec connSpec = new IMSConnectionSpec();
connSpec.setUser(yourUser);
connSpec.setPassword(yourPassword);

// When using imsudbLocal.rar or imsudbXA.rar
javax.resource.cci.Connection conn = ((com.ibm.ims.db.cci.
ConnectionFactoryImpl)cf).getConnection(connSpec);

// When using imsudbJLocal.rar or imsudbJXA.rar
javax.sql.Connection conn =
((com.ibm.ims.db.hybrid.IMSHybridDataSource)ds).getConnection(connSpec);
```

Note: Overriding the connection property values in this way does not alter the default values or require any modifications to the J2C connection factory parameters.

5. Use the connection that you established in the previous steps to access the IMS database.
 - For examples on accessing IMS data with the IMS Universal JCA/JDBC driver, see [“Accessing IMS data with the IMS Universal JCA/JDBC driver”](#) on page 592.
 - For more examples on accessing IMS data with other Java classes, see [“Accessing IMS data with the DLIInteractionSpec class”](#) on page 585 or [“Accessing IMS data with the SQLInteractionSpec class”](#) on page 589.
6. After your Java EE application is finished with the connection, close the connection with the `close()` method on the `Connection` interface:

```
conn.close();
```

Example code for connecting to an IMS database using the IMS Universal Database resource adapter in a managed environment

The following code sample shows the flow for connecting to an IMS database using the IMS Universal Database resource adapter from an Java EE application:

In this example, the JNDI facility is used to look up the `ConnectionFactory` object and override connection properties *user* and *password*:

```
//obtain the initial JNDI Naming context
Context initctx = new InitialContext();

//perform JNDI lookup to obtain the connection factory
javax.resource.cci.ConnectionFactory cf =
    (javax.resource.cci.ConnectionFactory)initctx.lookup("java:comp/env/imsdblocal");

//specify connection properties
IMSConnectionSpec connSpec = new IMSConnectionSpec();
connSpec.setUser("user");
connSpec.setPassword("password");

//create CCI connection
```

```
javax.resource.cci.Connection conn = ((com.ibm.ims.db.cci.  
ConnectionFactoryImpl)cf).getConnection(connSpec);
```

Related tasks

[“Configuring the IMS Universal drivers for SSL support” on page 668](#)

With type-4 connectivity, the IMS Universal drivers provide support for the Secure Sockets Layer (SSL) through the Java Secure Socket Extension (JSSE).

Related reference

[IMS Universal drivers support for the Common Client Interface \(Application Programming APIs\)](#)

Connecting using the IMS Universal JCA/JDBC driver in a managed environment

To access IMS databases using a JDBC programming interface in a managed (or three-tier) environment, you need to deploy the IMS Universal JCA/JDBC driver on your Java EE application server and configure the connection properties.

To configure and use the IMS Universal JCA/JDBC driver to access an IMS database:

1. Deploy the IMS Universal JCA/JDBC driver in WebSphere Application Server using the administrative console.
2. Create a connection factory for use with the IMS Universal JCA/JDBC driver in WebSphere Application Server through the administrative console.
 - a) Specify a name for the connection factory and a Java Naming and Directory Interface (JNDI) name. Set the connection factory interface as `javax.sql.DataSource`.
 - b) Set the following custom connection properties for the IMS Universal JCA/JDBC driver connection factory:

DatastoreName

The name of the IMS data store to access.

When using type-4 connectivity, the **DatastoreName** property must match either the name of the data store defined to ODBM or be blank. The data store name is defined in the ODBM CSLDCxxx PROCLIB member using either the `DATASTORE(NAME=name)` or `DATASTORE(NAME=name, ALIAS(NAME=aliasname))` parameter. If an alias is specified, you must specify the *aliasname* as the value of the **datastoreName** property. If the **DatastoreName** value is left blank (or not supplied), IMS Connect connects to any available instance of ODBM as it is assumed that data sharing is enabled between all datastores defined to ODBM.

When using type-2 connectivity, set the **DatastoreName** property to the IMS subsystem alias. This is not required to be set for the Java Dependent Region run time.

DatabaseName

The location of the database metadata representing the target IMS database.

The **DatabaseName** property can be specified in one of two ways, depending on whether the metadata is stored in the IMS catalog or as a static metadata class generated by the IMS Enterprise Suite Explorer for Development.

If your IMS system uses the IMS catalog, the **DatabaseName** property is the name of the PSB that your application uses to access the target IMS database.

If you are using the IMS Explorer for Development, the **databaseName** property is the fully qualified name of the Java metadata class generated by the IMS Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **DatabaseName** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

MetadataURL

The location of the database metadata representing the target IMS database.

This property is deprecated. Use **DatabaseName** instead.

The **MetadataURL** property is the fully qualified name of the Java metadata class generated by the IMS Enterprise Suite Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **MetadataURL** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

PortNumber

The TCP/IP server port number to be used to communicate with IMS Connect. The port number is defined using the DRDAPORT parameter on the ODACCESS statement in the IMS Connect configuration PROCLIB member. The default port number is 8888. Do not set this property when using type-2 connectivity.

DatastoreServer

The name or IP address of the data store server (IMS Connect). You can provide either the host name (for example, `dev123.svl.ibm.com`) or the IP address (for example, `192.166.0.2`). Do not set this property when using type-2 connectivity.

DriverType

The type of driver connectivity to use. The **DriverType** value must be "4" for type-4 connectivity or "2" for type-2 connectivity. If the driver is running on WebSphere Application Server for z/OS, you can also set the **DriverType** value to "2_CTX" for RRSLocalOption connectivity.

sslConnection

Optional. Indicates if this connection uses Secure Sockets Layer (SSL) for data encryption. Set this property to "true" to enable SSL, or to "false" otherwise. Do not set this property when using type-2 connectivity.

sslKeyStoreType

Optional. Specifies the format of the file that contains cryptographic objects needed to establish a secure socket connection. The valid values are "JKS" and "PKCS12". This value is only used when **sslConnection** is set to "true" and **sslKeyStoreType** is not specified. The **sslKeyStoreType** parameter defaults to "JKS".

sslSecureSocketProtocol

Optional. Specifies the cryptographic communication protocol for the new connection. Specify a protocol that is supported by the server and provides the highest level of security. The valid values are "SSL", "SSLv3", "TLSv1.1", and "TLSv1.2". This value is only used when **sslConnection** is set to "true". If **sslConnection** is set to "true" and **sslSecureSocketProtocol** is not specified, a default protocol will be determined at runtime by the JRE and the server.

sslTrustStoreLocation

Optional. Specifies the location of the cryptographic trust store file for the new connection. This value is only used when **sslConnection** is set to true.

sslTrustStorePassword

Optional. Specifies the password to access the cryptographic trust store file. This value is only used when **sslConnection** is set to true.

sslKeyStoreLocation

Optional. Specifies the location of the cryptographic key store file for the new connection. This value is only used when **sslConnection** is set to true.

sslKeyStorePassword

Optional. Specifies the password to access the cryptographic key store file. This value is only used when **sslConnection** is set to true.

initStatusGroup

Optional. When a connection is made and the PSB is allocated, this property will indicate that the driver should automatically issue an INIT STATUS GROUPA or INIT STATUS GROUPB if a value of 'A' or 'B' is provided. The default will not issue an INIT STATUS GROUP call.

loginTimeout

Optional. Specifies the number of seconds that the driver waits for a response from the server before timing out a connection initialization or server request. Set this property to a non-negative integer for the number of seconds. Set this property to 0 for an infinite timeout length. Do not set this property when using type-2 connectivity.

user

The user name for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

password

The password for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

signedCompare

Optional. When this property is set to "true", special SSAs are generated to support ranged queries over signed data types. If the property is set to "false", standard binary comparisons are performed based on the binary representation of the data type value. Setting the value to "false" can increase performance but might result in incorrect results. The default value for this property is "true".

flattenTables

Optional. Indicates whether to produce a flattened view of the database tables. A value of true exposes the sub-elements of a STRUCT or an ARRAY as additional columns of the table. The default value is false.

IMS Explorer flattens the copybook structures when you import the copybook. Although the copybook itself remains unchanged in the IMS catalog, the information about the structure of each table is altered for that particular connection.

The the flattenTables property allows you to query the fields in complex structures directly. For more information about support for flattening complex structures, see [“Support for flattening complex structures” on page 566](#).

Restriction: The flattenTables connection property supports static arrays and structures only. Dynamic arrays are not altered.

t2OutputBufferSize

Optional. The size of the output buffer in bytes for the results from a SELECT operation for a type-2 connection.

The minimum value for t2OutputBufferSize is 500000. If any value less than 500000 is set, this property value will be adjusted to 500000. There is no maximum bound. The default value is 1280000.

treatInvalidDecimalAsNull

Optional. Indicates whether to interpret certain Decimal values that appear invalid in Java applications (such as PACKEDDECIMAL and ZONEDDECIMAL with invalid sign bits) as null. By default, this property is "false", and a conversion exception is thrown when the Java applications are processing invalid values.

removeInvalidCaseFields

Optional. If you use `SELECT *` to access a map case, the JDBC Universal Driver returns the set of all map case segment fields. When this property is set to `true`, the results contain only map case fields that satisfy a `WHERE` clause with a `DEPENDINGON` field condition. To learn more, see [“Using the removeInvalidCaseFields property” on page 610](#).

currentSchema

Optional. Specifies the default schema name that is used to qualify unqualified database objects in dynamically prepared SQL statements.

dbViewLocation

Optional. Specifies the fully qualified path to a `databaseView` metadata class. You can use this property to include a metadata class that is not located in your project path.

dpsbOnCommit

Optional. Set this property to **true** to deallocate the PSB when a commit occurs.

Recommendation: Do not set this property to **true** except in a managed environment with integrated connection pooling.

fetchSize

Optional. Gives the client a hint about the number of rows to get from the database when more rows are needed. The number specified for this property only affects data retrieved with the current connection. If the value specified is 0, all of the applicable rows are returned.

The default value for this property is 0 for both managed and unmanaged connections.

llField

Optional. The default value for this property is **true**. You can modify the `LL` field value to change the length of a variable length segment instance. Set this property to **false** to hide the `LL` field data from the column list in the result set.

maxRows

Optional. Specifies the maximum number of rows to return in a query result set. The default value is 0, which returns all of the applicable rows in the result set.

expandArrayResultSet

Optional. This property enables IMS Universal JDBC driver to create an **ArrayResultSet** with the array elements as the actual fields in the `ResultSet` instead of a `Struct` object containing those elements as the only entry into that `ResultSet`. To learn more, see [“Using the expandArrayResultSet property” on page 611](#).

ssaOptimization

Optional. When set to `true`, the SSA is optimized to a fully qualified key if subfields of that key are provided. To learn more, see [Optimizing SQL statements to promote partial key subfields in the WHERE Clause \(Application Programming\)](#).

traceFile

Optional. Specifies the name of the trace file for the connection.

traceFileAppend

Optional. If the specified trace file exists, setting this property to `true` specifies that the trace data for the new connection must be appended to the existing trace file instead of overwriting it.

This property is ignored if no value is specified for **traceFile**.

traceDirectory

Optional. Specifies the file system directory where the trace file is located. By default, this path is the directory where the application is executed.

This property is ignored if no value is specified for **traceFile**.

traceLevel

Optional. Specifies which traces are enabled for the connection. The valid values for this property are defined in the Java API documentation for the `IMSDataSource` class.

By default, all traces are disabled.

This property is ignored if no value is specified for **traceFile**.

Trace level	traceLevel package value	traceLevel constant field in IMSDataSource	traceLevel demical value
All	com.ibm.ims.db.opendb.*	TRACE_ALL	-1
DL/I	com.ibm.ims.db.opendb.dli.*	TRACE_DLI	28
DRDA	com.ibm.ims.db.opendb.drda.*	TRACE_DRDA	1
JDBC	com.ibm.ims.db.opendb.jdbc.*	TRACE_JDBC	32
Java EE	com.ibm.ims.opendb.spi.* com.ibm.ims.db.opendb.cci.*	TRACE_JEE	192

3. In the deployment descriptor for your Java EE application, add a resource reference for the connection factory that was created in the previous step.

Set the name of the resource reference to the JNDI name of the connection factory and set the type to `javax.resource.cci.ConnectionFactory`.

4. In your Java EE application, create an initial JNDI naming context and get a `javax.sql.DataSource` instance for the IMS Universal JCA/JDBC hybrid driver using JNDI lookup.

The following code sample shows how to perform the JNDI lookup to obtain a `DataSource` instance, where the connection factory has the JNDI name "imsdblocal":

```
InitialContext ic = new InitialContext();
javax.sql.DataSource ds =
    (DataSource)ic.lookup("java:comp/env/imsdblocal");
```

5. Use the `getConnection` method on the `DataSource` instance to obtain a `java.sql.Connection` instance, as shown by the following code sample:

```
Connection con = ds.getConnection();
```

6. After your Java EE application has finished with the connection, close the connection using the `close` method on the `Connection` interface.

Example code for connecting to an IMS database using the IMS Universal JCA/JDBC driver

The following code sample shows the flow for connecting to an IMS database using the IMS Universal JCA/JDBC driver from an Java EE application:

```
//obtain the initial JNDI Naming context
InitialContext ic = new InitialContext();

//perform JNDI lookup to obtain the data source
javax.sql.DataSource ds =
    (DataSource)ic.lookup("java:comp/env/imsdblocal");

//specify connection properties
ds.setUser("myUserID");
ds.setPassword("myPassword");
```

```

props.put("sslConnection", "true");
props.put("loginTimeout", "10");

//create JDBC connection
java.sql.Connection con = ds.getConnection();

```

Related tasks

“Configuring the IMS Universal drivers for SSL support” on page 668

With type-4 connectivity, the IMS Universal drivers provide support for the Secure Sockets Layer (SSL) through the Java Secure Socket Extension (JSSE).

Related reference

[javax.sql.DataSource methods supported \(Application Programming APIs\)](#)

Sample EJB application using the IMS Universal Database resource adapter CCI programming interface

The following sample EJB bean demonstrates the basic programming flow for a JCA application using the IMS Universal Database resource adapter in a managed environment.

```

package client;

import java.sql.SQLException;
import javax.naming.InitialContext;
import javax.resource.ResourceException;
import javax.resource.cci.Connection;
import javax.resource.cci.ConnectionFactory;
import javax.resource.cci.Interaction;
import javax.resource.cci.ResultSet;
import javax.transaction.UserTransaction;
import com.ibm.ims.db.cci.SQLInteractionSpec;

/**
 * Bean implementation class for Enterprise Bean: StatefulBeanManaged
 */
public class BeanManagedSampleBean implements javax.ejb.SessionBean {

    private javax.ejb.SessionContext mySessionCtx;

    public void execute() throws Exception {
        InitialContext ic = new InitialContext();
        ConnectionFactory cf =
            (ConnectionFactory) ic.lookup("java:comp/env/MyMCF");
        Connection conn = null;
        UserTransaction ut = null;

        try {
            ut = this.mySessionCtx.getUserTransaction();
            ut.begin();

            conn = cf.getConnection();
            Interaction ix = conn.createInteraction();
            SQLInteractionSpec iSpec = new SQLInteractionSpec();

            // This query will return information for each person
            // in the phonebook with the last name WATSON
            iSpec.setSQL("SELECT * FROM " +
                "PCB01.PHONEBOOK WHERE LASTNAME='WATSON'");

            ResultSet rs = (ResultSet) ix.execute(iSpec, null);

            // Print out the first name of every person in the
            // phonebook with the last name WATSON
            while (rs.next()) {
                System.out.println(rs.getString("FIRSTNAME"));
            }

            rs.close();
            ix.close();
            ut.commit();
            conn.close();
        } catch (ResourceException e) {
            ut.rollback();
            conn.close();
        } catch (SQLException e) {

```

```

        ut.rollback();
        conn.close();
    }
}

/**
 * getSessionContext
 */
public javax.ejb.SessionContext getSessionContext() {
    return mySessionCtx;
}

/**
 * setSessionContext
 */
public void setSessionContext(javax.ejb.SessionContext ctx) {
    mySessionCtx = ctx;
}

/**
 * ejbCreate
 */
public void ejbCreate() throws javax.ejb.CreateException {
}

/**
 * ejbActivate
 */
public void ejbActivate() {
}

/**
 * ejbPassivate
 */
public void ejbPassivate() {
}

/**
 * ejbRemove
 */
public void ejbRemove() {
}
}

```

Accessing IMS data with the DLIIInteractionSpec class

Use the DLIIInteractionSpec class to retrieve, insert, update, and delete data from an IMS database using DL/I-like programming semantics with the IMS Universal Database resource adapter.

Before your application component can retrieve, insert, update, or delete data from an IMS database, you need to obtain a `javax.resource.cci.Connection` instance for the physical connection to the database.

To retrieve, insert, update, or delete data using the DLIIInteractionSpec class:

1. In your application component, create a new `javax.resource.cci.Interaction` instance using the `Connection.createInteraction` method.

For example, in the following code sample, *con* is a `javax.resource.cci.Connection` instance for an IMS database:

```
Interaction ix = con.createInteraction();
```

2. Create a new `com.ibm.ims.db.cci.DLIIInteractionSpec` instance.

```
DLIIInteractionSpec iSpec = new DLIIInteractionSpec();
```

3. Set the function to perform using the `DLIIInteractionSpec.setFunctionName` method, and specifying the function constant value listed in the table below as the input parameter.

Data operation to perform	setFunctionName value
Data retrieval	DLIIInteractionSpec.RETRIEVE
Data insertion	DLIIInteractionSpec.CREATE

Data operation to perform	setFunctionName value
Data update	DLIInteractionSpec.UPDATE
Data deletion	DLIInteractionSpec.DELETE

For example, the following code sample specifies a data retrieval operation:

```
iSpec.setFunctionName(DLIInteractionSpec.RETRIEVE);
```

- Set the PCB name using the `DLIInteractionSpec.setPCBName` method.


For example, the following code sample specifies the PCB to be used for this interaction as "PCB01":

```
iSpec.setPCBName("PCB01");
```

- Set the segment search argument (SSA) list using the `DLIInteractionSpec.setSSAList` method. The `setSSAList` method allows you to specify an SSA in a syntax similar to traditional DL/I.

- You can manually provide the SSA qualification statement as a string in the argument. The syntax is as follows:

Syntax for segment search argument qualification statement in the setSSAList method



segName

The name of the segment as defined in the Java metadata class generated by the IMS Enterprise Suite Explorer for Development.

cmdCode (optional)

All DL/I command codes except Q and C are supported.

fldName

The name of the field.

relOp

The SSA qualification statement's relational operator. Supported values are:

=

Equals

!=

Not equal

>

Greater than

>=

Greater than or equals

<

Less than

<=

Less than or equals

value

A string representation of the field value. If the value is character-based, the string has to be enclosed in quotation marks. If the value is numeric, it does not need to be enclosed in quotation marks. If the character-based value has quotation marks, use a single quote as an escape for the quote in the value. For example, if the value is "O'brian", you would enter it as "O'brian".

boolOp

Boolean operators for adding additional field-level qualifications. The supported Boolean operators are:

- logical AND (specified * or &)
- logical OR (specified + or |)
- independent AND (specified #)

The following code example shows how to set the segment search argument list to return the last patient admitted to all wards with more than five doctors and less than three nurses in hospital "ALEXANDRIA". The *L command means "last occurrence".

```
String ssalist =
"Hospital(HospName='ALEXANDRIA') Ward(Doccount>5 | Nurcount<3) Patient *L";
iSpec.setSSAList(ssalist);
```

- Instead of providing the string manually, you can use the `com.ibm.ims.db.cci.SSAListHelper` class to generate the string.

The following code example shows how to set the SSA qualification statement string using the `SSAListHelper`:

```
SSAListHelper sh = new SSAListHelper();
sh.addInitialQualification
("Hospital","HospName",SSAListHelper.EQUALS, "ALEXANDRIA");
sh.addInitialQualification("Ward","Doccount",
SSAListHelper.GREATER_THAN, 5);
sh.appendQualification("Ward",SSAListHelper.OR, "Nurcount",
SSAListHelper.LESS_THAN, 3);
sh.addCommandCode("Patient",SSAListHelper.CC_L);
iSpec.setSSA(sh.toString());
```

6. Create a `javax.resource.cci.RecordFactory` instance using the `ConnectionFactory.getRecordFactory` method.

For example, the following code sample creates a `ConnectionFactory` instance *rf*:

```
RecordFactory rf = cf.getRecordFactory();
```

This step is not needed for a DELETE operation.

7. Create a `javax.resource.cci.MappedRecord` instance using the `RecordFactory.getMappedRecord` method. For a RETRIEVE operation, pass the name of the record you want to create as an argument to this method. For a CREATE or UPDATE operation, the argument is the name of the segment to insert or update.

For example, in the following code sample for a RETRIEVE operation, *rf* is a `javax.resource.cci.RecordFactory` instance:

```
MappedRecord input = rf.createMappedRecord("myHospitalRecord");
```

This step is not needed for a DELETE operation because the `MappedRecord` is not used.

8. Specify the field to target for the data operation using the `MappedRecord.put` method. Pass the name of the field as the first argument to this method. For a CREATE or UPDATE operation, the `MappedRecord` is used to specify the field to insert or update as well as its values. Pass in the value of the field as the second argument. For a RETRIEVE operation, the `MappedRecord` is used to specify the field that the application is interested in retrieving as a result of the call, and the second argument in the `put` method call is ignored (you can pass in a null). If you do not specify any fields, the RETRIEVE operation will return all the fields of the leaf segment for that record, along with all the fields in segments which have SSAs specified with a *D command. For CREATE, UPDATE, and RETRIEVE operations, you can specify multiple fields by making multiple `MappedRecord.put` method calls.

For example, in the following code sample, *input* is a `javax.resource.cci.MappedRecord` instance and "HospCode" is the name of the field we want to retrieve:

```
input.put("HospCode", null);
```

This step is not needed for a DELETE operation as the `MappedRecord` is not used.

9. Execute the query by calling the `Interaction.execute` method. Pass the `DLIInteractionSpec` object and the `MappedRecord` object as arguments.

If the query is successful, the method returns a `Record` object with the query results. You can cast the `Record` instance to `javax.resource.cci.ResultSet` and process the results as tabular data in your application component. For example, in the following code sample, *results* is a `javax.resource.cci.ResultSet` instance, *ix* is a `javax.resource.cci.Interaction` instance, *iSpec* is a `com.ibm.ims.db.cci.DLIInteractionSpec` instance, and *input* is a `javax.resource.cci.MappedRecord` instance:

```
results = (ResultSet)ix.execute(iSpec, input);
```

Example code for IMS data operations using the DLIInteraction interface

The following complete code example shows how to use the `DLIInteraction` interface to retrieve fields from a WARD segment.

```
package client;

import java.sql.SQLException;
import javax.naming.InitialContext;
import javax.resource.ResourceException;
import javax.resource.cci.Connection;
import javax.resource.cci.ConnectionFactory;
import javax.resource.cci.Interaction;
import javax.resource.cci.ResultSet;
import javax.transaction.UserTransaction;
import com.ibm.ims.db.cci.DLIInteractionSpec;

/**
 * Bean implementation class for Enterprise Bean: StatefulBeanManaged
 */
public class BeanManagedSampleDLIBean implements javax.ejb.SessionBean {

    private javax.ejb.SessionContext mySessionCtx;

    public void execute() throws Exception {
        InitialContext ic = new InitialContext();
        ConnectionFactory cf =
            (ConnectionFactory) ic.lookup("java:comp/env/MyMCF");
        Connection conn = null;
        UserTransaction ut = null;

        try {
            ut = this.mySessionCtx.getUserTransaction();
            ut.begin();

            conn = cf.getConnection();
            Interaction ix = conn.createInteraction();

            DLIInteractionSpec iSpec = new DLIInteractionSpec();
            iSpec.setFunctionName("RETRIEVE");
            iSpec.setPCBName("PCB09");

            // This query will return the WARDNAME, PATCOUNT, DOCCOUNT,
            // and NURCOUNT fields for all WARDS with WARNNO = 51
            iSpec.setSSAList("WARD (WARDNO = '51')");

            // Create RecordFactory
            RecordFactory rf = cf.getRecordFactory();

            // Create Record
            MappedRecord input = rf.createMappedRecord("WARD");
            // Specify the fields to retrieve
            input.put("WARDNAME", null);
            input.put("PATCOUNT", null);
            input.put("DOCCOUNT", null);
            input.put("NURCOUNT", null);

            ResultSet results = (ResultSet) ix.execute(iSpec, input);
            while (results.next()) {
                System.out.println(results.getString("WARDNAME"));
                System.out.println(results.getString("PATCOUNT"));
                System.out.println(results.getString("DOCCOUNT"));
                System.out.println(results.getString("NURCOUNT"));
            }
        } catch (Exception e) {
            // Handle exception
        }
    }
}
```



```

        }

        rs.close();
        ix.close();
        ut.commit();
        conn.close();
    } catch (ResourceException e) {
        ut.rollback();
        conn.close();
    } catch (SQLException e) {
        ut.rollback();
        conn.close();
    }
}

/**
 * getSessionContext
 */
public javax.ejb.SessionContext getSessionContext() {
    return mySessionCtx;
}

/**
 * setSessionContext
 */
public void setSessionContext(javax.ejb.SessionContext ctx) {
    mySessionCtx = ctx;
}

/**
 * ejbCreate
 */
public void ejbCreate() throws javax.ejb.CreateException {
}

/**
 * ejbActivate
 */
public void ejbActivate() {
}

/**
 * ejbPassivate
 */
public void ejbPassivate() {
}

/**
 * ejbRemove
 */
public void ejbRemove() {
}
}

```

Related concepts

[“Connecting to IMS with the IMS Universal Database resource adapter” on page 573](#)

The IMS Universal Database resource adapter provides connectivity to IMS databases from a Java EE-managed environment.

Related reference

[IMS Universal drivers support for the Common Client Interface \(Application Programming APIs\)](#)

Accessing IMS data with the SQLInteractionSpec class

Use the SQLInteractionSpec class to retrieve, insert, update, and delete data from an IMS database using SQL queries with the IMS Universal Database resource adapter. The IMS Universal Database resource adapter supports the same SQL statement syntax and usage as the IMS Universal JDBC driver and has the same restrictions.

Before your application component can retrieve, insert, update, or delete data from an IMS database, you need to obtain a `javax.resource.cci.Connection` instance for the physical connection to the database.

To retrieve, insert, update, or delete data using the SQLInteractionSpec class:

1. In your application component, create a new `javax.resource.cci.Interaction` instance using the `Connection.createInteraction` method.

For example, in the following code sample, *con* is a `javax.resource.cci.Connection` instance for an IMS database:

```
Interaction ix = con.createInteraction();
```

2. Create a new `com.ibm.ims.db.cci.SQLInteractionSpec` instance.

```
SQLInteractionSpec iSpec = new SQLInteractionSpec();
```

3. Set the SQL query string using the `SQLInteractionSpec.setSQL` method.

In the query, you can specify the qualification column values in the WHERE clause with a ? parameter marker, meaning that the values will be provided later (similar to a `PreparedStatement` in JDBC).

- The following example shows how to specify the SELECT statement without using parameter markers, where *iSpec* is an instance of `SQLInteractionSpec`:

```
iSpec.setQuery("SELECT PATIENT.PATNAME, ILLNESS.ILLNAME "+
"FROM pcb01.HOSPITAL, pcb01.PATIENT, pcb01.ILLNESS " +
"WHERE HOSPITAL.HOSPNAME='SANTA TERESA'");
```

- The following example shows how to perform the SELECT statement with parameter markers, where *iSpec* is an instance of `SQLInteractionSpec`:

```
iSpec.setQuery("SELECT PATIENT.PatName, WARD.WardName "+
"FROM pcb01.HOSPITAL, pcb01.PATIENT, pcb01.WARD " +
"WHERE HOSPITAL.HospName=? AND WARD.DocCount>?");
```

4. Create a `javax.resource.cci.RecordFactory` instance using the `ConnectionFactory.getRecordFactory` method.

Creating a `RecordFactory` is not needed for SQL queries without parameter markers.

5. Create a `javax.resource.cci.IndexedRecord` instance using the `RecordFactory.getIndexedRecord` method. Pass the name of the record you want to create as an argument to this method.

For example, in the following code sample, *rf* is a `javax.resource.cci.RecordFactory` instance:

```
IndexedRecord input = rf.createIndexedRecord("myPatientRecord");
```

Creating a `IndexedRecord` is not needed for SQL queries without parameter markers.

6. If your query string uses parameter markers, use the `IndexedRecord.add` method to qualify the WHERE clause.

For example, using the same query string with parameter markers as in step 5, where *input* is an instance of `IndexedRecord`:

```
input.add(1, "Santa Teresa"); //HospName value is "Santa Teresa"
input.add(2, 5);              //DocCount value is greater than 5
```

7. Execute the query by calling the `Interaction.execute` method. Pass the `SQLInteractionSpec` object and the `IndexedRecord` object as arguments.

If your SQL query does not use parameter markers, the second argument in the `execute` method call is ignored (you can pass in a null). If the query is successful, the method returns a `Record` object with the query results. You can cast the `Record` instance to `javax.resource.cci.ResultSet` and process the results as tabular data in your application component. For example, in the following code sample, *results* is a `javax.resource.cci.ResultSet` instance, *ix* is a `javax.resource.cci.Interaction` instance, *iSpec* is a `com.ibm.ims.db.cci.SQLInteractionSpec` instance, and *input* is a `javax.resource.cci.IndexedRecord` instance:

```
results = (ResultSet)ix.execute(iSpec, input);
```

Example code for IMS data operations using the SQLInteractionSpec class

The following code example shows how to use the SQLInteractionSpec class to retrieve patient names from PATIENT records.

```
package client;

import java.sql.SQLException;
import javax.naming.InitialContext;
import javax.resource.ResourceException;
import javax.resource.cci.Connection;
import javax.resource.cci.ConnectionFactory;
import javax.resource.cci.Interaction;
import javax.resource.cci.ResultSet;
import javax.transaction.UserTransaction;
import com.ibm.ims.db.cci.SQLInteractionSpec;

/**
 * Bean implementation class for Enterprise Bean: StatefulBeanManaged
 */
public class BeanManagedSampleSQLBean implements javax.ejb.SessionBean {

    private javax.ejb.SessionContext mySessionCtx;

    public void execute() throws Exception {
        InitialContext ic = new InitialContext();
        ConnectionFactory cf =
            (ConnectionFactory) ic.lookup("java:comp/env/MyMCF");
        Connection conn = null;
        UserTransaction ut = null;

        try {
            ut = this.mySessionCtx.getUserTransaction();
            ut.begin();

            conn = cf.getConnection();
            Interaction ix = conn.createInteraction();
            SQLInteractionSpec iSpec = new SQLInteractionSpec();

            // This query will return the WARDNAME, PATCOUNT, DOCCOUNT,
            // and NURCOUNT fields for the WARD with WARDNO = 51
            iSpec.setSQL("SELECT WARDNAME, PATCOUNT, DOCCOUNT, " +
                "NURCOUNT FROM PCB09.WARD WHERE WARDNO='51'");

            ResultSet rs = (ResultSet) ix.execute(iSpec, null);

            while (rs.next()) {
                System.out.println(rs.getString("WARDNAME"));
                System.out.println(rs.getString("PATCOUNT"));
                System.out.println(rs.getString("DOCCOUNT"));
                System.out.println(rs.getString("NURCOUNT"));
            }

            rs.close();
            ix.close();
            ut.commit();
            conn.close();
        } catch (ResourceException e) {
            ut.rollback();
            conn.close();
        } catch (SQLException e) {
            ut.rollback();
            conn.close();
        }
    }

    /**
     * getSessionContext
     */
    public javax.ejb.SessionContext getSessionContext() {
        return mySessionCtx;
    }

    /**
     * setSessionContext
     */
    public void setSessionContext(javax.ejb.SessionContext ctx) {
        mySessionCtx = ctx;
    }
}
```

```

/**
 * ejbCreate
 */
public void ejbCreate() throws javax.ejb.CreateException {
}

/**
 * ejbActivate
 */
public void ejbActivate() {
}

/**
 * ejbPassivate
 */
public void ejbPassivate() {
}

/**
 * ejbRemove
 */
public void ejbRemove() {
}
}

```

Related concepts

[“Connecting to IMS with the IMS Universal Database resource adapter” on page 573](#)

The IMS Universal Database resource adapter provides connectivity to IMS databases from a Java EE-managed environment.

Related reference

[IMS Universal drivers support for the Common Client Interface \(Application Programming APIs\)](#)

[“SQL statement usage with the IMS Universal JDBC driver” on page 619](#)

The following usage rules apply to SQL statements passed to IMS with the IMS Universal JDBC driver.

Accessing IMS data with the IMS Universal JCA/JDBC driver

Use the IMS Universal JCA/JDBC driver if you require full use of the IMS Universal JDBC driver within a Java EE runtime environment.

Before your Java EE application component can retrieve, insert, update, or delete data from an IMS database, you need to obtain a `java.sql.Connection` instance for the physical connection to the database.

In your Java EE application component, code the application logic for the data operations you want to perform in the same way as for a JDBC application. The IMS Universal JCA/JDBC driver has the same SQL statement syntax support and usage restrictions as the IMS Universal JDBC driver.

Example EJB application using the IMS Universal JCA/JDBC driver

The following code sample shows a bean-managed EJB application that connects to an IMS database, retrieves a list of patient names using a SQL SELECT query, and modifies the patient information using a SQL UPDATE query.

```

package client;
import java.sql.SQLException;
import javax.naming.InitialContext;
import javax.sql.DataSource;
import java.sql.Connection;
import java.sql.Statement;
import java.sql.ResultSet;

import javax.transaction.UserTransaction;
import com.ibm.ims.db.cci.SQLInteractionSpec;

/**
 * Bean implementation class for Enterprise Bean: StatefulBeanManaged
 */
public class JDBCBeanManagedSampleSQLBean {
    private javax.ejb.SessionContext mySessionCtx;

```

```

public void execute() throws Exception {
    InitialContext ic = new InitialContext();
    DataSource ds =
        (DataSource) ic.lookup("java:comp/env/MyMCF");
    Connection conn = null;
    UserTransaction ut = null;

    try {
        ut = this.mySessionCtx.getUserTransaction();
        ut.begin();

        conn = ds.getConnection();

        Statement st = conn.createStatement();

        // List all of the patient names in the
        // SURG ward in the ALEXANDRIA hospital
        ResultSet rs = st.executeQuery("SELECT patname from " +
            "pcb01.hospital, ward, patient " +
            "where hospital.hospname = 'ALEXANDRIA' " +
            "and ward.wardname = 'SURG'");
        while (rs.next()) {
            System.out.println(rs.getString("patname"));
        }

        // Update the name of the patient with patient
        // number 0222 in ward 04 in the hospital
        // with code R1210010000A
        int updatedRecords = st.executeUpdate("UPDATE PCB01.PATIENT " +
            "SET PATNAME='UPDATED NAME' WHERE PATNUM='0222' " +
            "AND HOSPITAL_HOSPCODE='R1210010000A' AND WARD_WARDNO='04'");
        System.out.println("Updated " + updatedRecords + " Record(s)");

        rs.close();
        ut.commit();
        conn.close();
    } catch (SQLException e) {
        e.printStackTrace();
        ut.rollback();
        conn.close();
    }
}

/**
 * getSessionContext
 */
public javax.ejb.SessionContext getSessionContext() {
    return mySessionCtx;
}

/**
 * setSessionContext
 */
public void setSessionContext(javax.ejb.SessionContext ctx) {
    mySessionCtx = ctx;
}

/**
 * ejbCreate
 */
public void ejbCreate() throws javax.ejb.CreateException {
}

/**
 * ejbActivate
 */
public void ejbActivate() {
}

/**
 * ejbPassivate
 */
public void ejbPassivate() {
}

/**
 * ejbRemove
 */
public void ejbRemove() {
}

```

```
}  
}
```

Related reference

[“SQL statement usage with the IMS Universal JDBC driver” on page 619](#)

The following usage rules apply to SQL statements passed to IMS with the IMS Universal JDBC driver.

Programming with the IMS Universal JDBC driver

IMS provides a Java Database Connectivity (JDBC) driver for SQL-based database connectivity to access IMS databases over TCP/IP with the IMS Universal JDBC driver that is included in the IMS Universal drivers. The IMS Universal JDBC driver is based on the JDBC 4.0 standard.

JDBC is an application programming interface (API) that Java applications use to access relational databases or tabular data sources. The JDBC API is the industry standard for database-independent connectivity between the Java programming language and any database that has implemented the JDBC interface. The client uses the interface to query and update data in a database. It is the responsibility of the JDBC driver itself to implement the underlying (specific) access protocol for the specific database the driver is implemented for. Drivers are client-side adapters (they are installed in the client machine, not in the server) that convert requests from Java programs to a protocol that the database management system (DBMS) can understand.

IMS support for JDBC lets you write Java applications that can issue dynamic SQL calls to access IMS data and process the result set that is returned in tabular format. The IMS Universal JDBC driver is designed to support a subset of the SQL syntax with functionality that is limited to what the IMS database management system can process natively. Its DBMS-centric design allows the IMS Universal JDBC driver to fully leverage the high performance capabilities of IMS. The IMS Universal JDBC driver also provides aggregate function support, and ORDER BY and GROUP BY support.

Preparing to write a Java application with the IMS Universal drivers

Java application programs that use the IMS Universal drivers require the Java Development Kit (JDK) 7.0. Java programs that run in JMP and JBP regions require JDK 7.0 or later. Java application programs that use the IMS Universal drivers must have access to database metadata in order to interact with IMS databases. This metadata can either be accessed directly in the IMS catalog database or it can be generated as a Java metadata class with the IMS Enterprise Suite Explorer for Development.

Related reference

[IBM Java development kits on the developerWorks website](#)

Supported drivers for JDBC

The IMS Universal JDBC driver supports the type-2 and type-4 JDBC architectures.

The table below lists the IMS support available for the four types of JDBC driver architectures:

Table 100. JDBC driver architectures supported by IMS

JDBC driver architecture	Description	IMS support
Type-1	Drivers that implement the JDBC API as a mapping to another data access API, such as Open Database Connectivity (ODBC). Drivers of this type are generally dependent on a native library, which limits their portability.	IMS does not support a type-1 driver.

Table 100. JDBC driver architectures supported by IMS (continued)

JDBC driver architecture	Description	IMS support
Type-2	Drivers that are written partly in the Java programming language and partly in native code. The drivers use a native client library specific to the data source to which they connect. Because of the native code, their portability is limited. Java programs with type 2 JDBC connectivity can run on the same z/OS system or zSeries logical partition (LPAR) as the target IMS subsystem.	Use the IMS Universal JDBC driver with type-2 connectivity to access IMS from WebSphere Application Server for z/OS, IMS Java Dependent Regions (JDRs), and CICS.
Type-3	Drivers that use a pure Java client and communicate with a server using a database-independent protocol. The server then communicates the client's requests to the data source.	IMS does not support a type-3 driver.
Type-4	Drivers that are pure Java and implement the network protocol for a specific data source. The client connects directly to the data source.	Use the IMS Universal JDBC driver with type-4 connectivity to access the IMS subsystem via a TCP/IP network connection.

Connecting to IMS using the IMS Universal JDBC driver

You must first establish a connection to an IMS database before you can start sending queries and receiving results in your JDBC program.

Connecting to an IMS database using the JDBC DataSource interface

Using the DataSource interface is the preferred way to connect to IMS from your IMS Universal JDBC driver application.

Use the following procedure to create and use a DataSource interface in your IMS Universal JDBC driver application.

1. Create an instance of the `com.ibm.ims.jdbc.IMSDataSource` class.

This class is the IMS Universal JDBC driver implementation of the DataSource interface.

2. Set the following connection properties of the DataSource instance.

DatastoreName

The name of the IMS data store to access.

When using type-4 connectivity, the **DatastoreName** property must match either the name of the data store defined to ODBM or be blank. The data store name is defined in the ODBM CSLDCxxx PROCLIB member using either the `DATASTORE(NAME=name)` or `DATASTORE(NAME=name, ALIAS(NAME=aliasname))` parameter. If an alias is specified, you must specify the *aliasname* as the value of the **datastoreName** property. If the **DatastoreName** value is left blank (or not supplied), IMS Connect connects to any available instance of ODBM as it is assumed that data sharing is enabled between all datastores defined to ODBM.

When using type-2 connectivity, set the **DatastoreName** property to the IMS subsystem alias. This is not required to be set for the Java Dependent Region run time.

DatabaseName

The location of the database metadata representing the target IMS database.

The **DatabaseName** property can be specified in one of two ways, depending on whether the metadata is stored in the IMS catalog or as a static metadata class generated by the IMS Enterprise Suite Explorer for Development.

If your IMS system uses the IMS catalog, the **DatabaseName** property is the name of the PSB that your application uses to access the target IMS database.

If you are using the IMS Explorer for Development, the **databaseName** property is the fully qualified name of the Java metadata class generated by the IMS Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **DatabaseName** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

MetadataURL

The location of the database metadata representing the target IMS database.

This property is deprecated. Use **DatabaseName** instead.

The **MetadataURL** property is the fully qualified name of the Java metadata class generated by the IMS Enterprise Suite Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **MetadataURL** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

PortNumber

The TCP/IP server port number to be used to communicate with IMS Connect. The port number is defined using the DRDAPORT parameter on the ODACCESS statement in the IMS Connect configuration PROCLIB member. The default port number is 8888. Do not set this property when using type-2 connectivity.

DatastoreServer

The name or IP address of the data store server (IMS Connect). You can provide either the host name (for example, `dev123.svl.ibm.com`) or the IP address (for example, `192.166.0.2`). Do not set this property when using type-2 connectivity.

DriverType

The type of driver connectivity to use (value must be `IMSDatasource.DRIVER_TYPE_4` for type-4 connectivity or `IMSDatasource.DRIVER_TYPE_2` for type-2 connectivity).

user

The user name for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

password

The password for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

allMetadata

Optional. When this property is set to `true`, the DatabaseMetadata interface returns information for all resources in the IMS catalog. When the property is set to `false`, the DatabaseMetadata interface returns information for the allocated PSB. The default value for this property is `false`.

initStatusGroup

Optional. When a connection is made and the PSB is allocated, this property will indicate that the driver should automatically issue an `INIT STATUS GROUPA` or `INIT STATUS GROUPB` if a value of 'A' or 'B' is provided. The default will not issue an `INIT STATUS GROUP` call.

sslConnection

Optional. Indicates if this connection uses Secure Sockets Layer (SSL) for data encryption. Set this property to "true" to enable SSL, or to "false" otherwise. Do not set this property when using type-2 connectivity.

sslKeyStoreType

Optional. Specifies the format of the file that contains cryptographic objects needed to establish a secure socket connection. The valid values are "JKS" and "PKCS12". This value is only used when **sslConnection** is set to "true" and **sslKeyStoreType** is not specified. The **sslKeyStoreType** parameter defaults to "JKS".

sslSecureSocketProtocol

Optional. Specifies the cryptographic communication protocol for the new connection. Specify a protocol that is supported by the server and provides the highest level of security. The valid values are "SSL", "SSLv3", "TLSv1.1", and "TLSv1.2". This value is only used when **sslConnection** is set to "true". If **sslConnection** is set to "true" and **sslSecureSocketProtocol** is not specified, a default protocol will be determined at runtime by the JRE and the server.

sslTrustStoreLocation

Optional. Specifies the location of the cryptographic trust store file for the new connection. This value is only used when **sslConnection** is set to true.

sslTrustStorePassword

Optional. Specifies the password to access the cryptographic trust store file. This value is only used when **sslConnection** is set to true.

sslKeyStoreLocation

Optional. Specifies the location of the cryptographic key store file for the new connection. This value is only used when **sslConnection** is set to true.

sslKeyStorePassword

Optional. Specifies the password to access the cryptographic key store file. This value is only used when **sslConnection** is set to true.

loginTimeout

Optional. Specifies the number of seconds that the driver waits for a response from the server before timing out a connection initialization or server request. Set this property to a non-negative integer for the number of seconds. Set this property to 0 for an infinite timeout length. Do not set this property when using type-2 connectivity.

signedCompare

Optional. When this property is set to "true", special SSAs are generated to support ranged queries over signed data types. If the property is set to "false", standard binary comparisons are performed based on the binary representation of the data type value. Setting the value to "false" can increase performance but might result in incorrect results. The default value for this property is "true".

flattenTables

Optional. Indicates whether to produce a flattened view of the database tables. A value of true exposes the sub-elements of a STRUCT or an ARRAY as additional columns of the table. The default value is false.

IMS Explorer flattens the copybook structures when you import the copybook. Although the copybook itself remains unchanged in the IMS catalog, the information about the structure of each table is altered for that particular connection.

The the flattenTables property allows you to query the fields in complex structures directly. For more information about support for flattening complex structures, see [“Support for flattening complex structures” on page 566](#).

Restriction: The flattenTables connection property supports static arrays and structures only. Dynamic arrays are not altered.

t2OutputBufferSize

Optional. The size of the output buffer in bytes for the results from a SELECT operation for a type-2 connection.

The minimum value for t2OutputBufferSize is 500000. If any value less than 500000 is set, this property value will be adjusted to 500000. There is no maximum bound. The default value is 1280000.

treatInvalidDecimalAsNull

Optional. Indicates whether to interpret certain Decimal values that appear invalid in Java applications (such as PACKEDDECIMAL and ZONEDDECIMAL with invalid sign bits) as null. By default, this property is "false", and a conversion exception is thrown when the Java applications are processing invalid values.

removeInvalidCaseFields

Optional. If you use SELECT * to access a map case, the JDBC Universal Driver returns the set of all map case segment fields. When this property is set to true, the results contain only map case fields that satisfy a WHERE clause with a DEPENDINGON field condition. To learn more, see [“Using the removeInvalidCaseFields property” on page 610](#).

3. Optional: Set additional connection properties with the setProperties method of the IMSDataSource object.

currentSchema

Optional. Specifies the default schema name that is used to qualify unqualified database objects in dynamically prepared SQL statements.

dbViewLocation

Optional. Specifies the fully qualified path to a databaseView metadata class. You can use this property to include a metadata class that is not located in your project path.

dpsbOnCommit

Optional. Set this property to **true** to deallocate the PSB when a commit occurs.

Recommendation: Do not set this property to **true** except in a managed environment with integrated connection pooling.

fetchSize

Optional. Gives the client a hint about the number of rows to get from the database when more rows are needed. The number specified for this property only affects data retrieved with the current connection. If the value specified is 0, all of the applicable rows are returned.

The default value for this property is 0 for both managed and unmanaged connections.

llField

Optional. The default value for this property is **true**. You can modify the LL field value to change the length of a variable length segment instance. Set this property to **false** to hide the LL field data from the column list in the result set.

maxRows

Optional. Specifies the maximum number of rows to return in a query result set. The default value is 0, which returns all of the applicable rows in the result set.

expandArrayResultSet

Optional. This property enables IMS Universal JDBC driver to create an **ArrayResultSet** with the array elements as the actual fields in the ResultSet instead of a Struct object containing those elements as the only entry into that ResultSet. To learn more, see [“Using the expandArrayResultSet property”](#) on page 611.

ssaOptimization

Optional. When set to true, the SSA is optimized to a fully qualified key if subfields of that key are provided. To learn more, see [Optimizing SQL statements to promote partial key subfields in the WHERE Clause \(Application Programming\)](#).

traceFile

Optional. Specifies the name of the trace file for the connection.

traceFileAppend

Optional. If the specified trace file exists, setting this property to true specifies that the trace data for the new connection must be appended to the existing trace file instead of overwriting it.

This property is ignored if no value is specified for **traceFile**.

traceDirectory

Optional. Specifies the file system directory where the trace file is located. By default, this path is the directory where the application is executed.

This property is ignored if no value is specified for **traceFile**.

traceLevel

Optional. Specifies which traces are enabled for the connection. The valid values for this property are defined in the Java API documentation for the `IMSDataSource` class.

By default, all traces are disabled.

This property is ignored if no value is specified for **traceFile**.

Trace level	traceLevel package value	traceLevel constant field in IMSDataSource	traceLevel demical value
All	com.ibm.ims.db.opendb.*	TRACE_ALL	-1
DL/I	com.ibm.ims.db.opendb.dli.*	TRACE_DLI	28
DRDA	com.ibm.ims.db.opendb.drda.*	TRACE_DRDA	1
JDBC	com.ibm.ims.db.opendb.jdbc.*	TRACE_JDBC	32
Java EE	com.ibm.ims.opendb.spi.* com.ibm.ims.db.opendb.cci.*	TRACE_JEE	192

4. Establish a connection to the data source by calling the `getConnection` method on the `DataSource` object.
5. After your application has finished with the connection, close the connection using the `close` method on the `Connection` interface.

The following code example shows how to create a type-4 connection to an IMS database from your IMS Universal JDBC driver application using the `DataSource` interface:

```
import java.sql.*;
import javax.sql.*;
import com.ibm.ims.jdbc.*;

Connection conn = null;
```

```
// Create an instance of DataSource
IMSDataSource ds = new com.ibm.ims.jdbc.IMSDataSource();

// Set the URL of the fully qualified name of the Java metadata class
ds.setDatabaseName("class://BMP255.BMP255DatabaseView");

// Set the data store name
ds.setDatastoreName("IMS1");

// Set the data store server
ds.setDatastoreServer("ecdev47.svl.ibm.com");

// Set the port number
ds.setPortNumber(5555);

// Set the JDBC connectivity driver typ
ds.setDriverType(IMSDataSource.DRIVER_TYPE_4);

// Enable SSL for connection
ds.setSSLConnection("true");

// Set timeout for connection
ds.setLoginTimeout("10");

// Set user ID for connection
ds.setUser("myUserID");

// Set password for connection
ds.setPassword("myPassword");

// Create JDBC connection
conn = ds.getConnection();
```

Alternatively, you can set all of the connection properties as key value pairs in a `Properties` object and then set them simultaneously with the `IMSDataSource.setProperties` method:

```
Properties props = new Properties();
props.put( "user", "myUserID" );
props.put( "password", "MyPassword" );

IMSDataSource ds = new com.ibm.ims.jdbc.IMSDataSource();
ds.setProperties(props);
```

The typical usage of a `DataSource` object is for your system administrator to create and manage it separately. The program that creates and manages a `DataSource` object also uses the Java Naming and Directory Interface (JNDI) to assign a logical name to the `DataSource` object. The JDBC application that uses the `DataSource` object can then refer to the object by its logical name, and does not need any information about the underlying data source. In addition, your system administrator can modify the data source attributes, and you do not need to change your application program.

Recommendation: For maximum portability, use only the `DataSource` interface to obtain connections.

To obtain a connection using a `DataSource` object, given that the system administrator has already created the object and assigned a logical name to it:

1. From your system administrator, obtain the logical name of the data source to which you need to connect.
2. Create a `Context` object to use in the next step. The `Context` interface is part of the Java Naming and Directory Interface (JNDI), not JDBC.
3. In your application program, use JNDI to get the `DataSource` object that is associated with the logical data source name.
4. Use the `getConnection` method on the `DataSource` instance to obtain the connection.

The following code shows an example of the code that you need in your application program to obtain a connection using a `DataSource` object. In this example, the logical name of the data source that you need to connect to is `"jdbc/sampledb"`.

```
import java.sql.*;
import javax.naming.*;
import javax.sql.*;
...
```

```
Context ctx = new InitialContext();
DataSource ds = (DataSource)ctx.lookup("jdbc/sampledbs");
Connection con = ds.getConnection();
```

Related tasks

[“Configuring the IMS Universal drivers for SSL support” on page 668](#)

With type-4 connectivity, the IMS Universal drivers provide support for the Secure Sockets Layer (SSL) through the Java Secure Socket Extension (JSSE).

Related reference

[javax.sql.DataSource methods supported \(Application Programming APIs\)](#)

Connecting to an IMS database by using the JDBC DriverManager interface


A JDBC application can establish a connection to a data source using the JDBC DriverManager interface, which is part of the `java.sql` package.

The Java application first loads the JDBC driver by invoking the `Class.forName` method. After the application loads the driver, it connects to a database server by invoking the `DriverManager.getConnection` method. For example:


```
Connection conn = DriverManager.getConnection(url);
```

To connect to an IMS database by using the DriverManager interface in your IMS Universal JDBC driver application:

1. Load the IMS Universal JDBC driver with the DriverManager interface by invoking the `Class.forName` method with the argument `com.ibm.ims.jdbc.IMSDriver`.
2. Connect to the IMS database by invoking the `DriverManager.getConnection` method. The URL represents a data source, and indicates what type of JDBC connectivity you are using.
 - For type-4 connectivity, specify the URL in the following form:



- For type-2 connectivity, specify the URL in the following form:



The parts of the URL have the following meaning:

jdbc:ims:

Indicates that the connection is to an IMS database.

PortNumber

The TCP/IP server port number to be used to communicate with IMS Connect. The port number is defined using the DRDAPORT parameter on the ODACCESS statement in the IMS Connect configuration PROCLIB member. The default port number is 8888. Do not set this property when using type-2 connectivity.

MetadataURL

The location of the database metadata representing the target IMS database.

This property is deprecated. Use **DatabaseName** instead.

The **MetadataURL** property is the fully qualified name of the Java metadata class generated by the IMS Enterprise Suite Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **MetadataURL** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

DatabaseName

The location of the database metadata representing the target IMS database.

The **DatabaseName** property can be specified in one of two ways, depending on whether the metadata is stored in the IMS catalog or as a static metadata class generated by the IMS Enterprise Suite Explorer for Development.

If your IMS system uses the IMS catalog, the **DatabaseName** property is the name of the PSB that your application uses to access the target IMS database.

If you are using the IMS Explorer for Development, the **databaseName** property is the fully qualified name of the Java metadata class generated by the IMS Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **DatabaseName** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

DatastoreServer

The name or IP address of the data store server (IMS Connect). You can provide either the host name (for example, `dev123.svl.ibm.com`) or the IP address (for example, `192.166.0.2`). Do not set this property when using type-2 connectivity.

property

One of the following connection properties:

applName

Optional. Specifies the 1- to 8-character application name that is defined to RACF in the PTKTDATA class for JDBC applications that connect to IMS DB by using the SQL Batch utility. The value that is specified on this parameter is used by the SQL Batch utility to generate a RACF PassTicket to authenticate the user of the JDBC application to access IMS DB.

This parameter is valid only for the SQL Batch utility.

allMetadata

Optional. When this property is set to `true`, the DatabaseMetadata interface returns information for all resources in the IMS catalog. When the property is set to `false`, the DatabaseMetadata interface returns information for the allocated PSB. The default value for this property is `false`.

datastoreName

Optional. The name of the IMS data store to access.

When using type-4 connectivity, the **DatastoreName** property must match either the name of the data store defined to ODBM or be blank. The data store name is defined in the ODBM CSLDCxxx PROCLIB member using either the `DATASTORE(NAME=name)` or `DATASTORE(NAME=name, ALIAS(NAME=aliasname))` parameter. If an alias is specified, you must specify the *aliasname* as the value of the **datastoreName** property. If the **DatastoreName** value is left blank (or not supplied), IMS Connect connects to any available instance of ODBM as it is assumed that data sharing is enabled among all datastores defined to ODBM.

When using type-2 connectivity, set the **DatastoreName** property to the IMS subsystem alias. This is not required to be set for the Java Dependent Region run time.

initStatusGroup

Optional. When a connection is made and the PSB is allocated, this property will indicate that the driver should automatically issue an `INIT STATUS GROUPA` or `INIT STATUS GROUPB` if a value of 'A' or 'B' is provided. The default will not issue an `INIT STATUS GROUP` call.

loginTimeout

Optional. Specifies the number of seconds that the driver waits for a response from the server before timing out a connection initialization or server request. Set this property to a non-negative integer for the number of seconds. Set this property to 0 for an infinite timeout length. Do not set this property when using type-2 connectivity.

password

The password for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

sslConnection

Optional. Indicates if this connection uses Secure Sockets Layer (SSL) for data encryption. Set this property to "true" to enable SSL, or to "false" otherwise. Do not set this property when using type-2 connectivity.

sslKeyStoreType

Optional. Specifies the format of the file that contains cryptographic objects needed to establish a secure socket connection. The valid values are "JKS" and "PKCS12". This value is only used when **sslConnection** is set to "true" and **sslKeyStoreType** is not specified. The **sslKeyStoreType** parameter defaults to "JKS".

sslSecureSocketProtocol

Optional. Specifies the cryptographic communication protocol for the new connection. Specify a protocol that is supported by the server and provides the highest level of security. The valid values are "SSL", "SSLv3", "TLSv1.1", and "TLSv1.2". This value is only used when **sslConnection** is set to "true". If **sslConnection** is set to "true" and **sslSecureSocketProtocol** is not specified, a default protocol will be determined at runtime by the JRE and the server.

sslTrustStoreLocation

Optional. Specifies the location of the cryptographic trust store file for the new connection. This value is only used when **sslConnection** is set to true.

sslTrustStorePassword

Optional. Specifies the password to access the cryptographic trust store file. This value is only used when **sslConnection** is set to true.

sslKeyStoreLocation

Optional. Specifies the location of the cryptographic key store file for the new connection. This value is only used when **sslConnection** is set to true.

sslKeyStorePassword

Optional. Specifies the password to access the cryptographic key store file. This value is only used when **sslConnection** is set to true.

user

The user name for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

signedCompare

Optional. When this property is set to "true", special SSAs are generated to support ranged queries over signed data types. If the property is set to "false", standard binary comparisons are performed based on the binary representation of the data type value. Setting the value to

"false" can increase performance but might result in incorrect results. The default value for this property is "true".

flattenTables

Optional. Indicates whether to produce a flattened view of the database tables. A value of true exposes the sub-elements of a STRUCT or an ARRAY as additional columns of the table. The default value is false.

IMS Explorer flattens the copybook structures when you import the copybook. Although the copybook itself remains unchanged in the IMS catalog, the information about the structure of each table is altered for that particular connection.

The the flattenTables property allows you to query the fields in complex structures directly. For more information about support for flattening complex structures, see [“Support for flattening complex structures” on page 566](#).

Restriction: The flattenTables connection property supports static arrays and structures only. Dynamic arrays are not altered.

t2OutputBufferSize

Optional. The size of the output buffer in bytes for the results from a SELECT operation for a type-2 connection.

The minimum value for t2OutputBufferSize is 500000. If any value less than 500000 is set, this property value will be adjusted to 500000. There is no maximum bound. The default value is 1280000.

treatInvalidDecimalAsNull

Optional. Indicates whether to interpret certain Decimal values that appear invalid in Java applications (such as PACKEDDECIMAL and ZONEDDECIMAL with invalid sign bits) as null. By default, this property is "false", and a conversion exception is thrown when the Java applications are processing invalid values.

removeInvalidCaseFields

Optional. If you use SELECT * to access a map case, the JDBC Universal Driver returns the set of all map case segment fields. When this property is set to true, the results contain only map case fields that satisfy a WHERE clause with a DEPENDINGTON field condition. To learn more, see [“Using the removeInvalidCaseFields property” on page 610](#).

currentSchema

Optional. Specifies the default schema name that is used to qualify unqualified database objects in dynamically prepared SQL statements.

dbViewLocation

Optional. Specifies the fully qualified path to a databaseView metadata class. You can use this property to include a metadata class that is not located in your project path.

dpsbOnCommit

Optional. Set this property to **true** to deallocate the PSB when a commit occurs.

Recommendation: Do not set this property to **true** except in a managed environment with integrated connection pooling.

fetchSize

Optional. Gives the client a hint about the number of rows to get from the database when more rows are needed. The number specified for this property only affects data retrieved with the current connection. If the value specified is 0, all of the applicable rows are returned.

The default value for this property is 0 for both managed and unmanaged connections.

llField

Optional. The default value for this property is **true**. You can modify the LL field value to change the length of a variable length segment instance. Set this property to **false** to hide the LL field data from the column list in the result set.

maxRows

Optional. Specifies the maximum number of rows to return in a query result set. The default value is 0, which returns all of the applicable rows in the result set.

expandArrayResultSet

Optional. This property enables IMS Universal JDBC driver to create an **ArrayResultSet** with the array elements as the actual fields in the ResultSet instead of a Struct object containing those elements as the only entry into that ResultSet. To learn more, see [“Using the expandArrayResultSet property” on page 611](#).

ssaOptimization

Optional. When set to true, the SSA is optimized to a fully qualified key if subfields of that key are provided. To learn more, see [Optimizing SQL statements to promote partial key subfields in the WHERE Clause \(Application Programming\)](#).

traceFile

Optional. Specifies the name of the trace file for the connection.

traceFileAppend

Optional. If the specified trace file exists, setting this property to true specifies that the trace data for the new connection must be appended to the existing trace file instead of overwriting it.

This property is ignored if no value is specified for **traceFile**.

traceDirectory

Optional. Specifies the file system directory where the trace file is located. By default, this path is the directory where the application is executed.

This property is ignored if no value is specified for **traceFile**.

traceLevel

Optional. Specifies which traces are enabled for the connection. The valid values for this property are defined in the Java API documentation for the IMSDataSource class.

By default, all traces are disabled.

This property is ignored if no value is specified for **traceFile**.

Trace level	traceLevel package value	traceLevel constant field in IMSDataSource	traceLevel demical value
All	com.ibm.ims.db.opendb.*	TRACE_ALL	-1
DL/I	com.ibm.ims.db.opendb.dli.*	TRACE_DLI	28
DRDA	com.ibm.ims.db.opendb.drda.*	TRACE_DRDA	1
JDBC	com.ibm.ims.db.opendb.jdbc.*	TRACE_JDBC	32
Java EE	com.ibm.ims.opendb.spi.* com.ibm.ims.db.opendb.cci.*	TRACE_JEE	192

value

A valid value for the connection property.

To set the **sslConnection** and **loginTimeout** properties, use a `java.util.Properties` object. For example, the following sample code shows how to enable SSL and set the timeout value to 10 seconds:

```
Properties props = new Properties();
props.put("sslConnection", "true");
props.put("timeout", "10");
```

3. For type-4 connectivity, you must specify a user ID and password in one of the following ways: through the connection URL, through parameters, or through a `java.util.Properties` object.

To set the user ID and password for the connection through parameters, use the form of the `getConnection` method that specifies *user* and *password*. For example:

```
String url =
"jdbc:ims://tst.svl.ibm.com:8888/class://BMP2.BMP2DatabaseView";
String user = "myUserID";
String password = "MyPassword";
Connection conn = DriverManager.getConnection(url, user, password);
```

To set the user ID and password for the connection through a `java.util.Properties` object, use the form of the `getConnection` method that specifies a `java.util.Properties` object. For example:

```
Properties props = new Properties();
props.put( "user", "myUserID" );
props.put( "password", "MyPassword" );
String url =
"jdbc:ims://tst.svl.ibm.com:8888/class://BMP2.BMP2DatabaseView";
Connection conn = DriverManager.getConnection(url, props);
```

4. After your application has finished with the connection, close the connection using the `close` method on the `Connection` interface.

The following code example shows how to create a type-4 connection to an IMS database from your IMS Universal JDBC driver application using the `DriverManager` interface:

```
Connection conn = null;

// Create Properties object
Properties props = new Properties();

// Enable SSL for connection
props.put("sslConnection", "true");

// Set datastoreName for connection
props.put( "datastoreName", "IMS1" );

// Set timeout for connection
props.put("loginTimeout", "10");

// Set user ID for connection
props.put( "user", "myUserID" );

// Set password for connection
props.put( "password", "myPassword" );

// Set URL for the data source
Class.forName("com.ibm.ims.jdbc.IMSDriver");

// Create connection
conn = DriverManager.getConnection
("jdbc:ims://tst.svl.ibm.com:8888/class://BMP2.BMP2DatabaseView",
props);
```

Alternatively, you can specify the connection properties in the URL. For example:

```
String url="jdbc:ims://tst.svl.ibm.com:8888/class://"
+ "BMP2.BMP2DatabaseView:datastoreName=IMS1;"
+ "loginTimeout=10;sslConnection=true;user=myUserID;password=myPassword;";
Connection conn = DriverManager.getConnection(url);
```

Related tasks

[“Configuring the IMS Universal drivers for SSL support” on page 668](#)

With type-4 connectivity, the IMS Universal drivers provide support for the Secure Sockets Layer (SSL) through the Java Secure Socket Extension (JSSE).

Related reference

[SQL Batch utility \(Database Utilities\)](#)

[How to enable the IMS JDBC trace](#)

Sample application for the IMS Universal JDBC driver

The following sample Java application demonstrates the basic programming flow for a JDBC application using the IMS Universal JDBC driver.

The following example connects to an IMS database, retrieves a list of patient names using a SQL SELECT query, and modifies the patient information using a SQL UPDATE query.

```
package client;

import java.sql.Connection;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.Statement;

import com.ibm.ims.jdbc.IMSDataSource;

public class JDBCExample {

    public static void main(String[] args)
        throws SQLException {
        IMSDataSource ds = new IMSDataSource();
        ds.setDatabaseName("MYPDB");
        ds.setDatastoreName("IMS1");
        ds.setDatastoreServer("ec0123.my.host.com");
        ds.setPortNumber(5555);
        ds.setDriverType(IMSDatasource.DRIVER_TYPE_4);
        ds.setUser("myUserId");
        ds.setPassword("myPassword");

        Connection conn = null;

        try {
            conn = ds.getConnection();

            Statement st = conn.createStatement();

            // List all of the patient names in the
            // SURG ward in the ALEXANDRIA hospital
            ResultSet rs = st.executeQuery("SELECT patname from " +
                "pcb01.hospital, ward, patient " +
                "where hospital.hospname = 'ALEXANDRIA' " +
                "and ward.wardname = 'SURG'");
            while (rs.next()) {
                System.out.println(rs.getString("patname"));
            }

            // Update the name of the patient with patient
            // number 0222 in ward 04 in the hospital
            // with code R1210010000A
            int updatedRecords = st.executeUpdate("UPDATE PCB01.PATIENT " +
                "SET PATNAME='UPDATED NAME' WHERE PATNUM='0222' " +
                "AND HOSPITAL_HOSPCODE='R1210010000A' AND WARD_WARDNO='04'");
            System.out.println("Updated " + updatedRecords + " Record(s)");

            conn.commit();
            conn.close();
        } catch (SQLException e) {
```

```

        e.printStackTrace();
        if (!conn.isClosed()) {
            conn.rollback();
            conn.close();
        }
    }
}
}

```

Using ByteBuffers with the IMS Universal JDBC Driver

The IMS Universal JDBC driver allows the use of INSERT, SELECT, and UPDATE options for full segments using **java.nio.ByteBuffer** as an I/O area.

You can use ByteBuffers as a qualification within the WHERE clause and also use them as an I/O area to INSERT, SELECT, and UPDATE full segment I/O areas from an SQL Query. This feature can be accessed by using the new reserved field in your SQL query: **DFSIOSEG_{YOUR_SEGMENT_NAME}**. You do not have to define this field in your DBD source and it need not exist in the IMS Catalog. IMS Universal JDBC driver will automatically process this field when it is detected in the SQL Query.

Inserting data using ByteBuffers

You can insert a full segment I/O area using IMS Universal JDBC driver with a **PreparedStatement**. The following code sample illustrates inserting a full segment I/O area using ByteBuffers:

```

String preparedInsertStatement = "INSERT INTO WARD (DFSIOSEG_WARD, HOSPITAL_HOSPCODE) values
(?, 'R1210010000A')";
PreparedStatement pst = conn.prepareStatement(preparedInsertStatement);

//Prepare buffer for inserting a ward patient

ByteBuffer bufferInsert = ByteBuffer.allocate(900);
bufferInsert.position(0);
bufferInsert.putShort((short) 900);
bufferInsert.put("0010".getBytes("Cp1047"));
bufferInsert.put("COOL WARD".getBytes("Cp1047"));

// Set prepared value for the buffer and execute the PreparedStatement
pst.setObject(1, bufferInsert);
pst.execute();

```

Selecting data using ByteBuffers

You can retrieve full segment I/O areas by specifying the hidden DFSIOSEG field in your SQL query. After obtaining a **ResultSet** from your query, you can obtain your **ByteBuffer** object simply by getting an **Object** from your **ResultSet** using the new **DFSIOSEG** field as a parameter and casting it to a **ByteBuffer** object. If required, you can add additional fields in conjunction with the **DFSIOSEG** field. The following code sample illustrates how to retrieve a full segment I/O area using ByteBuffers:

```

String selectQuery = "SELECT DFSIOSEG_WARD, WARDNAME FROM WARD WHERE HOSPITAL_HOSPCODE =
'R1210010000A' AND WARDNO = '0010'";
ResultSet rs = st.executeQuery(selectQuery);
while(rs.next())
{
    String wardName = rs.getString("WARDNAME");
    ByteBuffer buffer = (ByteBuffer)rs.getObject("DFSIOSEG_WARD");
}

```

Updating data using ByteBuffers

Updating a full segment I/O area using the IMS Universal JDBC driver can also be achieved using a **PreparedStatement**. The following code sample illustrates how to update a full segment I/O area using ByteBuffers:

```

String preparedUpdateStatement = "UPDATE WARD SET DFSIOSEG_WARD = ?
WHERE HOSPITAL_HOSPCODE = 'R1210010000A' AND WARDNO = '0010'";
PreparedStatement pst = conn.prepareStatement(preparedInsertStatement);

```

```
// Prepare buffer for inserting a ward patient
ByteBuffer bufferUpdate = ByteBuffer.allocate(900);
bufferUpdate.position(0);
bufferUpdate.putShort((short) 900);
bufferUpdate.put("0010".getBytes("Cp1047"));
bufferUpdate.put("COOL WARD".getBytes("Cp1047"));

// Set prepared value for the buffer and execute the PreparedStatement
pst.setObject(1, bufferUpdate);
pst.execute();
```

Considerations for Byte Buffers and Type-2 Java applications

If your JDBC application is running in Type-2 mode, you have the option of obtaining a direct ByteBuffer to optimize the JVM storage use. You can obtain direct storage in the following ways:

- Directly using the **java.nio.ByteBuffer** class

```
ByteBuffer directBuffer = ByteBuffer.allocateDirect(10);
```

- Using the Universal Drivers Application object to obtain direct 31-bit storage

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);
```

Note: If an application is running in the 64-bit JVM, `ByteBuffer.allocateDirect()` will not work with your application and you must directly request 31-bit storage using the `Application.get31BitDirectByteBuffer()` method. `ByteBuffer.allocateDirect()` works in a 31-bit JVM.

It is important to release and free the buffers correctly in your application. You can free buffers that were obtained using the **Application** class in the following ways:

- Specifically freeing buffers

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);

// Doing work
app.free31BitDirectByteBuffer(directBuffer);
```

- IMS silently tracks all buffers allocated using the **Application** class and automatically releases all the checked out buffers during the `end()` procedure.

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);

// Doing work
app.end();
```

Related concepts

[“Using Byte Buffers with the IMS Universal DL/I driver” on page 663](#)

The IMS Universal DL/I driver includes several methods to handle complex data types that are not fully represented by the IMS Catalog.

Related reference

[“Hospital database example” on page 568](#)

The code examples for the IMS Universal drivers application programming topics use the Hospital database.

Using the `removeInvalidCaseFields` property

Use the optional **`removeInvalidCaseFields`** connection property to return only columns and fields in a valid map case.

In IMS, you can use DFSMAP and DFSCASE mapping definitions to create multiple field layouts for a segment. Each layout (or “map case”) has a unique ID, and only one layout can be valid at a time. Layouts are controlled by a defined control field, and a layout is valid when its ID matches the control field value.

If you use a SQL `SELECT *` statement and a `WHERE` clause with the IMS Universal JDBC driver to retrieve map case fields from a segment, the returned results include all fields by default, including `NULL` values. You can simplify your results by using the **`removeInvalidCaseFields`** property to remove case fields from the result set that do not satisfy the `DEPENDINGON` field condition in the `WHERE` clause.

This property only supports use of the equal sign (`=`) operator in the `WHERE` clause. If other operators are used, no case fields are removed from the returned result set.

Examples of supported cases are as follows:

- `SELECT * FROM SEGMENT WHERE CONTROLFIELD = 'VALUE'`
- `SELECT * FROM SEGMENT WHERE CONTROLFIELD = 'VALUE1' OR CONTROLFIELD = 'VALUE2'`
- `SELECT * FROM SEGMENT WHERE CONTROLFIELD1 = 'V1' AND CONTROLFIELD2 = 'V2'`

You must enable the **`removeInvalidCaseFields`** property before you can use it.

Enable the `removeInvalidCaseFields` property

The following code examples show four ways to enable the **`removeInvalidCaseFields`** property.

You can enable the **`removeInvalidCaseFields`** property in `server.xml` for use with the IMS Universal Driver:

```
<connectionFactory jndiName="HOSP_JDBC_T4_EXPANDED">
  <properties.imsudbJLocal databaseName="BMP255" datastoreName="IMS1" datastoreServer="localhost"
    driverType="4" password="myPassword" portNumber="{DRDA_PORT}" user="myUserID"
    removeInvalidCaseFields="true"/>
</connectionFactory>
```

You can enable the **`removeInvalidCaseFields`** in an `IMSManagedConnectionFactory` object:

```
IMSManagedConnectionFactory imsManagedConnectionFactory = new IMSManagedConnectionFactory();
imsManagedConnectionFactory.setremoveInvalidCaseFields(true);
```

You can use the Driver Manager to enable the **`removeInvalidCaseFields`** property via a JDBC URL:

```
String url="jdbc:ims://tst.svl.ibm.com:8888/class://"
    + "BMP2.BMP2DatabaseView:datastoreName=IMS1;"
    +
    "loginTimeout=10;sslConnection=true;user=myUserID;password=myPassword;removeInvalidCaseFields=true;";
Connection conn = DriverManager.getConnection(url);

or

Connection conn = null;

// Create Properties object
Properties props = new Properties();

// Enable removeInvalidCaseFields
props.put("removeInvalidCaseFields", "true");

// Create connection
```

```
conn = DriverManager.getConnection
("jdbc:ims://tst.svl.ibm.com:8888/class://BMP2.BMP2DatabaseView", props);
```

You can enable the **removeInvalidCaseFields** property on an IMSDataSource object:

```
IMSDataSource

IMSDataSource ds = new IMSDataSource();
Properties props = new Properties();
props.setProperty("removeInvalidCaseFields", "true");
props.setProperty("llField", "false");
ds.setProperties(props);

Or

IMSDataSource ds = new IMSDataSource();
ds.setremoveInvalidCaseFields(true);
```

Related concepts

[DFSMAP statement overview \(Database Administration\)](#)

[DFSCASE statement overview \(Database Administration\)](#)

Related reference

[DFSMAP statements \(System Utilities\)](#)

[DFSCASE statements \(System Utilities\)](#)

Using the expandArrayResultSet property

The IMS Universal JDBC driver allows for easier access to Array datatype fields using the **expandArrayResultSet** property.

This property enables IMS Universal JDBC driver to create an **ArrayResultSet** with the array elements as the actual fields in the **ResultSet** instead of a Struct object containing those elements as the only entry into that **ResultSet**.

This property indicates whether the IMS Universal JDBC driver would flatten an array when an **ArrayResultSet** is created. If this property is set to **true**, the array will be expanded to its individual array sub fields in the **ArrayResultSet**. If set to **false**, the application will receive an array object for each array instance that exists within the **ArrayResultSet**. The default value is set to **false**.

Programming example to use expandArrayResultSet property

The following is an example of obtaining a **ResultSet** using the **expandArrayResultSet** property.

```
Properties props = new Properties();
props.setProperty("expandArrayResultSet", "true");
ds.setProperties(props);
Connection conn = ds.getConnection();

Statement st = conn.createStatement(ResultSet.TYPE_FORWARD_ONLY, ResultSet.CONCUR_UPDATABLE);

String selectSql = "SELECT * FROM ARRAY WHERE HOSPITAL_HOSPCODE='R1210010000A' ";
st.executeQuery(selectSql);
rs = st.getResultSet();
while (rs.next()) {
    Array accountArray = rs.getArray("ACCOUNTS");
    ResultSet arrayResultSet = accountArray.getResultSet();

    // delete first row
    arrayResultSet.first();
    arrayResultSet.deleteRow();

    // move two relative to 1
    arrayResultSet.relative(2);

    // Start insert at position 3
    arrayResultSet.moveToInsertRow();
    arrayResultSet.updateString("TYPE", "VISA");
    arrayResultSet.updateString("ACCOUNT_DATA", "1111XXXXXXXX4444");
    arrayResultSet.updateString("EXPIRATION", "0112");
    arrayResultSet.updateString("SECURITY_CODE", "9111");
```

```

arrayResultSet.insertRow();

// move to last remembered position
arrayResultSet.moveToCurrentRow();

// go to the last row
arrayResultSet.last();

// go back 4 relative to position 5. should be at 1
arrayResultSet.relative(-4);

// Update string at position 3.
arrayResultSet.updateString("EXPIRATION", "2099");
arrayResultSet.updateRow();

// Commit all changes
rs.updateRow();
}
rs.close();
conn.close();

```

Enabling the `expandArrayResultSet` property

Enabling the `expandArrayResultSet` property in the `server.xml` for use with the IMS Universal Driver

```

<connectionFactory jndiName="HOSP_JDBC_T4_EXPANDED">
  <properties.imsudbJLocal databaseName="BMP255" datastoreName="IMS1" datastoreServer="localhost"
  driverType="4" password="myPassword" portNumber="{DRDA_PORT}" user="myUserID"
  expandArrayResultSet="true"/>
</connectionFactory>

```

Enabling the `expandArrayResultSet` in an `IMSManagedConnectionFactory` object.

```

IMSManagedConnectionFactory imsManagedConnectionFactory = new IMSManagedConnectionFactory();
imsManagedConnectionFactory.setExpandArrayResultSet(true);

```

Using the Driver Manager to enable the `expandArrayResultSet` property via a JDBC URL

```

String url="jdbc:ims://tst.svl.ibm.com:8888/class://"
+ "BMP2.BMP2DatabaseView:datastoreName=IMS1;"
+
"loginTimeout=10;sslConnection=true;user=myUserID;password=myPassword;expandArrayResultSet=true;"
;
Connection conn = DriverManager.getConnection(url);

or

Connection conn = null;

// Create Properties object
Properties props = new Properties();

/ Enable expandArrayResultSet
props.put("expandArrayResultSet", "true");

// Create connection
conn = DriverManager.getConnection
("jdbc:ims://tst.svl.ibm.com:8888/class://BMP2.BMP2DatabaseView",
props);

```

Enabling the `expandArrayResultSet` property on an `IMSDataSource` object.

```

IMSDataSource

IMSDataSource ds = new IMSDataSource();
Properties props = new Properties();
props.setProperty("expandArrayResultSet", "true");
props.setProperty("llField", "false");
ds.setProperties(props);

Or

IMSDataSource ds = new IMSDataSource();
ds.setExpandArrayResultSet(true);

```


Related concepts

[“Using Byte Buffers with the IMS Universal DL/I driver” on page 663](#)

The IMS Universal DL/I driver includes several methods to handle complex data types that are not fully represented by the IMS Catalog.

Related reference

[“Programming with the Callout API” on page 692](#)

Use the `IMSCallout` API to issue IMS synchronous callout requests from JMP or JBP applications that are running in a Java dependent region. Another option is to use the Java Message Service (JMS) API, which provides limited support.

Writing SQL queries to access an IMS database with the IMS Universal JDBC driver

Use the IMS Universal JDBC driver to connect to an IMS database for writing SQL queries.

The IMS catalog provides metadata to your application program. You can write SQL queries to access IMS data based on the metadata information available in the catalog database.

If your IMS system does not use the IMS catalog, you can use the IMS Enterprise Suite Explorer for Development instead.

The IMS Explorer for Development generates a Java database metadata class. This class contains the PSB and DBD metadata classes that the IMS Universal drivers use at runtime.

The metadata includes information about the IMS database, including segments, segment names, the segment hierarchy, fields, field types, field names, fields offsets, and field lengths.

The metadata is used by the IMS Universal JDBC driver to allocate program specification blocks (PSBs), issue DL/I calls, perform data transformation, and translate SQL queries to DL/I calls.

The following table summarizes the mapping between IMS database elements and relational database elements.

Table 101. Mapping between IMS database elements and relational database elements.

Hierarchical database elements in IMS	Equivalent relational database elements
Segment name	Table name
Segment instance	Table row
Segment field name	Column name
Segment unique key	Table primary key
Foreign key field	Table foreign key

Related concepts

[IMS Explorer for Development overview](#)

SQL keywords supported by the IMS JDBC drivers

The SQL support provided by the IMS Universal JDBC driver is based on the SQL-92 standard for relational database management systems.

If you use a SQL keyword as a name for a PCB, segment, or field, your JDBC application program will throw an error when it attempts an SQL query. These keywords are not case-sensitive.

The following SQL keywords are supported by the IMS JDBC drivers:

ABS
ACOS
ALL

ALTER
AND
AS
ASC
ASIN
ATAN
ATAN2
AVG
BETWEEN
CEIL
CEILING
COS
COSH
COT
COUNT
CREATE
DEGREES
DELETE
DESC
DISTINCT
DROP
EXP
FETCH
FIRST
FLOOR
FROM
GROUP BY
INNER
INSERT
INTO
JOIN
LN
LOG
LOG10
MAX
MIN
MOD
NULL
ON
ONLY
OR
ORDER BY
POWER
RADIANS
ROW
ROWS
SELECT
SET
SIGN
SIN
SINH

SQRT
SUM
TAN
TANH
UPDATE
VALUES
WHERE

Related reference

[“Portable SQL keywords restricted by the IMS Universal JDBC drivers” on page 616](#)

If you use any of the following SQL keywords as a name for a PCB, segment, or field, your JDBC application will receive an error when it attempts an SQL query. Instead, use the aliasing feature of the IMS Enterprise Suite Explorer for Development. These keywords are not case-sensitive.

SQL aggregate functions supported by the IMS JDBC drivers

The IMS Universal JDBC driver supports SQL aggregate functions and related keywords.

- AS
- AVG
- COUNT
- GROUP BY
- MAX
- MIN
- ORDER BY
 - ASC
 - DESC
- SUM

The `ResultSet` type for aggregate functions and ORDER BY and GROUP BY clauses is always `TYPE_SCROLL_INSENSITIVE`.

The following table shows the data types of the fields that are accepted by the aggregate functions, along with the resulting data type in the `ResultSet`.

Table 102. Supported SQL aggregate functions and their supported data types

Function	Argument type	Result type
SUM and AVG	Byte	Long
	Short	Long
	Integer	Long
	Long	Long
	BigDecimal	Double-precision floating point
	Single-precision floating point	Double-precision floating point
	Double-precision floating point	Double-precision floating point
MIN and MAX	Any type except BIT, BLOB, or BINARY	Same as argument type
COUNT	Any type	Long

Column names generated by aggregate functions

The `ResultSet` column name from an aggregate function is a combination of the aggregate function name and the field name separated by an underscore character (`_`). For example, the statement `SELECT MAX(age)` results in a column name `MAX_age`. Use this column name in all subsequent references—for example, `resultSet.getInt("MAX_age")`.

If the aggregate function argument field is table-qualified, the `ResultSet` column name is the combination of the aggregate function name, the table name, and the column name, separated by underscore characters (`_`). For example, `SELECT MAX(Employee.age)` results in a column name `MAX_Employee_age`.

Using the AS clause

You can use the `AS` keyword to rename the aggregate function column in the result set or any other field in the `SELECT` statement. You cannot use the `AS` keyword to rename a table in the `FROM` clause. When you use the `AS` keyword to rename the column, you must use this new name to refer to the column. For example, if you specify `SELECT MAX(age) AS oldest`, a subsequent reference to the aggregate function column is `resultSet.getInt("oldest")`.

If you are using the IMS Universal JDBC driver and you specified a `SELECT` query with column names renamed by an `AS` clause, you can only refer to the field in the resulting `ResultSet` by the `AS` rename. However, in the rest of your `SELECT` query, in the `WHERE`, `ORDER BY`, and `GROUP BY` clauses, you can use either the original column name or the `AS` rename.

Using the ORDER BY and GROUP BY clauses

Important: The field names that are specified in a `GROUP BY` or `ORDER BY` clause must match exactly the field name that is specified in the `SELECT` statement.

When using the IMS Universal JDBC driver, the following queries with the `ORDER BY` and `GROUP BY` clauses are valid:

```
SELECT HOSPNAME, COUNT(PATNAME) AS PatCount FROM PCB01.HOSPITAL, PATIENT
GROUP BY HOSPNAME ORDER BY HOSPNAME
```

```
SELECT HOSPNAME, COUNT(DISTINCT PATNAME) AS PatCount FROM PCB01.HOSPITAL,
PATIENT GROUP BY HOSPNAME ORDER BY HOSPNAME
```

Using the COUNT function with DISTINCT

When using the IMS Universal JDBC driver, the `COUNT` aggregate function can be qualified with the `DISTINCT` keyword. For example, the following query returns all hospital names listed in ascending order along with the number of distinct patient names from that hospital. The `COUNT` aggregate function generates a column name `COUNT_DISTINCT_PATNAME`.

```
SELECT HOSPNAME, COUNT(DISTINCT PATNAME)FROM PCB01.HOSPITAL, PATIENT
GROUP BY HOSPNAME ORDER BY HOSPNAME
```

Portable SQL keywords restricted by the IMS Universal JDBC drivers

If you use any of the following SQL keywords as a name for a PCB, segment, or field, your JDBC application will receive an error when it attempts an SQL query. Instead, use the aliasing feature of the IMS Enterprise Suite Explorer for Development. These keywords are not case-sensitive.

The keywords shown in the following table are reserved SQL keywords.

ABORT to CROSS	CURRENT to IS	JOIN to REAL	REFERENCES to WORK
ABORT	CURRENT	JOIN	REFERENCES
ANALYZE	CURSOR	LAST	RESET

ABORT to CROSS	CURRENT to IS	JOIN to REAL	REFERENCES to WORK
AND	DECIMAL	LEADING	REVOKE
ALL	DECLARE	LEFT	RIGHT
ALLOCATE	DEFAULT	LIKE	ROLLBACK
ALTER	DELETE	LISTEN	ROW
AND	DESC	LOAD	ROWS
ANY	DISTINCT	LOCAL	SELECT
ARE	DO	LOCK	SET
AS	DOUBLE	MAX	SETOF
ASC	DROP	MIN	SHOW
ASSERTION	END	MOVE	SMALLINT
AT	EXECUTE	NAMES	SUBSTRING
AVG	EXISTS	NATIONAL	SUM
BEGIN	EXPLAIN	NATURAL	TABLE
BETWEEN	EXTRACT	NCHAR	TO
BINARY	EXTEND	NEW	TRAILING
BIT	FALSE	NO	TRANSACTION
BOOLEAN	FETCH	NONE	TRIM
BOTH	FIRST	NOT	TRUE
BY	FLOAT	NOTIFY	UNION
CASCADE	FOR	NULL	UNIQUE
CAST	FOREIGN	NUMERIC	UNLISTEN
CHAR	FROM	ON	UNTIL
CHARACTER	FULL	ONLY	UPDATE
CHECK	GRANT	OR	USER
CLOSE	GROUP	ORDER	USING
CLUSTER	HAVING	OUTER	VACUUM
COLLATE	IN	PARTIAL	VALUES
COLUMN	INNER	POSITION	VARCHAR
COMMIT	INSERT	PRECISON	VARYING
CONSTRAINT	INT	PRIMARY	VERBOSE
COPY	INTERVAL	PRIVILEGES	VIEW
COUNT	INTERVAL	PROCEDURE	WHERE
CREATE	INTO	PUBLIC	WITH
CROSS	IS	REAL	WORK

Related reference

[“SQL keywords supported by the IMS JDBC drivers” on page 613](#)

The SQL support provided by the IMS Universal JDBC driver is based on the SQL-92 standard for relational database management systems.

Writing DDL statements to modify IMS resources with the IMS Universal JDBC driver

You can write a java application to create or modify an active DBD or PSB in an IMS system.

1. Create a connection to an IMS system by using the Universal Drivers.
2. Create a statement and execute the DDL query.
3. Execute a query called Commit DDL.

For example:

```
ds = FVTConnectionFactory.getIMSDataSource(alias, driverType,
host, port, userName,

password, url);
con = ds.getConnection();

Statement st = con.createStatement();
st.executeUpdate("CREATE DATABASE MYDB ACCESS SHSAM CCSID 'Cp1047'
VERSION '2.0'");
st.executeUpdate("CREATE TABLE MYTABLE (COLUMN1 DECIMAL(5,2)
INTERNALNAME COLUMN1 TYPE C BYTES 10 START 1) AMBIGUOUS INSERT LAST IN MYDB");
st.executeUpdate("CREATE TABLESPACE tb1 IN MYDB BLOCK PRIMARY
32768");
st.executeUpdate("COMMIT DDL");
```

Optimizing SQL statements to promote partial key subfields in the WHERE Clause

You can use a variety of SSA optimization techniques to improve access to your IMS data. To use subfield values to define a search range within a primary key range, set the optional `ssaOptimization` property to `true`.

A general SSA optimization technique is to use a SQL WHERE statement to qualify access to your IMS data. In the metadata grammar, define a searchable field as a column, then write a SQL statement with a WHERE clause that references the column. A Segment Search Argument (SSA) is generated that optimizes access by applying your WHERE qualification to the search field. For best performance, the WHERE clause must reference columns that supply one or more key values from the primary key or a secondary index.

You can use a similar technique to optimize your IMS data access with a *partial key*. To create a partial key, map multiple columns to an IMS primary key or secondary index so that each column references only a portion of the field. Then write a WHERE clause that references only a subset of the columns. Optimization occurs when the partial key is used to generate a range that specifies the lowest and highest key values within that subset of columns. When you write the WHERE clause, specify partial keys in the sequence in which they map to the IMS field. To learn more, see [Partial keys](#).

In some cases, partial key values only cover the beginning byte or bytes of the key rather than the entire offset range. The optional property **`ssaOptimization`** is available for such cases. When set to `true`, this optimization creates a search range on the key in the SSA list, and then fills the remaining bytes with the appropriate values. Rather than using low and high *key* values, this optimization uses the low and high *hexadecimal* values of the remaining bytes to create low and high bounds for the range. The result is a search for all the values that are in the range that begin with any given subfield value. To use this property, set **`ssaOptimization`** to **`true`**.

Example

In the following example, a searchable IMS key field named MYSAMPLE is 10 bytes long. MYSAMPLE is mapped to three columns, a WHERE clause is included in the SQL query, and the `ssaOptimization` property is set to true. IMS returns only the segments that match the WHERE qualification and that are within the generated SSA range.

Step 1. Map three columns to MYSAMPLE in the following sequence:

```
COLUMN1 = bytes 1-3  
COLUMN2 = bytes 4-5  
COLUMN3 = bytes 6-10
```

Step 2. Issue the following SQL query:

```
WHERE COLUMN1 = 'abc' and COLUMN2 = '00'
```

The following SSA is generated:

```
MYSAMPLE >= abc00(low values) & MYSAMPLE <= abc00(high values).
```

IMS returns segments where bytes 1–5 are equal to 'abc00' and where the bytes in COLUMN3 are within the range defined by the lowest and highest hexadecimal values for bytes 6–10.

The SSA is generated successfully and optimization occurs because COLUMNS 1 and 2 are specified in proper sequence. Optimization cannot occur if the WHERE clause references only COLUMN2 or COLUMN3 because the query processor cannot generate an SSA and the query processor or connector must retrieve all of the mapped segments and perform all filtering logic.

Related concepts

[“WHERE clause subfield support” on page 627](#)

When you use the IMS JDBC drivers to pass SQL statements, you can use the WHERE clause to list subfields of any field, as long as the field itself is searchable and is fully defined by the subfields.

Related information

[Keyed access techniques, SSA and IMS optimization](#)

SQL statement usage with the IMS Universal JDBC driver

The following usage rules apply to SQL statements passed to IMS with the IMS Universal JDBC driver.

Foreign key fields

In relational databases, hierarchies can be logically built by creating foreign key relationships between tables. In IMS, the hierarchies are explicit and are part of the database definition itself. The IMS Universal JDBC driver introduces the concept of *foreign keys* to capture these explicit hierarchies in a relational sense, which makes the SQL syntax for IMS equivalent to standard SQL.

When accessing IMS databases with the IMS Universal JDBC driver, every table that is not the root table in a hierarchic path will virtually contain the unique keys of all of its parent segments up to the root of the database. These keys are called *foreign key fields*.

For segments with secondary indexes, the secondary index is also the primary key of the segment. Foreign keys that correspond with the segment are derived from the name of the secondary index.

Restriction: Secondary indexes cannot be referenced in SELECT statements. They can only be referenced in WHERE clauses.

The purpose of the foreign key fields is to maintain referential integrity, similar to foreign keys in relational databases. This allows SQL SELECT, INSERT, UPDATE, and DELETE queries to be written against specific tables and columns located in a hierarchic path.

Remember: Foreign keys are maintained internally by the IMS Universal JDBC driver; the keys are not physically stored in the IMS database.

Hospital database example without a secondary index

For example, in the Hospital database, the HOSPITAL, WARD, and PATIENT tables are on the same hierarchic path. The JDBC application would view the tables as containing the following columns.

HOSPITAL table

Columns are:

- HOSPNAME
- HOSPCODE (primary key)

WARD table

Columns are:

- WARDNO (primary key)
- WARDNAME
- PATCOUNT
- NURCOUNT
- DOCCOUNT
- HOSPITAL_HOSPCODE (foreign key field referencing the HOSPCODE column in the HOSPITAL table)

PATIENT table

Columns are:

- PATNUM (primary key)
- PATNAME
- WARD_WARDNO (foreign key field referencing the WARDNO column in the WARD table)
- HOSPITAL_HOSPCODE (foreign key field referencing the HOSPCODE column in the HOSPITAL table)

The following queries show how SQL SELECT statements can use foreign keys, based on the previous database example. The following statement retrieves all columns from a PATIENT table derived from a child segment under HOSPITAL and WARD on a hierarchic path:

```
SELECT * FROM PCB01.PATIENT
WHERE HOSPITAL_HOSPCODE = 'H5140070000H'
AND WARD_WARDNO = '0023'
```

The following example shows an INSERT statement using foreign keys:

```
INSERT INTO PCB01.PATIENT (PATNUM, PATNAME,
WARD_WARDNO, HOSPITAL_HOSPCODE)
VALUES ('00345', 'John Doe', '0023', 'H5140070000H')
```

The following statements retrieve the hospital code and all ward names from a WARD table. These statements are all equivalent:

```
SELECT HOSPITAL.HOSPCODE, WARD.WARDNAME
FROM PCB01.HOSPITAL, PCB01.WARD
```

```
SELECT HOSPITAL_HOSPCODE, WARD.WARDNAME
FROM PCB01.WARD
```

```
SELECT WARD.HOSPITAL_HOSPCODE, WARD.WARDNAME
FROM PCB01.WARD
```

```
SELECT HOSPITAL_HOSPCODE, WARDNAME
FROM PCB01.WARD
```


The following statement will **fail** because the column HOSPITAL_HOSPCODE is not in the table HOSPITAL.:

```
SELECT HOSPITAL_HOSPCODE FROM PCB01.HOSPITAL
```

Hospital database example with a secondary index

The following example shows how the previous example would change if the HOSPITAL table had a secondary index.

HOSPITAL table

Columns are:

- HOSPNAME (secondary index and primary key)
- HOSPCODE

WARD table

Columns are:

- WARDNO (primary key)
- WARDNAME
- PATCOUNT
- NURCOUNT
- DOCCOUNT
- HOSPITAL_HOSPNAME (foreign key field referencing the HOSPNAME secondary index in the HOSPITAL table)

PATIENT table

Columns are:

- PATNUM (primary key)
- PATNAME
- WARD_WARDNO (foreign key field referencing the WARDNO column in the WARD table)
- HOSPITAL_HOSPNAME (foreign key field referencing the HOSPNAME secondary index in the HOSPITAL table)

CREATE statement usage

The CREATE statement is used to create a resource in IMS.

Each one of the create commands has their own separate list of keywords available. You can issue the following create commands:

- [CREATE DATABASE \(Application Programming APIs\)](#)
- [CREATE TABLE \(Application Programming APIs\)](#)
- [CREATE TABLESPACE \(Application Programming APIs\)](#)
- [CREATE PROGRAMVIEW \(Application Programming APIs\)](#)

Example

Issuing the CREATE DATABASE statement, looks similar to the following.

```
CREATE DATABASE MYDB ACCESS HIDAM VSAM LIKE BASEDB CCSID 'UTF-8'  
DATA CAPTURE CHANGES (EXIT1
```

ALTER statement usage

The ALTER statement is used to modify a resource in IMS.

Each one of the alter commands has their own separate list of keywords available. You can issue the following alter commands:

- [ALTER DATABASE \(Application Programming APIs\)](#)
- [ALTER TABLE \(Application Programming APIs\)](#)
- [ALTER TABLESPACE \(Application Programming APIs\)](#)

Example

Issuing the ALTER DATABASE statement, looks similar to the following.

```
ALTER DATABASE MYDB CCSID 'Cp1047' VERSION '1.8' DATXEXITYES
```

Note: This alter statement will only modify the given keywords, any other keywords that were originally defined will not change unless specified specifically in the ALTER statement.

DROP statement usage

The DROP statement is used to remove a resource in IMS.

Each one of the drop commands has their own separate list of keywords available. You can issue the following drop commands:

- [DROP DATABASE \(Application Programming APIs\)](#)
- [DROP PROGRAMVIEW \(Application Programming APIs\)](#)
- [DROP TABLE \(Application Programming APIs\)](#)
- [DROP TABLESPACE \(Application Programming APIs\)](#)

Example

Issuing the DROP DATABASE statement, looks similar to the following.

```
DROP DATABASE MYDB
```

SELECT statement usage

The SELECT statement is used to retrieve data from one or more tables. The result is returned in a tabular result set.

When using the SELECT statement with the IMS Universal JDBC driver:

- If you are selecting from multiple tables and the same column name exists in one or more of these tables, you must table-qualify the column or an ambiguity error will occur.
- The FROM clause must list all the tables you are selecting data from. The tables listed in the FROM clause must be in the same hierarchic path in the IMS database.
- In Java applications using the IMS JDBC drivers, connections are made to PSBs. Because there are multiple database PCBs in a PSB, queries must specify which PCB in a PSB to use. To specify which PCB to use, always qualify segments that are referenced in the FROM clause of an SQL statement by prefixing the segment name with the PCB name. You can omit the PCB name only if the PSB contains only one PCB.

Examples of valid IMS Universal JDBC driver SELECT queries

Selecting specified columns

The following statement retrieves the ward names and patient names from the WARD and PATIENT tables, respectively:

```
SELECT WARD.WARDNAME,PATIENT.PATNAME
FROM PCB01.WARD, PATIENT
```

Selecting all columns with * symbol

The following statement retrieves all columns for the PATIENT table:

```
SELECT *
FROM PCB01.PATIENT
```

The following statement retrieves the hospital name from the HOSPITAL table and all columns from the WARD table:

```
SELECT HOSPITAL.HOSPNAME, WARD.*
FROM PCB01.HOSPITAL, PCB01.WARD
```

Selecting with DISTINCT

The following statement retrieves all distinct patient names from the PATIENT table:

```
SELECT DISTINCT PATNAME
FROM PCB01.PATIENT
```

Selecting with ORDER BY

The ORDER BY clause is used to sort the rows. By default, results are sorted by ascending numerical or alphabetical order. The following statement retrieves all distinct hospital names, sorted in alphabetical order:

```
SELECT DISTINCT HOSPNAME FROM PCB01.HOSPITAL
ORDER BY HOSPNAME
```

The following statement retrieves all ward names sorted in alphabetical order, and the number of patients in each ward sorted in ascending numerical order. If two WARDNAME values in the ORDER BY compare are equal, the tiebreaker will be their corresponding PATCOUNT values (in this case, the row with the numerically smaller corresponding PATCOUNT value is displayed first).

```
SELECT WARDNAME, PATCOUNT FROM PCB01.WARD
ORDER BY WARDNAME, PATCOUNT
```

Use the DESC qualifier to sort the query result in descending numerical or reverse alphabetical order. The following statement retrieves all patient names in reverse alphabetical order:

```
SELECT PATNAME FROM PCB01.PATIENT
ORDER BY PATNAME DESC
```

Use the ASC qualifier to explicitly sort the query result in ascending numerical or reverse alphabetical order. The following statement retrieves all ward names sorted in ascending alphabetical order, and the number of patients in each ward sorted in descending numerical order:

```
SELECT WARDNAME, PATCOUNT FROM PCB01.WARD
ORDER BY WARDNAME ASC, PATCOUNT DESC
```

Selecting with GROUP BY

The GROUP BY clause is used to return results for aggregate functions, grouped by distinct column values. The following statement returns the aggregated sum of all doctors in every ward in a hospital, grouped by distinct ward names:

```
SELECT WARDNAME, SUM(DOCCOUNT)
FROM PCB01.WARD
WHERE HOSPITAL_HOSPCODE = 'H5140070000H'
GROUP BY WARDNAME
```

The following statement returns the hospital name, ward name, and the count of all patients in each ward in each hospital, grouped by distinct hospital names and sub-grouped by ward names:

```
SELECT HOSPNAME, WARDNAME, COUNT(PATNAME)
FROM PCB01.HOSPITAL, WARD, PATIENT
GROUP BY HOSPNAME, WARDNAME
```

Using the AS clause

Use the AS clause to rename the aggregate function column in the result set or any other field in the SELECT statement. The following statement returns the aggregate count of distinct patients in the PATIENT table with the alias of "PATIENTCOUNT":

```
SELECT COUNT(DISTINCT PATNAME)
AS PATIENTCOUNT
FROM PCB01.PATIENT
```

The following statement returns the aggregate count of distinct wards in all hospitals with the alias of "WARDCOUNT", sorted by the hospital names in alphabetical order, and grouped by distinct hospital names (under a renamed column alias "HOSPITALNAME"):

```
SELECT HOSPNAME AS HOSPITALNAME, COUNT(DISTINCT WARDNAME)
AS WARDCOUNT
FROM PCB01.HOSPITAL, WARD
GROUP BY HOSPNAME
ORDER BY HOSPNAME
```

INSERT statement usage

The INSERT statement is used to insert new rows into a table.

Foreign key fields enable the IMS Universal JDBC driver to properly position the new record (or segment instance) to be inserted in the hierarchic path using standard SQL processing, similar to foreign keys in a relational database. When inserting a record in a table at a non-root level, you must specify values for all the foreign key fields of the table.

Examples of valid IMS Universal JDBC driver INSERT statements

Inserting data at the root

The following statement inserts a new HOSPITAL record:

```
INSERT INTO PCB01.HOSPITAL (HOSPCODE, HOSPNAME)
VALUES ('R1210050000A', 'O' 'MALLEY CLINIC')
```

Inserting data into a specified table in a hierarchic path

When inserting a record in a table at a non-root level, you must specify values for all the foreign key fields of the table. The following statement inserts a new ILLNESS record under a specific HOSPITAL, WARD, and PATIENT table. In this example, the ILLNESS table has three foreign keys HOSPITAL_HOSPCODE, WARD_WARDNO, and PATIENT_PATNUM. The new record will be inserted if and only if there is a HOSPCODE in the HOSPITAL table with the value of 'H5140070000H', a WARD table with a WARDNO value of '01', and a PATIENT table with PATNUM value of 'R1210050000A'.

```
INSERT INTO PCB01.ILLNESS (HOSPITAL_HOSPCODE, WARD_WARDNO,
ILLNAME, PATIENT_PATNUM)
VALUES ('H5140070000H', '01', 'COLD', 'R1210050000A')
```

The following statement inserts a new WARD record under a specific HOSPITAL table. In this example, the WARD table has the foreign key HOSPITAL_HOSPCODE. The new record will be inserted if and only if there is a HOSPCODE in the HOSPITAL table with the value of 'H5140070000H'.

```
INSERT INTO PCB01.WARD (WARDNO, HOSPITAL_HOSPCODE, WARDNAME)
VALUES ('0001', 'H5140070000H', 'EMGY')
```

Inserting data in a searchable field with subfields

If a searchable field consists of subfields, you can insert data by setting all the subfield values such that the searchable field is completely populated.

Examples of invalid IMS Universal JDBC driver INSERT statements

Inserting a record at a non-root level without specifying foreign key fields

In this statement, the WARD_WARDNO foreign key field is missing. The query will fail because it violates the referential integrity constraint that all foreign keys must be provided with legal values.

```
INSERT INTO PCB01.PATIENT (HOSPITAL_HOSPCODE, PATNAME, PATNUM)
VALUES ('HW3201', 'JOHN O' 'CONNER', 'Z800')
```

UPDATE statement usage

The UPDATE statement is used to modify the data in a table.

Examples of valid IMS Universal JDBC driver UPDATE statements

Updating one column in a record

The following statement updates the root:

```
UPDATE HOSPITAL SET HOSPNAME = 'MISSION CREEK'
WHERE HOSPITAL.HOSPCODE = 'H001007'
```

Updating multiple columns in a specified record in a hierarchic path

Foreign keys allow the IMS Universal JDBC driver to maintain referential integrity by identifying the exact record (or segment instance) to update. The following statement updates a WARD record under a specific HOSPITAL. In this example, the WARD table has the foreign key HOSPITAL_HOSPCODE. The record will be updated if and only if there is a HOSPCODE in the HOSPITAL table with the value of 'H5140070000H'.

```
UPDATE WARD SET WARDNAME = 'EMGY',
DOCCOUNT = '2', NURCOUNT = '4',
WHERE HOSPITAL_HOSPCODE = 'H5140070000H'
AND WARDNO = '01'
```

Examples of invalid IMS Universal JDBC driver UPDATE statements

Updating a foreign key field

Making an UPDATE on a foreign key field is invalid for the IMS Universal JDBC driver. For example, the following UPDATE query will fail:

```
UPDATE WARD SET WARDNAME = 'EMGY',
HOSPITAL_HOSPCODE = 'H5140070000H'
WHERE WARDNO = '01'
```

DELETE statement usage

The DELETE statement is used to delete rows in a table. DELETE operations are cascaded to all child segments.

Examples of valid IMS Universal JDBC driver DELETE statements

Deleting an entire database

The following statement deletes the HOSPITAL database:

```
DELETE FROM pcb01.HOSPITAL
```

Deleting a root

The following statement deletes the root segment instance and all its children in the hierarchic path.

```
DELETE FROM pcb01.HOSPITAL
WHERE HOSPCODE = 'H5140070000H'
```

Deleting a single record

Foreign keys allow the IMS Universal JDBC driver to maintain referential integrity by identifying the exact record (or segment instance) to delete. The following statement deletes a single record from the WARD table. In this example, the WARD table has the foreign key HOSPITAL_HOSPCODE. The WARD

record will be deleted if and only if there is a HOSPCODE in the HOSPITAL table with the value of 'H5140070000H' and a WARD table with a WARDNO value of '0001'.

```
DELETE FROM pcb01.WARD
WHERE HOSPITAL_HOSPCODE = 'H5140070000H'
AND WARDNO = '0001'
```

Deleting multiple records

The following statement deletes multiple records from the PATIENT table. In this example, the PATIENT table has two foreign keys: HOSPITAL_HOSPCODE and WARD_WARDNO. The PATIENT record will be deleted if and only if there is a HOSPCODE in the HOSPITAL table with the value of 'H5140070000H', a WARD table with WARDNO value of '0001', and a PATIENT table with a PATNUM value greater than '0007'.

```
DELETE FROM pcb01.PATIENT
WHERE PATNUM > '0007'
AND HOSPITAL_HOSPCODE = 'H5140070000H'
AND WARD_WARDNO = '0001'
```

The following statement deletes all WARD segment instances in the entire database:

```
DELETE FROM pcb01.WARD
```

WHERE clause usage

Use the WHERE clause to select data conditionally in SQL SELECT, UPDATE, or DELETE statements.

When using the WHERE clause with the IMS Universal JDBC driver, use columns that are in any table listed in the FROM clause.

Recommendation: Qualify columns with table names to avoid the ambiguity that can arise if a column exists in multiple tables that are joined in the FROM clause.

The IMS JDBC drivers convert the WHERE clause in a SQL query to a segment search argument (SSA) list when querying a database. SSA rules restrict the type of conditions you can specify in the WHERE clause. The following restrictions apply:

- In general, compare columns to values, not other columns. With the introduction of foreign keys, you can compare one column to another column if one column is the foreign key and the other column is the primary key it is referencing. Example:

```
WHERE HOSPITAL_HOSPCODE = HOSPITAL.HOSPCODE
```

You can use the following operators between column names and values in the individual qualification statements:

- =
Equals
- !=
Not equal
- >
Greater than
- >=
Greater than or equals
- <
Less than
- <=
Less than or equals

The WHERE clause in the following example is valid because it compares a column to a value:

```
WHERE PAYMENTS.PATNUM='A415'
```

In contrast, the following WHERE clause is invalid and fails because it compares two columns:

```
WHERE PAYMENTS.PATNUM=PAYMENTS.AMOUNT
```

- Do not use parentheses. Qualification statements are evaluated from left to right. The order of evaluation for operators is the IMS evaluation order for segment search arguments.
- List all qualification statements for a table adjacently. For example, in the following valid WHERE clause, the qualified columns from the same PATIENT table are listed adjacently:

```
WHERE PATIENT.PATNAME='BOB' OR PATIENT.PATNUM='A342' AND WARD.WARDNO='52'
```

The following invalid WHERE clause fails because the columns from the HOSPITAL table are separated by the columns from the WARD table:

```
WHERE HOSPITAL.HOSPNAME='Santa Teresa' AND WARD.WARDNO='52'  
OR WARD.WARDNAME='CARD' AND HOSPITAL.HOSPCODE='90'
```

- The OR operator can be used only between qualification statements that contain columns from the same table. You cannot use the OR operator across tables. To combine qualification statements for different tables, use an AND operator. For example, the following invalid WHERE clause will fail:

```
WHERE WARD.WARDNO='03' OR PATIENT.PATNUM='A415'
```

However, the following WHERE clause is valid because the OR operator is between two qualification statements for the same table:

```
WHERE PATIENT.PATNUM='A409' OR PATIENT.PATNAME='Sandy'
```

- When using prepared statements, you can use the question mark (?) character, which is later filled in with a value. For example, the following WHERE clause is valid:

```
WHERE PAYMENTS.AMOUNT>?
```

- When using a column that is a virtual foreign key in the qualification statements, you must follow these rules when access to the database is random:
 - Specify a virtual foreign key that is only from one of the first child segments of the root segment (level 2). You cannot specify virtual foreign keys that are at lower levels in the hierarchy.
 - Use the OR operator to combine multiple qualification statements that use the virtual foreign key.
 - Specify qualifications statements that use only the virtual foreign key. For example, you cannot use the OR or the AND operator to specify another column from the table or another table.

For example, the following WHERE clause is valid:

```
WHERE HOSPITAL_HOSPCODE='ARS100100D' OR HOSPITAL_HOSPCODE='ARS100100D'
```

(Optional) **ssaOptimization**- This property indicates whether to optimize the SSA List for SQL queries that include columns that are a subset of a primary key or a secondary index in the WHERE clause. The SSA List would be modified to a ranged search on the key provided that the column/columns cover the beginning byte/bytes of the key offset range. The default value is **false**.

WHERE clause subfield support

When you use the IMS JDBC drivers to pass SQL statements, you can use the WHERE clause to list subfields of any field, as long as the field itself is searchable and is fully defined by the subfields.

For example, a DBD-defined field is named ADDRESS and is 30 bytes long. In a COBOL copybook, this field is split into the subfields CITY, STATE, and ZIPCODE, as illustrated by the following code.

```
01 ADDRESS  
  02 CITY PIC X(10)  
  02 STATE PIC X(10)  
  03 ZIP PIC X(10)
```

Without subfield support, the ADDRESS value in the WHERE clause would have to be padded manually and entered as follows:

```
WHERE ADDRESS = 'san jose  ca          95141      '
```

With subfield support, you can enter the WHERE clause as follows:

```
WHERE CITY = 'san jose'  
      AND STATE = 'ca'  
      AND ZIPCODE = '95141'
```

The IMS JDBC drivers convert the individual subfields and bundle them into the ADDRESS field before sending the SQL query to IMS.

The following rules and restrictions apply to WHERE clause subfield support:

- Parameter markers are supported for subfields. For example, for a prepared statement, the following WHERE clause entry is valid:

```
WHERE CITY = ? AND STATE = ? AND ZIPCODE = ?
```

- The only relational operator supported for subfields is "=" (equals operator).
- The only Boolean operator is "AND" for connecting subfields. The following WHERE clause entry is valid because the subfields are connected using only "AND" operators:

```
WHERE HOSPCODE=? OR CITY = ? AND STATE = ? AND ZIPCODE = ?
```

- All the subfields for a particular searchable field must be specified in the WHERE clause. You cannot omit any subfields of a field. For example, the following WHERE clause entry is invalid because the STATE subfield was not provided:

```
WHERE CITY = ? AND ZIPCODE = ?
```

- When specifying the subfields in a WHERE clause, all the subfields for a searchable field must be listed adjacent to each other. For example, the following WHERE clause entry is invalid because the listing of the subfields is not contiguous:

```
WHERE CITY = ? AND STATE = ? OR HOSPCODE=? AND ZIPCODE = ?
```

- You can enter subfields for multiple searchable fields in the WHERE clause. For example, if the PATNAME field was broken into LASTNAME and FIRSTNAME subfields, you can specify the subfields for ADDRESS and PATNAME as follows:

```
WHERE CITY = ? AND STATE = ? AND ZIPCODE = ?  
      OR LASTNAME = ? AND FIRSTNAME = ?
```

- When specifying the subfields in a WHERE clause across multiple tables, all the subfields for the searchable fields in each table must be listed together, before listing the subfields for the next table. For example, if the ADDRESS field was in the HOSPITAL table and the PATNAME field was in the PATIENT table, the following WHERE clause entry is invalid because not all the ADDRESS subfields have been listed for HOSPITAL:

```
WHERE HOSPITAL.CITY = ? AND HOSPITAL.ZIPCODE = ?  
      AND PATIENT.LASTNAME = ? AND PATIENT.FIRSTNAME = ?
```

Writing DL/I calls to access an IMS database with the IMS Universal JDBC driver

In addition to support for SQL queries, the IMS Universal JDBC driver also supports casts to DL/I objects.

The IMS Universal JDBC driver is normally used to provide an interface for an application program to make SQL queries to get IMS data. However, you can also cast to DL/I objects from the JDBC interface.

You might want to use this approach if you want to reuse a known-good DL/I call in a new application context, or if you want to obtain maximum performance for a simple query.

1. Create JDBC and PSB connection objects.
2. Obtain a connection to the IMS data.
3. Cast the JDBC connection to get a PSB handle from the IMS Universal DL/I driver.
4. Get a PCB object from your existing PSB connection object.
5. Allocate a data connection from your PSB object.
6. Build and submit your segment search argument (SSA) list.
7. Create Path and PathSet objects to contain the returned data.
8. Process the PathSet data.

This example demonstrates how to make the cast to the DL/I driver and obtain data. This examples accesses a JDBC connection from a JNDI datasource that was created in a JEE server with the IMS Universal Database Resource Adapter.

```
import java.sql.Connection;

import javax.naming.InitialContext;
import javax.sql.DataSource;

import com.ibm.ims.dli.PCB;
import com.ibm.ims.dli.PSB;
import com.ibm.ims.dli.Path;
import com.ibm.ims.dli.PathSet;
import com.ibm.ims.dli.SSAList;

public class JDBCtoDLI {
    public static void main(String args[]){

        Connection conn = null; // This is a JDBC Connection
        // This is the equivalent connection object to the JDBC connection for the IMS Java DL/I API
        PSB psb = null;
        try{
            // Lookup the JNDI DataSource that contains the IMS connection information.
            // The JNDI DataSource would be defined in the JEE
            // server with the IMS Universal Database Resource Adapters
            InitialContext ic = new InitialContext();
            DataSource ds = (DataSource) ic.lookup("myJNDIName");

            // Get a JDBC Connection from the DataSource
            conn = ds.getConnection();

            // Cast the JDBC Connection to the IMS ConnectImpl in order to retrieve
            // a handle to the PSB in the IMS Java DL/I API
            psb = ((com.ibm.ims.jdbc.ConnectionImpl) conn).getPSB();

            //Get a PCB using the PSB you just created
            PCB ivp1pcb = psb.getPCB("PHONEAP");
            //Allocate PSB1 to establish a connection to the data
            psb.allocate();
            System.out.println("PSB for IVPDB1 Allocated");

            //Do work on PSB1
            SSAList ssaList = ivp1pcb.getSSAList("PhoneBook");    //Create an SSA list to use in a DLI
call.

                                //This SSA list qualifies the entire
                                //PhoneBook segment. PhoneBook is an
                                //alias name for segment A1111111 which
                                //is specified in the database view
                                //(DFSIVPDBView.java).

            //Create a path object for later use
            Path path = null;
            PathSet ps = ivp1pcb.batchRetrieve(ssaList);    //This statement uses the PCB object
created

                                //above to do a batch retrieval of
                                //all the segment instances of PhoneBook.
                                //The data returned will be placed in
                                //a PathSet which is a collection of
                                //Path's containing the data you requested
```

```

System.out.println("Batch retrieved all segment instances of Phone Book");

System.out.println("FIRSTNAME\tLASTNAME\tEXTENSION\tZIPCODE");
System.out.println("-----");

/*
 * The following while loop process the PathSet by checking that there is
 * a next element (ps.hasNext), then it prints out the three fields that are defined in the
 * database view (FIRSTNAME, LASTNAME, EXTENSION) for segment PhoneBook. This will continue
 * until there are no more elements in the PathSet
 */
while(ps.hasNext()){
    path = ps.next();
    System.out.println(path.getString("FIRSTNAME").trim()+"\t\t"+
        path.getString("LASTNAME").trim()+"\t\t"+
        path.getString("EXTENSION").trim()+"\t\t"+
        path.getString("ZIPCODE").trim());
}

//INSERT a segment into the PhoneBook
path = ssaList.getPathForInsert("PhoneBook");
path.setString("LASTNAME", "LAST15");
path.setString("FIRSTNAME", "FIRST15");
path.setString("EXTENSION", "8-111-1515");
path.setString("ZIPCODE", "D15/R15");
ivp1pcb.insert(path, ssaList);
System.out.println("\nInserted New Phone Book Entry with LASTNAME equal to LAST15");

//Batch retrieve all segment instances of the PhoneBook (A1111111) segment
ps = ivp1pcb.batchRetrieve(ssaList);
System.out.println("\nBatch Retrieved all segment instances of Phone Book and verify LAST15
was inserted");

System.out.println("FIRSTNAME\tLASTNAME\tEXTENSION\tZIPCODE");
System.out.println("-----");
while(ps.hasNext()){
    path = ps.next();
    System.out.println(path.getString("FIRSTNAME").trim()+"\t\t"+
        path.getString("LASTNAME").trim()+"\t\t"+
        path.getString("EXTENSION").trim()+"\t\t"+
        path.getString("ZIPCODE").trim());
}

//UPDATE FIRSTNAME to NEWNAME where LASTNAME equals LAST15
ssaList.addInitialQualification(1, "LASTNAME", SSaList.EQUALS, "LAST15");
if(ivp1pcb.getUnique(path, ssaList, true)){
    path.setString("FIRSTNAME", "NEWNAME");
    if(16448==ivp1pcb.replace(path)){
        System.out.println("\nUpdated FIRSTNAME for segments with LASTNAME of LAST15");
    }
}

ssaList.removeAllQualificationStatements(1);
//Batch retrieve all segment instances of the PhoneBook (A1111111) segment
ps = ivp1pcb.batchRetrieve(ssaList);
System.out.println("\nBatch Retrieved all segment instances of Phone Book and " +
    "\nverify that the FISTNAME was updated for the entry LAST15");

System.out.println("FIRSTNAME\tLASTNAME\tEXTENSION\tZIPCODE");
System.out.println("-----");
while(ps.hasNext()){
    path = ps.next();
    System.out.println(path.getString("FIRSTNAME").trim()+"\t\t"+
        path.getString("LASTNAME").trim()+"\t\t"+
        path.getString("EXTENSION").trim()+"\t\t"+
        path.getString("ZIPCODE").trim());
}

//DELETE all segments where LASTNAME equals LAST15
ssaList.addInitialQualification(1, "LASTNAME", SSaList.EQUALS, "LAST15");
if(ivp1pcb.batchDelete(ssaList)==1){
    System.out.println("\nSegment with LASTNAME equal to LAST15 has been deleted");
}

ssaList.removeAllQualificationStatements(1);

//Batch retrieve all segment instances of the PhoneBook (A1111111) segment
ps = ivp1pcb.batchRetrieve(ssaList);
System.out.println("\nBatch Retrieved all segment instances of Phone Book and " +
    "\nverify that the segment with LASTNAME of LAST15 has been deleted");

System.out.println("FIRSTNAME\tLASTNAME\tEXTENSION\tZIPCODE");

```

```

        System.out.println("-----");
        while(ps.hasNext()){
            path = ps.next();
            System.out.println(path.getString("FIRSTNAME").trim()+"\t\t"+
                path.getString("LASTNAME").trim()+"\t\t"+
                path.getString("EXTENSION").trim()+"\t\t"+
                path.getString("ZIPCODE").trim());
        }

        //Commit the work
        psb.commit();
        System.out.println("\nPSB Committed");
        //Deallocate the PSB
        psb.deallocate();
        System.out.println("PSB deallocated");
        //Close the socket connection
        psb.close();
        System.out.println("Connection Closed");
        System.out.println("Open Database IVP Completed");
    } catch(Exception e) {
        e.printStackTrace();
        try {
            psb.deallocate();
            psb.close();
        } catch(Exception e1){
            e1.printStackTrace();
        }
    }
}
}
}
}

```

Related concepts

[“Programming with the IMS Universal DL/I driver” on page 641](#)

Use the IMS Universal DL/I driver when you need to write granular queries to access IMS databases directly from a Java client in a non-managed environment.

IMS Universal JDBC driver support for XML

You can write applications to store XML data in IMS databases or retrieve XML data from IMS databases by using the IMS Universal JDBC driver. Both the type-4 and type-2 drivers offer this support.

You can use the IMS Universal JDBC driver support for XML to complete the following operations:

- Retrieve XML data from an IMS database as a character large object (CLOB) through a SQL SELECT statement.
- Store XML data into an IMS database, through a SQL INSERT statement, by using either the `PreparedStatement.setClob` method or the `PreparedStatement.setCharacterStream` method.

The syntax for storing and retrieving XML data by using the IMS Universal JDBC driver is independent of how the XML data is physically stored in the IMS database. The interface is not sensitive to whether the data is stored in decomposed storage mode, intact storage mode, or both or whether the data is stored in an existing or new IMS database.

To use IMS Universal JDBC driver support for XML, you need to generate a runtime Java metadata class that corresponds to a program specification block (PSB). The Java metadata class must define the column fields in the IMS database for storing and retrieving XML data and identify the XML schema that describes the structure of the data.

IMS Version 12 and later includes new DBD source parameters (the `DATATYPE=XML` parameter of the `FIELD` statement and the `OVERFLOW` parameter of the `DFSMARSH` statement) that you can use to define XML-containing fields and overflow segments. If your IMS system uses the IMS catalog database, the IMS Enterprise Suite Explorer for Development can make a connection to the catalog to dynamically retrieve the needed metadata instead of generating a static metadata class. If your IMS system does not use the IMS catalog database, these field definitions are included in the static Java metadata class created with the IMS Explorer for Development.

Defining XML datatype column fields in the Java metadata class

To use IMS Universal JDBC driver support for XML, you need to define the XML datatype column fields for storing and retrieving XML data.

Note: 15.3 and later supports new DBD generation parameters for XML datatype definitions: the DATATYPE=XML parameter for the FIELD statement and the OVERFLOW segment definition for the DFSMARSH statement. If you use these parameters, the Java metadata class will already contain the XML definitions and you do not need to modify the class. If you are using the IMS catalog database, the metadata is available with a data connection instead of with a static metadata class.

To define an XML data type column field in the Java metadata class:

1. Generate the Java metadata class with the IMS Enterprise Suite Explorer for Development.
2. Generate the XML schema for the database with the IMS Enterprise Suite DLIModel utility plug-in or manually create the XML schema based on the DBD and data type mappings.
3. Specify the XML datatype column field by modifying the generated Java metadata class.

If you are storing or retrieving XML data in decomposed storage mode, define the XML datatype column field with the following DLTypeInfo constructor syntax. One or more XML datatype column fields can be defined in a segment.

```
public DLTypeInfo(String fieldName,
                  String XMLSchemaName,
                  DLTypeInfo.XML);
```

4. During database connection setup, pass the name of the Java metadata class to the IMS Universal JDBC driver.

The following example shows how to define XML column datatype fields in a Java metadata class for decomposed mode. In this example, an XML datatype column field named "HOSPXML" is defined that is associated with the "BMP255-PCB01.xsd" XML schema. Another XML datatype column field named "HXML" is defined that is associated with the "B.xsd" XML schema.

```
// The following describes Segment: HOSPITAL ("HOSPITAL") in PCB: PCB01 ("PCB01")
static DLTypeInfo[] PCB01HOSPITALArray= {
    new DLTypeInfo("HOSPLL", DLTypeInfo.CHAR, 1, 2, "HOSPLL"),
    new DLTypeInfo("HOSPCODE", DLTypeInfo.CHAR, 3, 12,
        "HOSPCODE", DLTypeInfo.UNIQUE_KEY),
    new DLTypeInfo("HOSPNAME", DLTypeInfo.CHAR, 15, 17, "HOSPNAME"),
    new DLTypeInfo("HOSPXML", "BMP255-PCB01.xsd", DLTypeInfo.XML),
    new DLTypeInfo("HXML", "B.xsd", DLTypeInfo.XML)
};
```

Storing XML data by using the IMS Universal JDBC driver

You can use the IMS Universal JDBC driver to store XML data into an IMS database through an SQL INSERT statement.

To store XML data in your IMS Universal JDBC driver application:

1. Specify the file path that contains the XML schema file (.xsd) that describes the input XML data structure by setting the <http://www.ibm.com/ims/schema-resolver/file/path> environment variable.

The following example shows how to programmatically set the environment variable. In this example, the file path `uxml/samples` indicates a relative path to the XML schema file. You can also specify an absolute file path.

```
System.setProperty("http://www.ibm.com/ims/schema-resolver/file/path",
    "uxml/samples");
```

2. Specify your XML data source.

If you are reading in XML data from an external source, such as a file, you must create a `java.io.Reader` object to wrap the input XML data. The following example shows how to create an `InputStreamReader` object to wrap an external file named `hospwashington.xml`. The

InputStreamReader object converts the bytes that are read from the input file from ASCII encoding to Unicode.

```
String doc = "hospswashington.xml";
InputStream fileStream = getClass().getResourceAsStream(doc);
if (fileStream == null) {
    throw new FileNotFoundException("Insert Document: '" + doc + "' was
    not found in classpath");
}
InputStreamReader fileReader = new InputStreamReader(fileStream, "ASCII");
```

3. Insert the XML data.

- a) Create a `java.sql.PreparedStatement` object representing the SQL INSERT call.

In the SQL INSERT statement, you must specify the name of the XML column to store the XML data. The column name must match the name that is defined in the Java metadata class.

The following example shows how to create a `PreparedStatement` object to insert data into the `hospxml` column in the HOSPITAL segment, using the `java.sql.Connection` instance `conn`.

```
String s = "INSERT INTO pcb01.HOSPITAL (hospxml) VALUES (?)";
PreparedStatement ps = conn.prepareStatement(s);
```

- b) Set the value of the XML data to insert into the `PreparedStatement` object.

The following table describes the methods and corresponding input data types that you can use to insert data in XML columns.

Table 103. Methods and data types for updating XML columns

Method	Input data type
<code>PreparedStatement.setCharacterStream</code>	<code>Reader</code>
<code>PreparedStatement.setClob</code>	<code>Clob</code>

In decomposed storage mode, XML data is stored with EBCDIC encoding. In intact storage mode, the default encoding is Unicode.

The following code sample shows how to insert XML data into the Hospital database.

```
package uxml.samples;

import java.io.*;
import java.sql.Clob;
import java.sql.Connection;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.Statement;
import java.sql.PreparedStatement;
import com.ibm.ims.jdbc.IMSDataSource;

public class StoreXMLSamples{

    public static void main(String argv[]) throws SQLException,IOException {
        IMSDataSource ds = new IMSDataSource();
        ds.setDatabaseName("class://uxml.samples.BMP255NewSyntaxDatabaseView");
        ds.setDatastoreName("IMS1");
        ds.setDatastoreServer("yourhost.yourdomain.com");
        ds.setPortNumber(5555);
        ds.setDriverType(IMSDatasource.DRIVER_TYPE_4);
        ds.setUser("myUserID");
        ds.setPassword("myPass");

        // Specify file path of XML schema
        System.setProperty("http://www.ibm.com/ims/schema-resolver/file/path",
            "uxml/samples");

        Connection conn = null;

        try {
            conn = ds.getConnection();
            Statement st = conn.createStatement();
            String doc = "hospswashington.xml";
```

```

StoreXMLSamples storeSample = new StoreXMLSamples();
InputStream fileStream =
    storeSample.getClass().getResourceAsStream(doc);
if (fileStream == null) {
    throw new FileNotFoundException("Insert Document: '" +
        doc + "' was not found in classpath");
}

// Convert XML document from ASCII to Unicode
InputStreamReader fileReader =
    new InputStreamReader(fileStream, "ASCII");

PreparedStatement ps =
    conn.prepareStatement("INSERT INTO pcb01.HOSPITAL" +
        " (hospxml) VALUES (?)");

ps.setCharacterStream(1, fileReader, -1);
int rows = ps.executeUpdate();
System.out.println("Inserted");
conn.commit();
conn.close();
} catch (SQLException e) {
    e.printStackTrace();
    if (!conn.isClosed()) {
        conn.rollback();
        conn.close();
    }
}
}
}
}

```

Retrieving XML data by using the IMS Universal JDBC driver

You can use the IMS Universal JDBC driver to retrieve XML data from an IMS database as a character large object (CLOB) through an SQL SELECT statement.

To retrieve XML data in your IMS Universal JDBC driver application:

1. Specify the file path that contains the XML schema file (.xsd) describing the input XML data structure by setting the `http://www.ibm.com/ims/schema-resolver/file/path` environment variable.

The following example shows how to programmatically set the environment variable. In this example, the file path `uxml/samples` indicates a relative path to the XML schema file. You can also specify an absolute file path.

```

System.setProperty("http://www.ibm.com/ims/schema-resolver/file/path",
    "uxml/samples");

```

2. Specify and execute an SQL SELECT statement to retrieve the XML data.

The database table in your SQL SELECT statement must include the XML column for retrieving the XML data. If you specify the column name explicitly in the SQL SELECT statement, the column name must match the name defined in the Java metadata class.

The following example shows how to obtain a `java.sql.ResultSet` object from an SQL SELECT call to retrieve the `hospxml` column in the `HOSPITAL` segment. In the example, `st` is a `java.sql.Statement` instance.

```

ResultSet rs = st.executeQuery("SELECT hospxml FROM PCB01.HOSPITAL");

```

3. Read the XML data from the `ResultSet` object after the retrieve call is made.

The XML data is stored in a `java.sql.Clob` object in the `ResultSet`.

The following code sample shows how to retrieve XML data from the Hospital database.

```

package uxml.samples;

import java.io.*;
import java.sql.Clob;
import java.sql.Connection;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.Statement;
import com.ibm.ims.jdbc.IMSDataSource;

```

```

public class RetrieveXMLSamples{

    public static void main(String argv[]) throws SQLException,IOException {
        IMSDataSource ds = new IMSDataSource();
        ds.setDatabaseName("class://uxml.samples.BMP255NewSyntaxDatabaseView");
        ds.setDatastoreName("IMS1");
        ds.setDatastoreServer("yourhost.yourdomain.com");
        ds.setPortNumber(5555);
        ds.setDriverType(IMSDataSource.DRIVER_TYPE_4);
        ds.setUser("myUserId");
        ds.setPassword("myPass");

        // Specify file path of XML schema
        System.setProperty("http://www.ibm.com/ims/schema-resolver/file/path",
            "uxml/samples");

        Connection conn = null;

        try {
            conn = ds.getConnection();

            Statement st = conn.createStatement();

            ResultSet rs = st.executeQuery("SELECT hospxml FROM PCB01.HOSPITAL");

            StringWriter sw = new StringWriter();

            while (rs.next()) {

                Clob clob = rs.getClob(1);
                Reader reader = clob.getCharacterStream();
                char[] buffer = new char[1000];
                int read = reader.read(buffer);
                while (read != -1) {
                    sw.write(buffer,0,read);
                    read = reader.read(buffer);
                }
            }

            String result = sw.toString();
            System.out.println(result);
            System.out.println();

            conn.commit();
            conn.close();
        } catch (SQLException e) {
            e.printStackTrace();
            if (!conn.isClosed()) {
                conn.rollback();
                conn.close();
            }
        }
    }
}

```

Data transformation support for JDBC

The IMS JDBC drivers provide data transformation on behalf of client applications. When provided with the information from the IMS catalog database or Java database metadata class, the libraries are able to internally convert data from one datatype to another. The IMS Universal DL/I driver also includes an extensible user data type converter for translating custom data types.

Supported JDBC data types

The following table lists the supported Java data types for each JDBC data type.

Table 104. Supported JDBC data types

JDBC data type	Java data type	Length
ARRAY	java.lang.Array	Application-defined
BIGINT	long	8 bytes

Table 104. Supported JDBC data types (continued)

JDBC data type	Java data type	Length
BINARY	byte[]	1 - 32 KB
BIT	Boolean	1 byte
CHAR	java.lang.String	1 - 32 KB
CLOB	java.sql.Clob	Application-defined
DATE	java.sql.Date	Application-defined
DOUBLE	double	8 bytes
FLOAT	float	4 bytes
INTEGER	int	4 bytes
PACKEDDECIMAL	java.math.BigDecimal	1 - 10 bytes
SMALLINT	short	2 bytes
STRUCT	java.lang.Struct	Application-defined
TIME	java.sql.Time	Application-defined
TIMESTAMP	java.sql.Timestamp	Application-defined
TINYINT	byte	1 byte
ZONEDDECIMAL	java.math.BigDecimal	1 - 19 bytes

Methods for retrieving and converting data types

With the IMS Universal JDBC driver, you can use the `ResultSet` interface (`java.sql.ResultSet`) to retrieve and convert the data from the type that is defined in the database metadata to the type that is required by your Java application. Similarly, with the IMS Universal DL/I driver, you can use the `Path` interface to perform data retrieval and conversion to Java data types.

The following table shows the available `get` methods in the `ResultSet` interface (for the IMS Universal JDBC driver) or the `Path` interface (for the IMS Universal DL/I driver) for accessing data of a certain Java data type.

The "No Truncation or Data Loss" column indicates the data types that are designed to be accessed with the given `getXXX` method. No truncation or data loss occurs when using those methods for those data types. The data types that are in the "Legal without Data Integrity" column are all other legal calls; however, data integrity cannot be ensured when using the given `getxxx` method to access those data types. If a data type is not in either column, using the given `getXXX` method for that data type will result in an exception.

Table 105. *ResultSet.getXXX* and *Path.getXXX* methods to retrieve data types

ResultSet.getXXX Method or Path.getXXX Method	Data Type (any not listed result in an exception)	
	No Truncation or Data Loss	Legal without Data Integrity
getBytes	TINYINT UTINYINT	SMALLINT INTEGER BIGINT FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getShort	SMALLINT USMALLINT	TINYINT INTEGER BIGINT FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getInt	INTEGER INTEGER	TINYINT SMALLINT BIGINT FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getLong	BIGINT UBIGINT	TINYINT SMALLINT INTEGER FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹

Table 105. *ResultSet.getXXX* and *Path.getXXX* methods to retrieve data types (continued)

ResultSet.getXXX Method or Path.getXXX Method	Data Type (any not listed result in an exception)	
	No Truncation or Data Loss	Legal without Data Integrity
getFloat	FLOAT	TINYINT SMALLINT INTEGER BIGINT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getDouble	DOUBLE	TINYINT SMALLINT INTEGER BIGINT FLOAT BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getBoolean	BIT	TINYINT SMALLINT INTEGER BIGINT FLOAT DOUBLE CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getString	CHAR VARCHAR	TINYINT SMALLINT INTEGER BIGINT FLOAT DOUBLE BIT PACKEDDECIMAL ¹ ZONEDDECIMAL ¹ BINARY DATE TIME TIMESTAMP

Table 105. *ResultSet.getXXX* and *Path.getXXX* methods to retrieve data types (continued)

ResultSet.getXXX Method or Path.getXXX Method	Data Type (any not listed result in an exception)	
	No Truncation or Data Loss	Legal without Data Integrity
getBigDecimal	BINARY ³ PACKEDDECIMAL ¹ ZONEDDECIMAL ¹	TINYINT SMALLINT INTEGER BIGINT FLOAT DOUBLE BIT CHAR VARCHAR
getClob	CLOB ²	all others result in an exception
getBytes	BINARY	SMALLINT INTEGER BIGINT FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL1 ZONEDDECIMAL1 DATE TIME TIMESTAMP
getDate	DATE	CHAR VARCHAR TIMESTAMP
getTime	TIME	CHAR VARCHAR TIMESTAMP
getTimestamp	TIMESTAMP	CHAR VARCHAR DATE TIME

Note:

1. PACKEDDECIMAL and ZONEDDECIMAL are data type extensions for IMS Universal JDBC driver and the IMS Universal DL/I driver. All other types are standard SQL types defined in SQL92. **Restriction:** PACKEDDECIMAL and ZONEDDECIMAL data types do not support the Sign Leading or Sign Separate modes. For these two data types, sign information is always stored with the Sign Trailing method.
2. The CLOB data type is supported only for the retrieval and storage of XML data.
3. The BINARY data type is valid only for decimal data used with a binary type converter.

If the field type is either PACKEDDECIMAL or ZONEDDECIMAL, the type qualifier is the COBOL PICTURE string that represents the layout of the field. All COBOL PICTURE strings that contain valid combinations

of 9s, Ps, Vs, and Ss are supported. Expansion of PICTURE strings is handled automatically. For example, '9(5)' is a valid PICTURE string. For zoned decimal numbers, the decimal point can also be used in the PICTURE string. PIC 9(06)V99 COMP and PIC 9(06)V99 COMP-4 are valid PICTURE clauses for BINARY decimal data.

If the field contains DATE, TIME, or TIMESTAMP data, the type qualifier specifies the format of the data. For example, a type qualifier of *ddMMyyyy* indicates that the data is formatted as follows:

```
11122015 is December 11, 2015
```

For DATE and TIME types, all formatting options in the `java.text.SimpleDateFormat` class are supported.

For the TIMESTAMP type, the formatting option 'f' is available for nanoseconds. TIMESTAMP can contain up to nine 'f's and replaces the 'S' options for milliseconds. Instead, 'fff' indicates milliseconds of precision. An example TIMESTAMP format is as follows:

```
yyyy-mm-dd hh:mm:ss.ffffffffff
```

COBOL copybook types that map to Java data types

Because data in IMS is not strongly typed, you can use COBOL copybook types to map your IMS data to Java data types.

The following table describes how COBOL copybook types are mapped to both `DLTypeInfo` constants in the `DLIDatabaseView` class and Java data types.

Table 106. Mapping from COBOL formats to DLTypeInfo constants and Java data types

Copybook format	DLTypeInfo constant	Java data type
PIC X	CHAR	<code>java.lang.String</code>
PIC 9 BINARY ¹	See "DLTypeInfo constants and Java data types based on the PICTURE clause". ²	See "DLTypeInfo constants and Java data types based on the PICTURE clause". ²
COMP-1	FLOAT	<code>float</code>
COMP-2	DOUBLE	<code>double</code>
PIC 9 COMP-3 ³	PACKEDDECIMAL	<code>java.math.BigDecimal</code>
PIC 9 DISPLAY ⁴	ZONEDDECIMAL	<code>java.math.BigDecimal</code>

Notes:

1. Synonyms for BINARY data items are COMP and COMP-4. A PIC 9(06)V99 statement with COMP or COMP-4 is used for binary decimal data.
2. For BINARY data items, the DLTypeInfo constant and Java type depend on the number of digits in the PICTURE clause. The table "DLTypeInfo constants and Java data types based on the PICTURE clause" describes the type based on PICTURE clause length.
3. PACKED-DECIMAL is a synonym for COMP-3.
4. If the USAGE clause is not specified at either the group or elementary level, it is assumed to be DISPLAY.

The following table shows the DLTypeInfo constants and the Java data types based on the PICTURE clause.

Table 107. *DLTypeInfo* constants and Java data types based on the *PICTURE* clause

Digits in PICTURE clause	Storage occupied	DLTypeInfo constant	Java data type
1 through 2	1 byte	TINYINT UTINYINT	byte
1 through 4	2 bytes	SMALLINT USMALLINT	short
5 through 9	4 bytes	INTEGER UIINTEGER	int
10 through 18	8 bytes	BIGINT UBIGINT	long

The following table shows examples of specific copybook formats mapped to *DLTypeInfo* constants.

Table 108. *Copybook* formats mapped to *DLTypeInfo* constants

Copybook format	DLTypeInfo constant
PIC X(25)	CHAR
PIC 9(02) COMP	UTINYINT
PIC S9(04) COMP	SMALLINT
PIC 9(04) COMP	USMALLINT
PIC S9(06) COMP-4	INTEGER
PIC 9(06) COMP-4	UIINTEGER
PIC 9(06)V99 COMP or COMP-4	BINARY
PIC S9(12) BINARY	BIGINT
PIC 9(12) BINARY	UBIGINT
COMP-1	FLOAT
COMP-2	DOUBLE
PIC S9(06)V99	ZONEDDECIMAL
PIC 9(06).99	ZONEDDECIMAL
PIC S9(06)V99 COMP-3	PACKEDDECIMAL

Programming with the IMS Universal DL/I driver

Use the IMS Universal DL/I driver when you need to write granular queries to access IMS databases directly from a Java client in a non-managed environment.

Because of the fundamental differences between hierarchical databases and relational databases, sometimes the JDBC API does not provide access to the full set of IMS databases features. The IMS Universal DL/I driver is closely related to the traditional IMS DL/I database call interface that is used with other programming languages for writing applications in IMS, and provides a lower-level access to IMS database functions than the JDBC API. By using the IMS Universal DL/I driver, you can build segment search arguments (SSAs) and use the methods of the program communication block (PCB) object to read,

insert, update, delete, or perform batch operations on segments. You can gain full navigation control in the segment hierarchy.

Preparing to write a Java application with the IMS Universal drivers

Java application programs that use the IMS Universal drivers require the Java Development Kit (JDK) 7.0. Java programs that run in JMP and JBP regions require JDK 7.0 or later. Java application programs that use the IMS Universal drivers must have access to database metadata in order to interact with IMS databases. This metadata can either be accessed directly in the IMS catalog database or it can be generated as a Java metadata class with the IMS Enterprise Suite Explorer for Development.

Basic steps in writing a IMS Universal DL/I driver application

In general, to write a application program with the IMS Universal DL/I driver, you need to complete the following tasks.

To write an IMS Universal DL/I driver application, follow these steps.

1. Import the `com.ibm.ims.dli` package that contains the IMS Universal DL/I driver classes, interfaces, and methods.
2. Connect to an IMS database subsystem.
3. Obtain a program specification block (PSB), which contains one or more PCBs.
4. Obtain a PCB handle, which defines an application's view of an IMS database and provides the ability to issue database calls to retrieve, insert, update, and delete database information.
5. Obtain an unqualified segment search argument list (SSAList) of one or more segments in the database hierarchy.
6. Add qualification statements to specify the segments targeted by DL/I calls.
7. If retrieving data, mark the segment fields to be returned.
8. Execute DL/I calls to the IMS database.
9. Handle errors that are returned from the DL/I programming interface.
10. Disconnect from the IMS database subsystem.

Related tasks

[“Retrieving data in a IMS Universal DL/I driver application” on page 649](#)

The IMS Universal DL/I driver provides support for data retrieval that mirrors DL/I semantics.

Related reference

[“Generating the runtime Java metadata class” on page 567](#)

To connect to an IMS database using the IMS Universal drivers, you need to include on your Java classpath the Java metadata class that provides the database view.

Java packages for IMS Universal DL/I driver support

Before you can invoke IMS Universal DL/I driver methods, you must access all or parts of various Java packages that contain those methods.

You can do that by either importing the packages or specific classes, or by using the fully-qualified class names. You might need the following packages or classes for your IMS Universal DL/I driver application:

`com.ibm.ims.dli`

Contains the core classes, interfaces, and methods for the IMS Universal DL/I driver.

`com.ibm.ims.base`

Contains exception classes for errors that are returned by DL/I or IMS.

Connecting to an IMS database by using the IMS Universal DL/I driver

Before you can execute DL/I calls from your IMS Universal DL/I driver application, you must connect to an IMS database.

The IMS Universal DL/I driver application can establish a connection to an IMS database using the PSB interface, which is part of the `com.ibm.ims.dli` package. Pass the connection properties using an `IMSConnectionSpec` instance.

To connect to an IMS database by using the IMS Universal DL/I driver:

1. Create an `IMSConnectionSpec` instance by calling the `createIMSConnectionSpec` method in the `IMSConnectionSpecFactory` class.
2. Set the following connection properties for the `IMSConnectionSpec` instance.

DatastoreName

The name of the IMS data store to access.

When using type-4 connectivity, the **DatastoreName** property must match either the name of the data store defined to ODBM or be blank. The data store name is defined in the ODBM CSLDCxxx PROCLIB member using either the `DATASTORE(NAME=name)` or `DATASTORE(NAME=name, ALIAS(NAME=aliasname))` parameter. If an alias is specified, you must specify the *aliasname* as the value of the **datastoreName** property. If the **DatastoreName** value is left blank (or not supplied), IMS Connect connects to any available instance of ODBM as it is assumed that data sharing is enabled between all datastores defined to ODBM.

When using type-2 connectivity, set the **DatastoreName** property to the IMS subsystem alias. This is not required to be set for the Java Dependent Region run time.

DatabaseName

The location of the database metadata representing the target IMS database.

The **DatabaseName** property can be specified in one of two ways, depending on whether the metadata is stored in the IMS catalog or as a static metadata class generated by the IMS Enterprise Suite Explorer for Development.

If your IMS system uses the IMS catalog, the **DatabaseName** property is the name of the PSB that your application uses to access the target IMS database.

If you are using the IMS Explorer for Development, the **databaseName** property is the fully qualified name of the Java metadata class generated by the IMS Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **DatabaseName** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

MetadataURL

The location of the database metadata representing the target IMS database.

This property is deprecated. Use **DatabaseName** instead.

The **MetadataURL** property is the fully qualified name of the Java metadata class generated by the IMS Enterprise Suite Explorer for Development. The URL must be prefixed with `class://` (for example, `class://com.foo.BMP255DatabaseView`).

In a J2C Connection Factory environment, the **MetadataURL** property can be overridden for an individual connection without affecting the default value specified for the resource adapter.

PortNumber

The TCP/IP server port number to be used to communicate with IMS Connect. The port number is defined using the `DRDAPORT` parameter on the `ODACCESS` statement in the IMS Connect configuration PROCLIB member. The default port number is 8888. Do not set this property when using type-2 connectivity.

DatastoreServer

The name or IP address of the data store server (IMS Connect). You can provide either the host name (for example, dev123.svl.ibm.com) or the IP address (for example, 192.166.0.2). Do not set this property when using type-2 connectivity.

DriverType

The type of driver connectivity to use (value must be `IMSConnectionSpec.DRIVER_TYPE_4` for type-4 connectivity or `IMSConnectionSpec.DRIVER_TYPE_2` for type-2 connectivity).

allMetadata

Optional. When this property is set to `true`, the `DatabaseMetadata` interface returns information for all resources in the IMS catalog. When the property is set to `false`, the `DatabaseMetadata` interface returns information for the allocated PSB. The default value for this property is `false`.

initStatusGroup

Optional. When a connection is made and the PSB is allocated, this property will indicate that the driver should automatically issue an `INIT STATUS GROUPA` or `INIT STATUS GROUPB` if a value of 'A' or 'B' is provided. The default will not issue an `INIT STATUS GROUP` call.

sslConnection

Optional. Indicates if this connection uses Secure Sockets Layer (SSL) for data encryption. Set this property to "true" to enable SSL, or to "false" otherwise. Do not set this property when using type-2 connectivity.

sslKeyStoreType

Optional. Specifies the format of the file that contains cryptographic objects needed to establish a secure socket connection. The valid values are "JKS" and "PKCS12". This value is only used when **sslConnection** is set to "true" and **sslKeyStoreType** is not specified. The **sslKeyStoreType** parameter defaults to "JKS".

sslSecureSocketProtocol

Optional. Specifies the cryptographic communication protocol for the new connection. Specify a protocol that is supported by the server and provides the highest level of security. The valid values are "SSL", "SSLv3", "TLSv1.1", and "TLSv1.2". This value is only used when **sslConnection** is set to "true". If **sslConnection** is set to "true" and **sslSecureSocketProtocol** is not specified, a default protocol will be determined at runtime by the JRE and the server.

sslTrustStoreLocation

Optional. Specifies the location of the cryptographic trust store file for the new connection. This value is only used when **sslConnection** is set to true.

sslTrustStorePassword

Optional. Specifies the password to access the cryptographic trust store file. This value is only used when **sslConnection** is set to true.

sslKeyStoreLocation

Optional. Specifies the location of the cryptographic key store file for the new connection. This value is only used when **sslConnection** is set to true.

sslKeyStorePassword

Optional. Specifies the password to access the cryptographic key store file. This value is only used when **sslConnection** is set to true.

loginTimeout

Optional. Specifies the number of seconds that the driver waits for a response from the server before timing out a connection initialization or server request. Set this property to a non-negative

integer for the number of seconds. Set this property to 0 for an infinite timeout length. Do not set this property when using type-2 connectivity.

user

The user name for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

password

The password for the connection to IMS Connect provided by your RACF administrator. Do not set this property when using type-2 connectivity.

dbViewLocation

Optional. Specifies the fully qualified path to a databaseView metadata class. You can use this property to include a metadata class that is not located in your project path.

treatInvalidDecimalAsNull

Optional. Indicates whether to interpret certain Decimal values that appear invalid in Java applications (such as PACKEDDECIMAL and ZONEDDECIMAL with invalid sign bits) as null. By default, this property is "false", and a conversion exception is thrown when the Java applications are processing invalid values.

3. Pass the connection request properties to the PSBFactory class to create the PSB instance.
When the PSB instance is created successfully, a connection is established to the database.
4. When you are finished with a connection to the IMS database from a IMS Universal DL/I driver application, you must close the connection to the database by calling the `close` method on the PSB instance.

Example: type-4 Connection

The following code example shows how to create a type-4 connection to an IMS database from your IMS Universal DL/I driver application:

```
IMSConnectionSpec connSpec = IMSConnectionSpecFactory.createIMSConnectionSpec();
connSpec.setDatastoreName("SYS1");
connSpec.setDatastoreServer("9.876.543.21");
connSpec.setPortNumber(8888);
connSpec.setDatabaseName("class://testdb.jdbo.HospitalDatabaseView");
connSpec.setSSLConnection(true);
connSpec.setLoginTimeout(10);
connSpec.setUser("usr");
connSpec.setPassword("usripwd");
connSpec.setDriverType(IMSConnectionSpec.DRIVER_TYPE_4);
PSB psb = PSBFactory.createPSB(connSpec);
```

Related tasks

[“Configuring the IMS Universal drivers for SSL support” on page 668](#)

With type-4 connectivity, the IMS Universal drivers provide support for the Secure Sockets Layer (SSL) through the Java Secure Socket Extension (JSSE).

IMS Universal DL/I driver interfaces for executing DL/I operations

In a traditional IMS application, you make DL/I calls to insert, update, delete, or retrieve data. To perform the same functions in a IMS Universal DL/I driver application, you invoke methods.

Methods are defined in the following interfaces:

- The program specification block (PSB) interface is used to connect to IMS databases. Use the PSB interface to obtain a handle to any program communication block (PCB) that is contained in the PSB. The PCB handle is used to access the particular database that is referenced by the PCB.
- The PCB interface represents a cursor position in an IMS database. The PCB interface supports DL/I message call functions, including Get Unique (GU), Get Next (GN), Get Next Within Parent (GNP), Insert (ISRT), Replace (REPL), and Delete (DLET). The PCB interface can obtain an unqualified list of segment

search arguments and perform batch retrieve, update, and delete operations. You can also use the PCB interface to return the application interface block (AIB) that is associated with the most recent DL/I call.

- The `SSAList` interface represents a list of segment search arguments (SSAs) used to specify the segments to target in a particular database call. Use the `SSAList` interface to construct the SSAs, and to set the command codes and lock class for the SSAs. You can set an initial qualification statement and append additional qualifiers, based on the values of the segment fields, to restrict which segments to target in the DL/I call. You can also specify which fields to return from a database retrieve call.
- The `Path` interface represents a database record for the purpose of a DL/I retrieval or update operation. The `Path` interface can be viewed as the concatenation of all of the segment instances in a specific database hierarchic path, starting from the highest level segment that is nearest the root segment to the lowest level segment. Use the `Path` interface to set or retrieve the value of any segment field that is located in the hierarchic path.
- The `PathSet` interface provides access to a collection of `Path` objects that are returned by a batch retrieve operation.
- The `AIB` interface and the database PCB (DBPCB) interface return useful information that was returned by IMS as a result of a DL/I call.
- The `GSAMPCB` interface represents a GSAM PCB and is essentially a cursor position in a GSAM database. This interface provides data access to GSAM databases with calls that are similar to DL/I calls.
- The `RSA` interface represents a GSAM database record search argument that is the key to a cursor position in the GSAM database.

Specifying segment search arguments using the `SSAList` interface

The `SSAList` interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Use the `SSAList` interface to construct each segment search argument (SSA) in the list and to set the command codes and lock class for the SSAs. Each SSA in the `SSAList` can be unqualified or qualified.

In addition, your application can specify which segment fields are to be returned from a database retrieve call by using the `markFieldForRetrieval` or the `markAllFieldsForRetrieval` methods. Following the IMS default, all of the fields in the lowest level segment specified by the `SSAList` are initially marked for retrieval.

- For non-batch DL/I data retrieval or update operations, use the `getPathForRetrieveReplace` method.
- For a DL/I insert call, use the `getPathForInsert` method.
- For a batch update operation, use the `getPathForBatchUpdate` method.

The following examples demonstrate how to specify segment search arguments using the `SSAList` interface. The examples are based on the Hospital database.

Creating an unqualified `SSAList`

This example returns a `Path` that consists of all fields in the segment "DOCTOR":

```
SSAList ssaList = pcb.getSSAList("HOSPITAL","DOCTOR");
Path path = ssaList.getPathForRetrieveReplace();
pcb.getUnique(path, ssaList, false);
```

In the previous example, the `ssaList` represents all segments along the hierarchic path from the topmost segment ("HOSPITAL") to the lowest segment ("DOCTOR"). The `ssaList` will look like this:

```
HOSPITALb
WARDbbbb
PATIENTbb
ILLNESSbb
TREATMNTb
DOCTORbbb
```

Creating a qualified SSAList

An SSAList can be qualified to filter the segments on the hierarchic path to be retrieved or updated. The general steps to create a qualified SSAList are:

1. Obtain an unqualified SSAList from the PCB using the `getSSAList` method.
2. Use the `addInitialQualification` method to specify the initial search criteria for a segment on the SSAList returned from a `getSSAList` method. For each segment represented in the SSAList, you can make one call to specify an initial qualification for that segment. The segment can be referenced by name or by using the 1-based offset of the SSA representing that segment within the SSAList. If you use more than one `addInitialQualification` statement for a segment, an exception will be thrown. The relational operator (**relationalOp**) parameter in the `addInitialQualification` method indicates the conditional criteria that the segment must meet in order to be qualified. Valid relational operators are:
 - EQUALS
 - GREATER_OR_EQUAL
 - GREATER_THAN
 - LESS_OR_EQUAL
 - LESS_THAN
 - NOT_EQUAL
3. To specify additional search criteria, use the `appendQualification` method. For each segment, you can make multiple calls to the `appendQualification` method to add more than one qualification statement. The Boolean operator (**booleanOp**) parameter in the `appendQualification` method indicates how this qualification is logically connected to the previous qualification. Valid Boolean operators are:
 - AND
 - OR
 - INDEPENDENT_AND
4. You can also qualify a SSAList by setting DL/I command codes and lock classes. The supported DL/I command codes include:
 - CC_A: The A command code (clear positioning).
 - CC_C: The C command code (concatenated key). Use the `addConcatenatedKey` method to add a concatenated key to a segment.
 - CC_D: The D command code (path call)
 - CC_F: The F command code (first occurrence)
 - CC_G: The G command code (prevent randomization).
 - CC_L: The L command code (last occurrence)
 - CC_N: The N command code (path call ignore)
 - CC_O: The O command code (contain field names or segment position and length).
 - CC_P: The P command code (set parentage)
 - CC_U: The U command code (maintain position at this level)
 - CC_V: The V command code (maintain position at this level and all superior levels)

You can use a lock class to prevent another program from updating a segment until your program reaches a commit point. Use the `addLockClass` method to add a lock class to a segment. The supported lock class letters are "A" to "J". The behavior of a lock class is the same as using a "Q" command code with that lock class letter.

The following code example demonstrates how to specify and use a qualified SSAList with a single initial qualification statement to retrieve data:

```
SSAList ssaList = pcb.getSSAList("HOSPITAL","DOCTOR");
ssaList.addInitialQualification("PATIENT","PATNAME",SSAList.EQUALS,"ANDREA SMITH");
ssaList.markFieldForRetrieval("ILLNESS","ILLNAME",true);
ssaList.markFieldForRetrieval("TREATMNT","TREATMNT",true);
Path path = ssaList.getPathForRetrieveReplace();
pcb.getUnique(path, ssaList, false);
```

For the code example above, the ssaList will look like this:

```
HOSPITALb
WARDbbbb
PATIENTb(PATNAMEbEQANDREAbSMITHbbbb)
ILLNESSb*D
TREATMNT*D
```

The code example above retrieves the following information for all records where PATNAME is "ANDREA SMITH":

- The ILLNAME field in the ILLNESS segment.
- The TREATMNT field in the TREATMNT segment.
- All fields in the DOCTOR segment (by default, IMS returns all fields in the lowest level segment specified by the SSAList).

The following code example demonstrates how to specify a qualified SSAList with a multiple qualification statements to retrieve data:

```
SSAList ssaList = pcb.getSSAList("HOSPITAL","WARD");
ssaList.addInitialQualification("WARD","NURCOUNT",SSAList.GREATER_THAN,4);
ssaList.appendQualification("WARD",SSAList.AND,"DOCCOUNT",SSAList.GREATER_THAN, 2);
```

For the code example above, the ssaList will look like this:

```
HOSPITALb
WARDbbbb(NURCOUNTGT4&DOCCOUNTGT2;)
```

The following example shows how to specify a qualified SSAList with the command code CC_L (which means "last occurrence") to find the most recently admitted patient in the "SANTA TERESA" hospital:

```
SSAList ssaList = pcb.getSSAList("HOSPITAL","PATIENT");
ssaList.addInitialQualification
    ("HOSPITAL","HOSPNAME",SSAList.EQUALS,"SANTA TERESA");
ssaList.addCommandCode("PATIENT",SSAList.CC_L);
Path path = ssaList.getPathForRetrieveReplace();
pcb.getUnique(path, ssaList, false);
```

For the code example above, the ssaList will look like this:

```
HOSPITAL (HOSPNAMEEQSANTAbTERESAbbbbb)
WARDbbbb
PATIENTb*L
```

Debugging an SSAList

If you need to debug by identifying whether the segment search arguments in your SSAList are correct, use the buildSSAListInBytes method to build the SSAList in DL/I format for further debugging:

```
byte[][] ssaListInBytes = ssaList.buildSSAListInBytes();
```

You can iterate over each segment search argument in the byte array that is returned and print it out to make sure the segment search argument is what it should be.

Related concepts

[“Segment search arguments \(SSAs\)” on page 168](#)

Segment search arguments (SSAs) specify information for IMS to use in processing a DL/I call. Regardless of the datatype for the field specified in a SSA, the SSA treats the field as a binary type and does a binary comparison.

Related reference

[“SSA coding formats” on page 223](#)

Use the following formats to code segment search arguments in assembler language, C language, COBOL, Pascal, and PL/I.

[“Hospital database example” on page 568](#)

The code examples for the IMS Universal drivers application programming topics use the Hospital database.

[Command code reference \(Application Programming APIs\)](#)

Retrieving data in a IMS Universal DL/I driver application

The IMS Universal DL/I driver provides support for data retrieval that mirrors DL/I semantics.

The following are the general steps to retrieve segments from the database:

1. Obtain an `SSAList` instance from the PCB instance representing the database.
2. Optionally, you can add qualification statements to the `SSAList` instance.
3. Specify the segment fields to retrieve. Use the `markFieldForRetrieval` method to mark a single field, or use the `markAllFieldsForRetrieval` method to mark all the fields for a segment. Following the IMS default, all of the fields in the lowest-level segment specified by the `SSAList` instance are initially marked for retrieval. When one or more of the fields in the lowest level segment specified in an `SSAList` instance is marked for retrieval, only the explicitly marked fields are retrieved.
4. Get a `Path` instance by using the `SSAList` instance from the previous steps and calling the `getPathForRetrieveReplace` method. When a retrieve call is made, the resulting `Path` object will contain all the fields that have been marked for retrieval.
5. Call a DL/I retrieve operation using one of the following methods from the PCB interface:

Java API for DL/I retrieve method	Usage
<code>getUnique</code>	Retrieves a specific unique segment. This method provides the same functionality as the DL/I Get Unique (GU) database call. If the <code>isHoldCall</code> parameter is set to true, the call behaves as a DL/I Get Hold Unique (GHU) database call.
<code>getNext</code>	Retrieves the next segment in a <code>Path</code> . This method provides the same functionality as the DL/I Get Next (GN) database call. If the <code>isHoldCall</code> parameter is set to true, the call behaves as a DL/I Get Hold Next (GHN) database call.
<code>getNextWithinParent</code>	Retrieves the next segment within the same parent. This method provides the same functionality as the DL/I Get Next Within Parent (GNP) database call. If the <code>isHoldCall</code> parameter is set to true, the call behaves as a DL/I Get Hold Next Within Parent (GHNP) database call.
<code>batchRetrieve</code>	Retrieves multiple segments with a single call. See "Batch data retrieval in a Java API for DL/I application" for more information about how to use this method.

6. Read the values of the retrieved fields out of the `Path` object after the retrieve call is made.

IMS Universal DL/I driver data retrieval example

The following code fragment illustrates how to use the `getUnique` method and `getNext` method to retrieve the hospital name (HOSPNAME), ward name (WARDNAME), patient count (PATCOUNT), nurse count (NURCOUNT), and doctor count (DOCCOUNT) fields from the Hospital database:

```
import com.ibm.ims.dli.*;

public class HospitalDLIReadClient {

    public static void main(String[] args) {
        PSB psb = null;
        PCB pcb = null;
        SSAList ssalist = null;
        Path path = null;
        PathSet pathSet = null;

        try {
            // establish a database connection
            IMSConnectionSpec connSpec
                = IMSConnectionFactory.createIMSConnectionSpec();
            connSpec.setDatastoreName("IMS1");
            connSpec.setDatastoreServer("ecdev123.svl.ibm.com");
            connSpec.setPortNumber(5555);
            connSpec.setMetadataURL("class://BMP266.BMP266DatabaseView");
            connSpec.setUser("usr");
            connSpec.setPassword("password");
            connSpec.setDriverType(IMSConnectionSpec.DRIVER_TYPE_4);
            psb = PSBFactory.createPSB(connSpec);
            System.out.println("**** Created a connection to the IMS database");

            pcb = psb.getPCB("PCb01");
            System.out.println("**** Created PCB object");

            // specify the segment search arguments
            ssalist = pcb.getSSAList("HOSPITAL", "WARD");
            // add the initial qualification
            ssalist.addInitialQualification("HOSPITAL", "HOSPCODE",
                SSAList.GREATER_OR_EQUAL, 444);
            // specify the fields to retrieve
            ssalist.markFieldForRetrieval("HOSPITAL", "HOSPNAME", true);
            ssalist.markAllFieldsForRetrieval("WARD", true);
            ssalist.markFieldForRetrieval("WARD", "WARDNO", false);
            System.out.println("**** Created SSAList object");

            // obtain a Path containing the segments that match the SSAList criteria
            path = ssalist.getPathForRetrieveReplace();
            System.out.println("**** Created Path object");

            // issue a DL/I GU call to retrieve the first segment on the Path
            if (pcb.getUnique(path, ssalist, true) {
                System.out.println("HOSPNAME: " + path.getString("HOSPITAL", "HOSPNAME"));
                System.out.println("WARDNAME: " + path.getString("WARD", "WARDNAME"));
                System.out.println("PATCOUNT: " + path.getInt("WARD", "PATCOUNT"));
                System.out.println("NURCOUNT: " + path.getInt("WARD", "NURCOUNT"));
                System.out.println("DOCCOUNT: " + path.getShort("WARD", "DOCCOUNT"));
            }

            // issue multiple DL/I GN calls until there are no more segments to retrieve
            while (pcb.getNext(pat, ssalist, true) {
                System.out.println("HOSPNAME: " + path.getString("HOSPITAL", "HOSPNAME"));
                System.out.println("WARDNAME: " + path.getString("WARD", "WARDNAME"));
                System.out.println("PATCOUNT: " + path.getInt("WARD", "PATCOUNT"));
                System.out.println("NURCOUNT: " + path.getInt("WARD", "NURCOUNT"));
                System.out.println("DOCCOUNT: " + path.getShort("WARD", "DOCCOUNT"));
            }

            // close the database connection
            psb.close();
            System.out.println("**** Disconnected from IMS database");

        } catch (DLIException e) {
            System.out.println(e);
            System.exit(0);
        }
    }
}
```

Related concepts

[“Specifying segment search arguments using the SSAList interface” on page 646](#)

The `SSAList` interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Related tasks

[“Basic steps in writing a IMS Universal DL/I driver application” on page 642](#)

In general, to write a application program with the IMS Universal DL/I driver, you need to complete the following tasks.

[“Batch data retrieval in a IMS Universal DL/I driver application” on page 651](#)

Use the `batchRetrieve` method to retrieve multiple segments in a single call.

Related reference

[“Methods for retrieving and converting data types” on page 636](#)

With the IMS Universal JDBC driver, you can use the `ResultSet` interface (`java.sql.ResultSet`) to retrieve and convert the data from the type that is defined in the database metadata to the type that is required by your Java application. Similarly, with the IMS Universal DL/I driver, you can use the `Path` interface to perform data retrieval and conversion to Java data types.

Batch data retrieval in a IMS Universal DL/I driver application

Use the `batchRetrieve` method to retrieve multiple segments in a single call.

Instead of the client application making multiple GU and GN calls, IMS will perform all of the GU and GN processing and will deliver the results back to the client in a single batch network operation. The fetch size property determines how much data is sent back on each batch network operation.

To perform a batch data retrieval operation:

1. Obtain an `SSAList` instance from the PCB instance that represents the database.
2. Optionally, you can add qualification statements to the `SSAList` instance.
3. Specify the segment fields to retrieve. Use the `markFieldForRetrieval` method to mark a single field, or use the `markAllFieldsForRetrieval` method to mark all the fields for a segment. Following the IMS default, all of the fields in the lowest-level segment specified by the `SSAList` instance are initially marked for retrieval.
4. Optionally, set the fetch size property.
The fetch size gives a hint to the IMS Universal DL/I driver as to the number of records to fetch from the database in a single batch operation. See "Improving query performance by setting fetch size" for more information.
5. Call the `batchRetrieve` method with the `SSAList` instance above as a parameter.
The `batchRetrieve` method returns a `PathSet` that contains a list of records that satisfy the criteria specified by the `SSAList`.
6. Read the values of the retrieved fields out of the `Path` object after the retrieve call is made.

IMS Universal DL/I driver batch data retrieval example

The following code fragment illustrates how to use the `batchRetrieve` method to retrieve the hospital name (HOSPNAME), ward name (WARDNAME), patient count (PATCOUNT), nurse count (NURCOUNT), and doctor count (DOCCOUNT) fields from the Hospital database:

```
import com.ibm.ims.dli.*;

public class HospitalDLIReadClient {

    public static void main(String[] args) {
        PSB psb = null;
        PCB pcb = null;
        SSAList ssalist = null;
        Path path = null;
        PathSet pathSet = null;
```

```

try {
    // establish a database connection
    IMSConnectionSpec connSpec
        = IMSConnectionSpecFactory.createIMSConnectionSpec();
    connSpec.setDatastoreName("IMS1");
    connSpec.setDatastoreServer("ecdev123.svl.ibm.com");
    connSpec.setPortNumber(5555);
    connSpec.setMetadataURL("class://BMP266.BMP266DatabaseView");
    connSpec.setUser("usr");
    connSpec.setPassword("password");
    connSpec.setDriverType(IMSConnectionSpec.DRIVER_TYPE_4);

    psb = PSBFactory.createPSB(connSpec);
    System.out.println("**** Created a connection to the IMS database");

    pcb = psb.getPCB("PCb01");
    System.out.println("**** Created PCB object");

    // specify the segment search arguments
    ssalist = pcb.getSSAList("HOSPITAL", "WARD");
    // add the initial qualification
    ssalist.addInitialQualification("HOSPITAL", "HOSPCODE",
        SSAList.GREATER_OR_EQUAL, 444);
    // specify the fields to retrieve
    ssalist.markFieldForRetrieval("HOSPITAL", "HOSPNAME", true);
    ssalist.markAllFieldsForRetrieval("WARD", true);
    ssalist.markFieldForRetrieval("WARD", "WARDNO", false);
    System.out.println("**** Created SSAList object");

    // issue the database call to perform a batch retrieve operation
    pathSet = pcb.batchRetrieve(ssalist);
    System.out.println("**** Batch Retrieve returned without exception");
    System.out.println("**** Created PathSet object");

    while(pathSet.hasNext()){
        path = pathSet.next();

        System.out.println("HOSPNAME: " + path.getString("HOSPITAL", "HOSPNAME"));
        System.out.println("WARDNAME: " + path.getString("WARD", "WARDNAME"));
        System.out.println("PATCOUNT: " + path.getInt("WARD", "PATCOUNT"));
        System.out.println("NURCOUNT: " + path.getInt("WARD", "NURCOUNT"));
        System.out.println("DOCCOUNT: " + path.getShort("WARD", "DOCCOUNT"));
    }
    System.out.println("**** Fetched all rows from PathSet");

    // close the database connection
    psb.close();
    System.out.println("**** Disconnected from IMS database");
} catch (DLIException e) {
    System.out.println(e);
    System.exit(0);
}
}
}

```

Related concepts

[“Specifying segment search arguments using the SSAList interface” on page 646](#)

The `SSAList` interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Related reference

[“Methods for retrieving and converting data types” on page 636](#)

With the IMS Universal JDBC driver, you can use the `ResultSet` interface (`java.sql.ResultSet`) to retrieve and convert the data from the type that is defined in the database metadata to the type that is required by your Java application. Similarly, with the IMS Universal DL/I driver, you can use the `Path` interface to perform data retrieval and conversion to Java data types.

Improving query performance by setting fetch size

You can optimize query performance by setting the number of records to retrieve in batch retrieval mode.

In the IMS Universal DL/I driver, a list of rows is represented by a `Path` instance containing one or more segments that match the segment search argument criteria specified by an `SSAList`. The *fetch size* is the number of rows physically retrieved from the IMS database per network call. This is set for you internally. You can also set the fetch size using the `setFetchSize` method from the `PCB` interface. Setting the

fetch size allows a single request to return multiple rows at a time, so that each application request to retrieve the next row does not always result in a network request. If the fetch size was n and the IMS Universal DL/I driver application requires more than the previous n number of rows during a batch retrieve operation, another network call will be made on behalf of the application to retrieve the next n number of rows that match the segment search argument criteria.

Methods for retrieving and converting data types

With the IMS Universal JDBC driver, you can use the `ResultSet` interface (`java.sql.ResultSet`) to retrieve and convert the data from the type that is defined in the database metadata to the type that is required by your Java application. Similarly, with the IMS Universal DL/I driver, you can use the `Path` interface to perform data retrieval and conversion to Java data types.

The following table shows the available get methods in the `ResultSet` interface (for the IMS Universal JDBC driver) or the `Path` interface (for the IMS Universal DL/I driver) for accessing data of a certain Java data type.

The "No Truncation or Data Loss" column indicates the data types that are designed to be accessed with the given `getXXX` method. No truncation or data loss occurs when using those methods for those data types. The data types that are in the "Legal without Data Integrity" column are all other legal calls; however, data integrity cannot be ensured when using the given `getxxx` method to access those data types. If a data type is not in either column, using the given `getXXX` method for that data type will result in an exception.

Table 109. `ResultSet.getXXX` and `Path.getXXX` methods to retrieve data types

ResultSet.getXXX Method or Path.getXXX Method	Data Type (any not listed result in an exception)	
	No Truncation or Data Loss	Legal without Data Integrity
getBytes	TINYINT UTINYINT	SMALLINT INTEGER BIGINT FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getShort	SMALLINT USMALLINT	TINYINT INTEGER BIGINT FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹

Table 109. *ResultSet.getXXX* and *Path.getXXX* methods to retrieve data types (continued)

ResultSet.getXXX Method or Path.getXXX Method	Data Type (any not listed result in an exception)	
	No Truncation or Data Loss	Legal without Data Integrity
getInt	INTEGER UNSIGNED_INTEGER	TINYINT SMALLINT BIGINT FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getLong	BIGINT UNSIGNED_BIGINT	TINYINT SMALLINT INTEGER FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getFloat	FLOAT	TINYINT SMALLINT INTEGER BIGINT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getDouble	DOUBLE	TINYINT SMALLINT INTEGER BIGINT FLOAT BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹

Table 109. *ResultSet.getXXX* and *Path.getXXX* methods to retrieve data types (continued)

ResultSet.getXXX Method or Path.getXXX Method	Data Type (any not listed result in an exception)	
	No Truncation or Data Loss	Legal without Data Integrity
getBoolean	BIT	TINYINT SMALLINT INTEGER BIGINT FLOAT DOUBLE CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹
getString	CHAR VARCHAR	TINYINT SMALLINT INTEGER BIGINT FLOAT DOUBLE BIT PACKEDDECIMAL ¹ ZONEDDECIMAL ¹ BINARY DATE TIME TIMESTAMP
getBigDecimal	BINARY ³ PACKEDDECIMAL ¹ ZONEDDECIMAL ¹	TINYINT SMALLINT INTEGER BIGINT FLOAT DOUBLE BIT CHAR VARCHAR
getClob	CLOB ²	all others result in an exception
getBytes	BINARY	SMALLINT INTEGER BIGINT FLOAT DOUBLE BIT CHAR VARCHAR PACKEDDECIMAL ¹ ZONEDDECIMAL ¹ DATE TIME TIMESTAMP

Table 109. *ResultSet.getXXX* and *Path.getXXX* methods to retrieve data types (continued)

ResultSet.getXXX Method or Path.getXXX Method	Data Type (any not listed result in an exception)	
	No Truncation or Data Loss	Legal without Data Integrity
getDate	DATE	CHAR VARCHAR TIMESTAMP
getTime	TIME	CHAR VARCHAR TIMESTAMP
getTimestamp	TIMESTAMP	CHAR VARCHAR DATE TIME

Note:

1. PACKEDDECIMAL and ZONEDDECIMAL are data type extensions for IMS Universal JDBC driver and the IMS Universal DL/I driver. All other types are standard SQL types defined in SQL92. **Restriction:** PACKEDDECIMAL and ZONEDDECIMAL data types do not support the Sign Leading or Sign Separate modes. For these two data types, sign information is always stored with the Sign Trailing method.
2. The CLOB data type is supported only for the retrieval and storage of XML data.
3. The BINARY data type is valid only for decimal data used with a binary type converter.

If the field type is either PACKEDDECIMAL or ZONEDDECIMAL, the type qualifier is the COBOL PICTURE string that represents the layout of the field. All COBOL PICTURE strings that contain valid combinations of 9s, Ps, Vs, and Ss are supported. Expansion of PICTURE strings is handled automatically. For example, '9(5)' is a valid PICTURE string. For zoned decimal numbers, the decimal point can also be used in the PICTURE string. PIC 9(06)V99 COMP and PIC 9(06)V99 COMP-4 are valid PICTURE clauses for BINARY decimal data.

If the field contains DATE, TIME, or TIMESTAMP data, the type qualifier specifies the format of the data. For example, a type qualifier of *ddMMyyyy* indicates that the data is formatted as follows:

```
11122015 is December 11, 2015
```

For DATE and TIME types, all formatting options in the `java.text.SimpleDateFormat` class are supported.

For the TIMESTAMP type, the formatting option 'f' is available for nanoseconds. TIMESTAMP can contain up to nine 'f's and replaces the 'S' options for milliseconds. Instead, 'fff' indicates milliseconds of precision. An example TIMESTAMP format is as follows:

```
yyyy-mm-dd hh:mm:ss.ffffffffff
```

Creating and inserting data in a IMS Universal DL/I driver application

Use the `create` or the `insert` methods in the PCB interface to add a new segment to the database.

In the IMS Universal DL/I driver, the `insert` and `create` methods provide functionality similar to the DL/I ISRT call. The `insert` methods will return an IMS status code indicating the results of the DL/I operation, whereas the `create` method returns the number of segments created (this will always return 1). An exception is thrown if a key field is not set.

The following are the general steps to add a new segment to the database:

1. Obtain an SSAList instance from the PCB instance representing the database.
2. Optionally, you can add qualification statements to the SSAList.
3. Get a Path instance by using the SSAList instance from the previous steps and calling the `getPathForInsert` method. The `getPathForInsert` method takes the name of an existing segment on the SSAList as a parameter. The parameter indicates the name of the segment type for the new segment. For instance, to add a new patient segment, you would pass the segment name `PATIENT` as the parameter.
4. Using the Path instance from the step above, set the field values for the new segment.
5. Call the `insert` or `create` method to add the new segment.

IMS Universal DL/I driver create and insert example

The following code fragment illustrates how to use the `create` and `insert` methods to add a new patient and illness segment in the database where the hospital name is "SANTA TERESA" and the ward name is "GENERAL".

```
SSAList ssaList = pcb.getSSAList("HOSPITAL","PATIENT");
ssaList.addInitialQualification("HOSPITAL","HOSPNAME",
SSAList.EQUALS,"SANTA TERESA");
ssaList.addInitialQualification("WARD","WARDNAME",
SSAList.EQUALS,"GENERAL");

Path path = ssaList.getPathForInsert("PATIENT");
path.setString("PATIENT", "PATNUM", "0088");
path.setString("PATIENT", "PATNAME", "JACK KIRBY");
int i = pcb.create(path, ssaList); // returns i = 1 if successful
System.out.println(i);

SSAList ssaList2 = pcb.getSSAList("HOSPITAL","ILLNESS");
ssaList2.addInitialQualification("HOSPITAL","HOSPNAME",
    SSAList.EQUALS,"SANTA TERESA");
ssaList2.addInitialQualification("WARD","WARDNAME",
    SSAList.EQUALS,"GENERAL");
ssaList2.addInitialQualification("PATIENT","PATNUM",
    SSAList.EQUALS,"0088");

Path path2 = ssaList2.getPathForInsert("ILLNESS");
path2.setString("ILLNAME", "APPENDICITIS");
short status = pcb.insert(path2, ssaList2);
```

The following code example shows another way to do this in a single call:

```
SSAList ssaList = pcb.getSSAList("HOSPITAL","ILLNESS");
ssaList.addInitialQualification("HOSPITAL","HOSPNAME",
    SSAList.EQUALS,"SANTA TERESA");
ssaList.addInitialQualification("WARD","WARDNAME",
    SSAList.EQUALS,"GENERAL");
ssaList.addCommandCode("PATIENT", SSAList.CC_D);

Path path = ssaList.getPathForInsert("PATIENT");
path.setString("PATIENT", "PATNUM", "0088");
path.setString("PATIENT", "PATNAME", "JACK KIRBY");
path.setString("ILLNAME", "APPENDICITIS");

int i = pcb.create(path, ssaList); // returns i = 1 if successful
```

Important: To persist changes made to the database, your application must call the `PSB.commit` method prior to deallocating the PSB, otherwise the changes are rolled back up to the last point `commit` was called.

Related concepts

[“Specifying segment search arguments using the SSAList interface” on page 646](#)

The `SSAList` interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Updating data in a IMS Universal DL/I driver application

Use the `replace` methods in the PCB interface to update an existing segment in the database.

In the IMS Universal DL/I driver, the `replace` methods provide functionality similar to the DL/I REPL call. The `replace` methods will return an IMS status code indicating the results of the DL/I operation.

The following are the general steps to update an existing segment in the database:

1. Obtain an `SSAList` instance from the PCB instance representing the database.
2. Optionally, you can add qualification statements to the `SSAList` instance. See "Specifying segment search arguments using the `SSAList` interface" for more information.
3. Get a `Path` instance by using the `SSAList` instance from the previous steps and calling the `getPathForRetrieveReplace` method.
4. Using the `Path` instance from the step above, set the field values to update for the segment.
5. Perform a Hold operation before issuing the `replace` call. The Hold operation can be a `getUnique`, `getNext`, or `getNextWithinParent` method call.
6. Call the `replace` method to update the segment.

IMS Universal DL/I driver update example

The following code fragment illustrates how to use the `replace` method to update a patient's name in patient records where the patient name is "ANDREA SMITH", the ward name is "SURG", and the hospital name is "ALEXANDRIA".

```
SSAList ssaList = pcb.getSSAList("HOSPITAL","PATIENT");
ssaList.addInitialQualification("HOSPITAL","HOSPNAME",SSAList.EQUALS,"ALEXANDRIA");
ssaList.addInitialQualification("WARD","WARDNAME",SSAList.EQUALS,"SURG");
ssaList.addInitialQualification("PATIENT","PATNAME",SSAList.EQUALS,"ANDREA SMITH");

Path path = ssaList.getPathForRetrieveReplace();
if(pcb.getUnique(path, ssaList, true)){
    path.setString("PATNAME", "ANDREA TAYLOR");
    pcb.replace(path);
}
while(pcb.getNext(path, ssaList, true){
    path.setString("PATNAME", "ANDREA TAYLOR");
    pcb.replace(path);
}
```

Note: To persist changes made to the database, your application must call the `commit` method prior to deallocating the PSB, otherwise the changes are rolled back up to the last point the `commit` method was called.

Related concepts

[“Specifying segment search arguments using the `SSAList` interface” on page 646](#)

The `SSAList` interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Making batch data updates in IMS Universal DL/I driver applications

Use the `batchUpdate` method in the PCB interface to update multiple existing segments in the database with one call.

The following are the general steps to update multiple existing segments in the database with a single call:

1. Obtain an `SSAList` instance from the PCB instance representing the database.
2. Optionally, you can add qualification statements to the `SSAList` method. See "Specifying segment search arguments using the `SSAList` interface" for more information.

3. Get a Path instance by using the SSAList instance from the previous steps and calling the `getPathForBatchUpdate` method.
4. Using the Path instance from the step above, set the field values to update for the segments.
5. Call the `batchUpdate` method to update the segments.

The following code fragment illustrates how to use the `batchUpdate` method to modify a patient's name. The SSAList instance is set to update only records where the patient name is "ANDREA SMITH", the ward name is "SURG", and the hospital name is "ALEXANDRIA". The `getPathForBatchUpdate` method is called to obtain a Path containing the PATIENT segment and its child segments. Finally, the `batchUpdate` method is called to change the value of the patient name field to "ANDREA TAYLOR".

```
SSAList ssaList = pcb.getSSAList("HOSPITAL","PATIENT");
ssaList.addInitialQualification("HOSPITAL","HOSPNAME",SSAList.EQUALS,"ALEXANDRIA");
ssaList.addInitialQualification("WARD","WARDNAME",SSAList.EQUALS,"SURG");
ssaList.addInitialQualification("PATIENT","PATNAME",SSAList.EQUALS,"ANDREA SMITH");
Path path = ssaList.getPathForBatchUpdate("PATIENT");
path.setString("PATNAME", "ANDREA TAYLOR");
pcb.batchUpdate(path, ssaList);
```

Important: To persist changes made to the database, your application must call the `commit` method prior to deallocating the PSB, otherwise the changes are rolled back up to the last point the `commit` method was called.

Related concepts

“Specifying segment search arguments using the SSAList interface” on page 646

The SSAList interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Deleting data in a IMS Universal DL/I driver application

Use the `delete` method in the PCB interface to delete existing segments in the database.

In the IMS Universal DL/I driver, the delete methods provide functionality similar to the DL/I DLET call. The delete call must be preceded by a HOLD operation. Deleting a segment causes all its child segments to be deleted. The delete method will return an IMS status code indicating the results of the DL/I operation.

The following are the general steps to delete existing segments in the database:

1. Obtain an unqualified SSAList instance from the PCB instance representing the database.
2. Optionally, you can add qualification statements to the SSAList instance.
See "Specifying segment search arguments using the SSAList interface" for more information.
3. Get a Path instance by using the SSAList instance from steps 1 and 2 and calling the `getPathForRetrieveReplace` method.
4. Perform a Hold operation before issuing the replace call. The Hold operation can be a `getUnique`, `getNext`, or `getNextWithinParent` method call.
5. You can delete all the segments on the Path retrieved by step 3 or delete a subset of the segments.
 - To delete all the segments on the Path, call the `PCB.delete` method with no arguments.
 - If the Path retrieved by step 3 returned multiple segments from the database and you do not want to delete all the segments on the Path, use the `PCB.delete` method that takes an SSAList argument and pass in an unqualified SSAList for the segment where you want the deletion to begin. An exception is thrown if a qualified SSAList is provided as an argument.

IMS Universal DL/I driver delete examples

The following code fragment illustrates how delete all segments in a Path. Calling the `delete` method with no arguments removes all PATIENT segments and its dependent segments (ILLNESS, TREATMNT,

DOCTOR, BILLING) where the patient name is "ANDREA SMITH", the ward name is "SURG", the hospital name is "ALEXANDRIA", and the patient number is "PatientNo7".

```
SSAList ssaList = pcb.getSSAList("HOSPITAL", "ILLNESS");
ssaList.addInitialQualification("HOSPITAL", "HOSPNAME", SSAList.EQUALS, "ALEXANDRIA");
ssaList.addInitialQualification("WARD", "WARDNAME", SSAList.EQUALS, "SURG");
ssaList.addInitialQualification("PATIENT", "PATNAME", SSAList.EQUALS, "ANDREA SMITH");
ssaList.addCommandCode("PATIENT", SSAList.CC_D);
Path path = ssaList.getPathForRetrieveReplace();
if (pcb.getUnique(path, ssaList, true)) {
    if (path.getString("PATIENT", "PATNUM").equals("PatientNo7")) {
        pcb.delete();
    }
}
while (pcb.getNext(path, ssaList, true)) {
    if (path.getString("PATIENT", "PATNUM").equals("PatientNo7")) {
        pcb.delete();
    }
}
```

The following code fragment illustrates how to use delete with an unqualified SSAList. Calling the delete method with an unqualified SSAList removes all ILLNESS segments and its dependent segments (TREATMNT, DOCTOR) where the patient name is "ANDREA SMITH", the ward name is "SURGICAL", the hospital name is "ALEXANDRIA", and the patient number is "PatientNo7".

```
SSAList ssaList = pcb.getSSAList("HOSPITAL", "ILLNESS");
ssaList.addInitialQualification("HOSPITAL", "HOSPNAME", SSAList.EQUALS, "ALEXANDRIA");
ssaList.addInitialQualification("WARD", "WARDNAME", SSAList.EQUALS, "SURGICAL");
ssaList.addInitialQualification("PATIENT", "PATNAME", SSAList.EQUALS, "ANDREA SMITH");
ssaList.markAllFieldsForRetrieval("PATIENT", true);
Path path = ssaList.getPathForRetrieveReplace();
SSAList illnessSSAList = pcb.getSSAList("ILLNESS");
if (pcb.getUnique(path, ssaList, true)) {
    if (path.getString("PATIENT", "PATNUM").equals("PatientNo7")) {
        pcb.delete(illnessSSAList);
    }
}
while (pcb.getNext(path, ssaList, true)) {
    if (path.getString("PATIENT", "PATNUM").equals("PatientNo7")) {
        pcb.delete(illnessSSAList);
    }
}
```

Important: To persist changes made to the database, your application must call the commit method prior to deallocating the PSB, otherwise the changes are rolled back up to the last point the commit method was called.

Related concepts

[“Specifying segment search arguments using the SSAList interface” on page 646](#)

The SSAList interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Making batch data deletions in a IMS Universal DL/I driver application

Use the batchDelete method in the PCB interface to delete multiple existing segments in the database with one call.

The following are the general steps to delete multiple existing segments in the database with a single call:

1. Obtain an unqualified SSAList instance from the PCB instance representing the database.
2. Optionally, you can add qualification statements to the SSAList.
See "Specifying segment search arguments using the SSAList" for more information.
3. Call the batchDelete method to delete the segments specified by the SSAList in the previous steps.

The following code fragment illustrates how to use the `batchDelete` method to remove a patient's records. The `SSAList` instance is set to restrict the deletion operation to remove only records where the patient name is "ANDREA SMITH", the ward name is "SURG", and the hospital name is "ALEXANDRIA".

```
SSAList ssaList = pcb.getSSAList("HOSPITAL", "PATIENT");
ssaList.addInitialQualification("HOSPITAL", "HOSPNAME", SSAList.EQUALS, "ALEXANDRIA");
ssaList.addInitialQualification("WARD", "WARDNAME", SSAList.EQUALS, "SURG");
ssaList.addInitialQualification("PATIENT", "PATNAME", SSAList.EQUALS, "ANDREA SMITH");
pcb.batchDelete(ssaList);
```

Important: To persist changes made to the database, your application must call the `PSB.commit` method prior to deallocating the PSB, otherwise the changes are rolled back up to the last point commit was called.

Related concepts

[“Specifying segment search arguments using the SSAList interface” on page 646](#)

The `SSAList` interface represents a set of a list of segment search arguments used to specify the segments to target in a particular database call.

Making an Inquiry (INQY) DL/I call without a PSB object

A new API is introduced to allow issuing INQY DL/I calls directly without the need for a PSB object.

All Type-2 users of the IMS Universal JDBC driver can use the application interface to call a method that issues the desired Inquiry (INQY) DL/I call and return an appropriate object with the returned data from the call, without requiring the PSB metadata. This functionality is delivered through APAR PH19041/UI74514.

The following code examples illustrate how the application interface enables issuing INQY DL/I calls and return the object with the returned data from the call:

- **INQY DBQUERY** : A user can issue an INQY DBQUERY call using the application interface. A call can be issued by calling the **`inqyDBQuery()`** method from the application interface. A **`DBQueryInfo`** object is returned containing the call's returned information.

```
Application app = ApplicationFactory.createApplication();
DBQueryInfo info = app.inqyDBQuery();
info.getReturnCode();
info.getReasonCode();
info.getStatusCode();
```

Returned data: `DBQueryInfo` object

```
public int getReasonCode()
public String getReasonCodeHex()
public int getReturnCode()
public String getReturnCodeHex()
public int getStatusCode()
public String getStatusCodeChars()
```

- **INQY ENVIRON**: A user can issue an INQY ENVIRON call using the application interface. A call can be issued by calling the **`inqyEnviron()`** method from the application interface. An **`EnvironInfo`** object is returned containing the call's returned information.

```
Application app = ApplicationFactory.createApplication();
EnvironInfo info = app.inqyEnviron();
info.getControlRegionType();
```

Returned data: `EnvironInfo` object

```
public int getReasonCode()
public String getReasonCodeHex()
public int getReturnCode()
public String getReturnCodeHex()
public int getStatusCode()
```

```

public String getStatusCodeChars()
public String getRRSIndicator()
public String getCatalogIndicator()

```

- **INQY MSGINFO:** A user can issue an INQY MSGINFO call using the application interface. A call can be issued by calling the **inqyMessageInfo()** method from the application interface. A **MessageInfo** object is returned containing the call's returned information.

```

Application app = ApplicationFactory.createApplication();
MessageInfo info = app.inqyMessageInfo();
info.getIMSID();

```

Returned data: MessageInfo object:

```

public int getReasonCode()
public String getReasonCodeHex()
public int getReturnCode()
public String getReturnCodeHex()
public int getStatusCode()
public String getStatusCodeChars()

```

- **INQY FIND:** A user can issue an INQY FIND call using the application interface. A call can be issued by calling the **inqyFind()** method from the application interface. A PCB name must be entered as a string to issue the call with. A **FindInfo** object is returned containing the call's returned information.

```

Application app = ApplicationFactory.createApplication();
FindInfo info = app.inqyFind(this.expectedPCBName);
assertEquals(0, info.getReturnCode());
assertEquals(0, info.getReasonCode());
assertEquals(16448, info.getPCBStatusCode());

```

Returned data: FindInfo object

```

public short getPCBStatusCode()
public String getPCBStatusCodeChar()
public int getReasonCode()
public String getReasonCodeHex()
public int getReturnCode()
public String getReturnCodeHex()
public int getStatusCode()
public String getStatusCodeChars()

```

- **INQY PROGRAM:** A user can issue an INQY PROGRAM call using the application interface. A call can be issued by calling the **inqyProgram()** method from the application interface. A **ProgramInfo** object is returned containing the call's returned information.

```

Application app = ApplicationFactory.createApplication();
ProgramInfo info = app.inqyProgram();
info.getProgramName();

```

Returned data : ProgramInfo object

```

public String getProgramName()
public int getReasonCode()
public String getReasonCodeHex()
public int getReturnCode()
public String getReturnCodeHex()
public int getStatusCode()
public String getStatusCodeChars()

```

Note: To learn more about the Inquiry (INQY) call and the subfunctions, see [INQY call \(Application Programming APIs\)](#)

Related tasks

[“Writing DL/I calls to access an IMS database with the IMS Universal JDBC driver” on page 628](#)

In addition to support for SQL queries, the IMS Universal JDBC driver also supports casts to DL/I objects.

Using Byte Buffers with the IMS Universal DL/I driver

The IMS Universal DL/I driver includes several methods to handle complex data types that are not fully represented by the IMS Catalog.

The IMS Universal DL/I driver includes several methods that implement ByteBuffer input and output (I/O) objects instead of the traditional Path objects. Using ByteBuffers as I/O areas allows an application to retain full control in managing and interpreting data inserted or retrieved from IMS. These methods are useful when dealing with complex data types that are not fully represented by the IMS Catalog.

```
getUnique();
getNext();
getNextWithinParent();
insert();
create();
replace();
```

Note:

The ByteBuffer support is extended to support the IMS Message Queue interface.

Programming examples using ByteBuffers and the IMS Universal DL/I Driver

The following example shows an example of inserting multiple segments along the same hierarchical path. The hierarchical path consists of three segments that are children of the root segment Hospital.

```
// DEFINE SSA1 for INSERTS
SSAList ssaList = pcb.getSSAList("HOSPITAL", "ILLNESS");
ssaList.addInitialQualification("HOSPITAL", "HOSPCODE", SSAList.EQUALS, "R1210010000A");
ssaList.addCommandCode("WARD", SSAList.CC_D);
ssaList.addCommandCode("PATIENT", SSAList.CC_D);
ssaList.addCommandCode("ILLNESS", SSAList.CC_D);

// Allocate a ByteBuffer
ByteBuffer bufferInsert = ByteBuffer.allocate(1000);

// WARD
bufferInsert.position(0);
bufferInsert.putShort((short) 900);
bufferInsert.put("0010".getBytes("Cp1047"));
bufferInsert.put("NEW WARD".getBytes("Cp1047"));

// PATIENT
bufferInsert.position(900);
bufferInsert.putShort((short) 37);
bufferInsert.put("9999".getBytes("Cp1047"));
bufferInsert.put("JACKSON".getBytes("Cp1047"));

// ILLNESS
bufferInsert.position(937);
bufferInsert.putShort((short) 900);
bufferInsert.put("FRECKLES".getBytes("Cp1047"));

// Insert ByteBuffer
pcb.insert(bufferInsert, ssaList);
```

The following example issues a GU call with a hold on the segment to replace that retrieved segment.

```
// DEFINE SSA2 FOR RETRIEVE/REPLACE
SSAList ssaList2 = pcb.getSSAList("HOSPITAL", "PATIENT");
ssaList2.addInitialQualification("HOSPITAL", "HOSPCODE", SSAList.EQUALS, "R1210010000A");

// Retrieve the segment with the ByteBuffer as an output object
ByteBuffer bufferRetrieve = ByteBuffer.allocate(900);
pcb.getUnique(bufferRetrieve, ssaList2, true);

// Replace
bufferReplace = ByteBuffer.allocate(42);
bufferReplace.putShort((short) 27);
bufferReplace.put("9999".getBytes("Cp1047"));
```

```
bufferReplace.put("FRANKLIN".getBytes("Cp1047"));
pcb.replace(bufferReplace);
```

Considerations for ByteBuffers and Type-2 Java applications

If your JDBC application is running in Type-2 mode, you have the option of obtaining a direct ByteBuffer to optimize JVM storage use. You can obtain direct storage in the following ways:

- Directly using the **java.nio.ByteBuffer** class

```
ByteBuffer directBuffer = ByteBuffer.allocateDirect(10);
```

- Using the Universal Drivers Application object to obtain direct 31-bit storage

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);
```

Note: If an application is running in the 64-bit JVM, `ByteBuffer.allocateDirect()` will not work with your application and you must directly request 31-bit storage using the `Application.get31BitDirectByteBuffer()` method. `ByteBuffer.allocateDirect()` works in a 31-bit JVM.

It is important to release and free the buffers correctly in your application. You can free buffers that were obtained using the **Application** class in the following ways:

- Specifically freeing buffers

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);

// Doing work
app.free31BitDirectByteBuffer(directBuffer);
```

- IMS silently tracks all buffers allocated using the **Application** class and automatically releases all the checked out buffers during the `end()` procedure.

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);
// Doing work
app.end();
```

Related concepts

[“Using ByteBuffers with the IMS Universal JDBC Driver” on page 608](#)

The IMS Universal JDBC driver allows the use of INSERT, SELECT, and UPDATE options for full segments using **java.nio.ByteBuffer** as an I/O area .

Related reference

[“Using the expandArrayResultSet property” on page 611](#)

The IMS Universal JDBC driver allows for easier access to Array datatype fields using the **expandArrayResultSet** property.

[“Hospital database example” on page 568](#)

The code examples for the IMS Universal drivers application programming topics use the Hospital database.

Inspecting the PCB status code and related information using the **com.ibm.ims.dli.AIB** interface

To inspect the PCB status code, return code, reason code, error code extension, and related information after a data access call by the IMS Universal drivers, use the **com.ibm.ims.dli.AIB** interface provided by the IMS Universal DL/I driver.

Typically, the IMS Universal drivers throws an exception if a call is not successful. The non-blank status codes that do not generate an exception are GD, GE, GB, GA, GK, QC, QD, and CF. In error cases, the generated exception contains the most pertinent information such as PCB status code, return code,

reason code, and error code extension. To inspect the PCB status code and related information for all the non-error cases, use the `com.ibm.ims.dli.AIB` interface.

The AIB instance contains all the data attributes of an IMS application interface block. The AIB instance also contains a reference to a `com.ibm.ims.dli.DBPCB` instance, which contains all the data attributes of a PCB instance.

When your IMS Universal drivers application makes a data access call to IMS, the application is internally making a DL/I call using a `com.ibm.ims.dli.PCB` instance. After each DL/I call that your application issues, IMS places a two-character status code in the DBPCB instance stored in the AIB instance for that PCB instance.

The `com.ibm.ims.dli.IMSStatusCodes` class contains constants for the IMS status codes. Use this helper class for comparison checking of the status code from the DL/I call.

The following code example shows how to access the AIB instance from a IMS Universal DL/I driver application:

```
try {
    psb = PSBFactory.createPSB(connSpec);
} catch (DLIException e) {
    AIB aib = e.getAib();
    if (aib != null) {
        String sc = aib.getDBPCB().getStatusCodeChars();
        String retcode = aib.getReturnCodeHex();
        String reascode = aib.getReasonCodeHex();
        System.out.println("Status code: " + sc + " Return Code: "
            + retcode + " Reason Code: " + reascode);
    }
}
```

The following code example shows how to access the AIB instance from a JDBC application. Note that you need to import the `com.ibm.ims.dli` package to use the AIB and DLIException objects in your Java code.

```
try {
    resultSet.updateString(1, "Harry Houdini");
} catch (SQLException e) {
    Throwable t = e.getCause();
    if (t != null && t instanceof DLIException) {
        com.ibm.ims.dli.DLIException de = (com.ibm.ims.dli.DLIException) t;
        com.ibm.ims.dli.AIB aib = de.getAib();
        if (aib != null) {
            String sc = aib.getDBPCB().getStatusCodeChars();
            String retcode = aib.getReturnCodeHex();
            String reascode = aib.getReasonCodeHex();
            System.out.println("Status code: " + sc + " Return Code: "
                + retcode + " Reason Code: " + reascode);
        }
    }
}
```

Related reference

[DL/I status code explanations \(Messages and Codes\)](#)

Committing or rolling back DL/I transactions

The IMS Universal DL/I driver provides support for local transactions with the `commit` and `rollback` methods.

A local transaction consists of a unit of work with several units of recovery. A IMS Universal DL/I driver application can commit or roll back changes to the database within a unit of recovery. In the IMS Universal DL/I driver, the local transaction is scoped to the PSB instance. No explicit call is needed to begin a local transaction. A unit of work starts when the application allocates a PSB object and obtains a connection to the database by calling the `PSB.allocate` method.

After the unit of work starts, the application makes DL/I calls to access the database and create, replace, insert, or delete data. The application commits the current unit of recovery by using the `PSB.commit` method. The commit operation instructs the database to commit all changes to the database that are

made from the point when the unit of work started, or from the point after the last commit or rollback method call, whichever was most recent.

Important: To persist changes made to the database, your application must call the commit method prior to deallocating the PSB, otherwise the changes are rolled back up to the last point the commit method was called.

The application can also end the unit of recovery by calling a roll back operation using the PSB.rollback method. Calling a roll back operation causes the database to undo all changes to the database made from the start of the unit of work, or from the point after the most recent commit or rollback call.

If the PSB.commit method or the PSB.rollback method is called and the PSB instance is not deallocated, a new unit of recovery is started. The overall unit of work ends when the PSB instance is deallocated. If the PSB.commit method or the PSB.rollback method are called while they are not currently in a unit of work (either before the PSB instance is allocated or after it is deallocated), an exception is thrown.

Local transaction with a single PSB

The following example code shows a local transaction for a single PSB.

```
IMSConnectionSpec connSpec = IMSConnectionSpecFactory.createIMSConnectionSpec();
connSpec.setDatastoreName("IMS1");
connSpec.setDatastoreServer("ecdev123.svl.ibm.com");
connSpec.setPortNumber(5555);
connSpec.setMetadataURL("class://BMP266.BMP266DatabaseView");
connSpec.setUser("usr");
connSpec.setPassword("password");
connSpec.setDriverType(IMSConnectionSpec.DRIVER_TYPE_4);

PSB psb = PSBFactory.createPSB(connSpec);
psb.allocate(); // new unit of work begins
PCB pcb = psb.getPCB("Pcb01");

SSAList ssa = pcb.getSSAList("HOSPITAL");
Path path = ssa.getPathForInsert("HOSPITAL");
path.setString("HOSPCODE", "R12100200000A");
path.setString("HOSPNAME", "SANTA TERESA");
pcb.insert(path);
psb.commit(); // or use psb.rollback() to undo the insert.
// The unit of recovery ends.
```

In this example, the application makes a connection to the database and allocates a PSB. The application obtains a PCB and specifies the path to insert a new HOSPITAL record. The application then performs a DL/I operation to insert the new record into the database. At this point, the application commits the insert operation and the new record is written to the database. Alternatively, the application can roll back the insert operation to return the database to the previous state before the insert call was made. This ends the current unit of recovery.

Local transaction with multiple PSBs

When two or more PSB objects are allocated by an application, separate local transactions for each PSB may run concurrently. The following example code shows multiple local transactions with two PSBs.

```
IMSConnectionSpec connSpec = IMSConnectionSpecFactory.createIMSConnectionSpec();
connSpec.setDatastoreName("IMS1");
connSpec.setDatastoreServer("ecdev123.svl.ibm.com");
connSpec.setPortNumber(5555);
connSpec.setMetadataURL("class://BMP266.BMP266DatabaseView");
connSpec.setUser("usr");
connSpec.setPassword("password");
connSpec.setDriverType(IMSConnectionSpec.DRIVER_TYPE_4);

// create a connection to MyDB
PSB psb = PSBFactory.createPSB(connSpec);
psb.allocate(); // new unit of work begins for psb

// create another connection to MyDB.
```

```
// Note: This does not need be be a connection to the same database.
PSB psb2 = PSBFactory.createPSB(connSpec);
psb2.allocate();

pcb = psb.getPCB("PCb01");
SSAList ssa = pcb.getSSAList("HOSPITAL");
Path path = ssa.getPathForInsert("HOSPITAL");
path.setString("HOSPCODE", "R1210020000A");
path.setString("HOSPNAME", "SANTA TERESA");
pcb.insert(path);
psb.commit(); // or use psb.rollback() to undo the insert.
// The unit of recovery for psb ends

pcb2 = psb2.getPCB("PCb01");
SSAList ssa2 = pcb2.getSSAList("HOSPITAL");
Path path2 = ssa2.getPathForInsert("HOSPITAL");
path2.setString("HOSPCODE", "R1210010000A");
path2.setString("HOSPNAME", "ALEXANDRIA");
pcb2.insert(path2);
psb2.rollback(); //or use psb2.commit() to commit the insert.
// The unit of recovery for psb2 ends

psb2.deallocate(); // unit of work ends for psb2
psb.deallocate(); // unit of work ends for psb
```

In this example, the application makes two connections to the same database. A PSB is allocated for the first connection and the application performs a DL/I operation to insert a new HOSPITAL record with hospital name "SANTA TERESA" into the database. Another PSB is allocated to the second connection and an insert operation is made for a new HOSPITAL record with hospital name "ALEXANDRIA". The application then commits the changes for the first PSB and writes the new record with hospital name "SANTA TERESA" to the database. The application issues a roll back statement for the second PSB, undoing the previous insert operation for the record with hospital name "ALEXANDRIA". Only one new record is inserted to the database: the HOSPITAL record with hospital name "SANTA TERESA".

Accessing dynamic arrays with the Universal Drivers using the DBArrayElementSet class

The IMS Universal DL/I driver enables accesses and updates to dynamic arrays with the DBArrayElementSet class.

A dynamic array is an array in which the number of repeating array elements can vary from one instance of a segment type to another. It is defined by coding FIELD statements in the input control statements of the DBD Generation utility. To access dynamic arrays through IMS Universal DL/I driver, use the DBArrayElementSet class, which supports creating, receiving, and updating data structures that represent DBArray objects.

The following example demonstrates how a DBArrayElementSet object accesses and updates a dynamic array that has 15 repeating occurrences.

```
ssaList = pcb.getSSAList("HOSPITAL");
ssaList.addInitialQualification(1, "HOSPCODE", SSAList.EQUALS, "R1210010000A");
path = ssaList.getPathForRetrieveReplace();
pcb.getUnique(path, ssaList, true);
DBArray dbArray = path.getArray("DYNAMICARRAY");
DBArrayElementSet dbArrayElementSet = dbArray.getElements();
dbArrayElementSet.absolute(0);

dbArrayElementSet.prepareElementToAdd();
dbArrayElementSet.setString("DYNAMICFIELD1", "RR");
dbArrayElementSet.setString("DYNAMICFIELD2", "RRRR");
dbArrayElementSet.setString("DYNAMICFIELD3", "DDDD");
dbArrayElementSet.addElement();
dbArrayElementSet.setString("DYNAMICFIELD1", "XX");
dbArrayElementSet.setString("DYNAMICFIELD2", "VVVV");
dbArrayElementSet.setString("DYNAMICFIELD3", "ZZZZ");
dbArrayElementSet.addElement();
dbArrayElementSet.absolute(17);
dbArrayElementSet.setString("DYNAMICFIELD1", "SS");
dbArrayElementSet.setString("DYNAMICFIELD2", "SSSS");
dbArrayElementSet.setString("DYNAMICFIELD3", "DDDD");
dbArrayElementSet.addElement();
dbArrayElementSet.setString("DYNAMICFIELD1", "XX");
```

```

dbArrayElementSet.setString("DYNAMICFIELD2", "VVVV");
dbArrayElementSet.setString("DYNAMICFIELD3", "ZZZZ");
dbArrayElementSet.addElement();

path.setInt("OCCURSFIELD", 19);
path.setArray("DYNAMICARRAY", dbArray);
pcb.replace(path);

```

In this example, the array field is named DYNAMICARRAY and contains three elements named DYNAMICFIELD1, DYNAMICFIELD2, DYNAMICFIELD3. The array field has a DEPENDSON field for setting the number of occurrences of an element with the name OCCURSFIELD. Four more occurrences are added to the array structure, making the total occurrence number 19. A replace (REPL) call is made to update the path.

You might find the following notes helpful when using the previous example:

- The `prepareElementToAdd()` function moves the `DBArrayElementSet` object to an insert row buffer that can be used to set the array's fields. Your current cursor's position will be stored. You will remain on the insert row buffer until the `moveToCurrentElement()` function is called that restores your cursor's position.

Note: Only dynamic arrays can use the `prepareElementToAdd()` function.

- You can move the `DBArrayElementSet` object's cursor by calling the following functions:

- `previous()`
- `next()`
- `last()`
- `first()`
- `beforeFirst()`
- `afterLast()`
- `absolute(int position)`
- `relative(int position)`

Moving the cursor's position cancels any pending data on the insert row buffer that is not added by the `addElement()` function.

- You can use setters such as `setString()`, `setInt()`, and so on to complete the following tasks:
 - Updating the array field value at the current cursor position
 - Adding the array field value into the insert row buffer if the array field value is initialized
- You can use getters such as `getString()`, `getInt()`, and so on to complete the following tasks:
 - Retrieving the array field value at the current cursor position
 - Retrieving the array field value from the insert row buffer if the array field value is initialized
- The DEPENDSON field must always be set into the path object when you add or update a `DBArrayElementSet` object and must reflect the total number of elements.

Related tasks

[Defining a dynamic array to IMS \(Database Administration\)](#)

Configuring the IMS Universal drivers for SSL support

With type-4 connectivity, the IMS Universal drivers provide support for the Secure Sockets Layer (SSL) through the Java Secure Socket Extension (JSSE).

This information applies to type-4 connectivity only. You can use SSL support in your Java applications in either a container-managed environment with the IMS Universal Database resource adapter, or in a stand-alone environment with the IMS Universal JDBC driver and the IMS Universal DL/I driver.

Configuring the IMS Universal Database resource adapter for SSL support in a container-managed environment

To enable SSL in a container-managed environment for the IMS Universal Database resource adapter, you need to configure the SSL certificate and key management settings from your WebSphere Application Server administrative console.

Prerequisites:

- You must first set up the IBM z/OS Communications Server Application Transparent Transport Layer Security (AT-TLS) to enable SSL support on the z/OS system for IMS Connect.
- You also need to retrieve the client certificate (.crt) to your local file system where WebSphere Application Server is installed. To retrieve the certificate, from TSO, browse the OMVSID.CERTAUTH.CERT member. Copy its contents into a text file on your local file system, and remove any trailing spaces. Name the file `hostname.crt`.

To configure the IMS Universal Database resource adapter for SSL support:

1. Open the WebSphere Application Server administrative console.
2. From the left pane, expand **Security -> SSL certificate and key management**.
3. Click **Key stores and certificates**.
4. Click **NodeDefaultTrustStore**.
5. Click **Signer certificates**.
6. Click **Add**.
7. In the **Alias** field, type a name that helps you remember that this certificate is associated with (extracted from) the server key ring file that was created when you set up AT-TLS to enable SSL on IMS Connect.
8. In the **File name** field, type the fully qualified path to the .crt file located on your local file system.
9. Click **OK** and then click **Save**.

The trusted certificate is picked up automatically and used during the SSL handshaking process at run time.

Related tasks

[Setting up AT-TLS SSL for IMS Connect \(System Definition\)](#)

Configuring IMS Universal drivers for SSL support in a stand-alone environment

To enable SSL in a stand-alone environment for the IMS Universal drivers, you need to generate and configure an SSL keystore.

Prerequisites:

- You must first set up the IBM z/OS Communications Server Application Transparent Transport Layer Security (AT-TLS) to enable SSL support on the z/OS system for IMS Connect.
- You also need to retrieve the client certificate (.crt) to your local file system. To retrieve the certificate, from TSO, browse the OMVSID.CERTAUTH.CERT member. Copy its contents into a text file on your local file system, and remove any trailing spaces. Name the file `hostname.crt`.

To configure the IMS Universal DL/I driver or the IMS Universal JDBC driver for SSL support:

1. Generate a new SSL keystore by using the Java Keytool provided by the Java SDK.

This keystore file will be used as a truststore by the JRE during SSL handshaking when it creates an SSL connection to IMS. Save the keystore (.ks) file on your local file system and record its location. Set the password for the keystore and record it.

Keystore files can contain public/private key pairs that are generated on the local system as well as public keys (in the form of certificates) that are received from remote communicating peers. When the

keystore is accessed to retrieve a certificate of a communicating peer for use during SSL handshaking, the keystore file is referred to as a *truststore*.

2. Verify that the certificate has not been tampered with before importing the certificate (.crt) file into the keystore as a trusted self-signed certificate.

You can do this with the Keytool by viewing the fingerprint of the local certificate and comparing it to the original that was extracted from the key ring file on the host.

3. Set the fully qualified path to the keystore file as the value for the system property `javax.net.ssl.trustStore` and set the keystore password as the value for the system property `javax.net.ssl.trustStorePassword`.

Optionally, to troubleshoot any SSL-related problems, you can turn on the SSL client-side trace by setting the system property `javax.net.debug=all`. To specify the system properties from the command line, enter:

```
java -Djavax.net.debug=all -Djavax.net.ssl.trustStore=myTruststore  
-Djavax.net.ssl.trustStorePassword=myTruststorePassword MyApp
```

Related tasks

[Setting up AT-TLS SSL for IMS Connect \(System Definition\)](#)

Tracing IMS Universal drivers applications

To obtain data for diagnosing problems with the IMS Universal drivers, you can collect trace data.



Attention: Collecting tracing data must be enabled only for diagnosing problems with the IMS Universal drivers. Enabling traces in a production environment on a regular basis is not recommended.

Use one of the following procedures to enable tracing.

Turning on automatic tracing in JRE logging.properties file

The recommended method is to enable the trace by setting the trace level for the IMS Universal drivers loggers in the `logging.properties` file of your Java Runtime Environment (JRE). Using this method, the application does not need to be recompiled. The file is located on the install path of your JRE, under `\jre\lib\logging.properties`. The recommended trace level is `FINEST`.

To set the trace for all IMS Universal drivers loggers, add the following line to the `logging.properties` file:

```
com.ibm.ims.db.opendb = FINEST
```

To send the trace output to a file, make the following changes to your `logging.properties` file:

1. Make sure that the file includes the following line and no other `handlers=` lines, unless they are commented out. A pound sign (#) at the beginning of a line indicates that the line is a comment.

```
handlers = java.util.logging.FileHandler, java.util.logging.ConsoleHandler
```

2. Add the following lines:

```
java.util.logging.FileHandler.level = FINEST  
java.util.logging.FileHandler.pattern = c:/UniversalDriverTrace.txt  
java.util.logging.FileHandler.formatter = java.util.logging.SimpleFormatter
```

Configuring Java EE tracing

Tracing can be turned on from your Java EE application server. In WebSphere Application Server, this is configured through the administrative console. The IMS Universal Database resource adapter must be deployed on WebSphere Application Server before tracing can be configured.

In WebSphere Application Server Liberty, tracing is configured through the `server.xml` configuration file.

- To get the most detailed trace from the IMS Universal Database resource adapter, follow these steps:
 1. Start the WebSphere Application Server administration console.
 2. Select **Troubleshooting**.
 3. Select **Logs and Trace**.
 4. Select your application server from the table.
 5. Under **General Properties**, select **Diagnostic Trace**.
 6. Under **Additional Properties**, select **Change Log Detail Levels**.
 7. Select the **Runtime** tab.
 - Make sure that the **Save runtime changes to configuration as well** check box is turned **ON**.
 - Under the **Change Log Details Levels** section, select the component **com.ibm.ims.db.opendb**. This brings up the **Message and Trace levels** menu.
 - Select the message level **FINEST**.
 - Click **Apply**.
 8. To save these changes for the next time the application server is started, click the **Save** link at the top of the page. WebSphere Application Server does not need to be restarted.
- To get the most detailed trace from the WebSphere Application Server Liberty, follow these steps:
 1. In the `server.xml` configuration file, add a `<logging>` element that contains a `traceSpecification` attribute and the String argument "com.ibm.ims.db.opendb.*" :

```
<logging traceSpecification="*=info:com.ibm.ims.db.opendb.*=finest"/>
```
 2. The default directory for the `trace.log` file can be found in the `logsk/` directory of the server instance. No restart is required for WebSphere Application Server Liberty.

Programmatically enabling tracing

You can also programmatically turn on tracing in your IMS Universal drivers application. This requires the application to be recompiled.

1. Import the `java.util.logging` package in your application and create a logger by calling the `Logger.getLogger` method with the String argument "com.ibm.ims.db.opendb".
2. In your application, you can set the level of tracing for the logger by using the `Logger.setLevel` method. The recommended trace level is `Level.FINEST`.

The following sample code shows how programmatic trace is enabled for any IMS Universal drivers application.

```
private static final Logger universalLogger
    = Logger.getLogger("com.ibm.ims.db.opendb");
universalLogger.setLevel(Level.FINEST);
FileHandler fh
    = new FileHandler("C:/UniversalTrace.txt");
fh.setFormatter(new SimpleFormatter());
fh.setLevel(Level.FINEST);
universalLogger.addHandler(fh);
```

Related tasks

[IMS Universal drivers: WebSphere Application Server Liberty type-4 connections \(Communications and Connections\)](#)

[IMS Universal drivers: WebSphere Application Server Liberty type-2 connections \(Communications and Connections\)](#)

Chapter 41. Programming Java dependent regions

Use these topics to design, write, and maintain application programs for running in the Java dependent regions.

Related concepts

[“IMS solutions for Java development overview” on page 553](#)

You can write Java applications to access IMS databases and process IMS transactions by using the drivers and resource adapters of the IMS solutions for Java development.

Overview of the IMS Java dependent regions

The IMS Java dependent regions are two types of IMS dependent regions that provide a Java Virtual Machine (JVM) environment for Java applications: Java message processing (JMP) regions and Java batch processing (JBP) regions.

Important: You can host Java applications on the mainframe to access IMS from the following z/OS environments:

- JMP and JBP regions
- WebSphere Application Server for z/OS
- Db2 for z/OS stored procedures
- CICS

Use the JMP or JBP regions to host your Java application if your application is required to run in an IMS dependent region.

JMP and JBP regions can run applications written in Java, object-oriented COBOL, object-oriented PL/I, or a combination of these languages.

To access IMS message queues from your JMP and JBP applications, use the IMS Java dependent region resource adapter. To access IMS databases from your JMP and JBP applications, you can also use these IMS Universal drivers: the IMS Universal JDBC driver and the IMS Universal DL/I driver.

In addition to IMS databases, you can access Db2 for z/OS databases from your JMP and JBP applications by using the JDBC driver for Db2 for z/OS (JCC driver version 3.57.91).

JMP and JBP regions can load and run with a 31-bit or 64-bit Java virtual machine (JVM). Use the JVM=64 or JVM=31 parameter to switch between the 64-bit and 31-bit addressing mode. If no value is specified on JVM=, the 31-bit addressing mode is used.

Java message processing (JMP) regions

JMP regions are like message processing program (MPP) regions, but JMP regions allow the scheduling only of Java programs. In the PSB source associated with the Java program, the option LANG=JAVA must be specified. A JMP application is started when there is a message in the queue for the JMP application and IMS schedules the message to be processed. JMP applications, like MPP applications, are executed through transaction codes submitted by users at terminals and from other applications. Each transaction code represents a transaction that the JMP application processes.

A single application can also be started from multiple transaction codes. JMP applications, like MPP applications, are flexible in how they process transactions and where they send the output. JMP applications send any output messages back to the message queues and process the next message with the same transaction code. The program continues to run until there are no more messages with the same transaction code. JMP applications share the following characteristics:

- They are small.
- They can produce output that is needed immediately.
- They can access IMS or DB2 data in a DB/DC environment and DB2 data in a DCCTL environment.

Java batch processing (JBP) regions

JBP regions run flexible programs that perform batch-type processing online and can access the IMS message queues for output, like non-message-driven batch message processing (BMP) applications. JBP applications are started by submitting a job with JCL or from TSO. JBP applications are like BMP applications, except that they cannot read input messages from the IMS message queue. For example, there is no IN= parameter in the startup procedure. Like BMP applications, JBP applications can use symbolic checkpoint and restart calls to restart the application after an abend. JBP applications can access IMS or Db2 for z/OS data in a DB/DC or DBCTL environment and Db2 for z/OS data in a DCCTL environment.

Related concepts

[“Programming with the IMS Universal drivers” on page 559](#)

Use these topics to design, write, and maintain application programs for IMS 15.3 using the IMS Universal drivers.

Related tasks

[“IBM Enterprise COBOL for z/OS interoperability with JMP and JBP applications” on page 702](#)

With the IBM Enterprise COBOL for z/OS support for COBOL and Java language interoperability, you can write Java and Object-Oriented (OO) COBOL applications that execute in a Java dependent region and invoke existing COBOL programs.

[“Accessing Db2 for z/OS databases from JMP or JBP applications” on page 703](#)

A JMP or JBP application can access Db2 for z/OS using the latest JDBC driver for Db2 for z/OS.

Programming with the IMS Java dependent region resource adapter

IMS provides a set of Java APIs called the IMS Java dependent region resource adapter to develop Java applications to run on the IMS Java dependent regions.

The IMS Java dependent region resource adapter provides Java application programs running in JMP or JBP regions with similar DL/I functionality to that provided in message processing program (MPP) and non-message driven BMP regions, such as:

- Accessing IMS message queues to read and write messages
- Performing program switches
- Commit and rollback processing
- Accessing GSAM databases
- Database recovery (CHKP/XRST)

Use the IMS Java dependent region resource adapter together with the type-2 IMS Universal JDBC driver or type-2 IMS Universal DL/I driver to perform database operations, including GSAM database access.

The following figure shows a Java application that is running in a JMP or JBP region. Database access and message processing requests are passed to the IMS Java dependent region resource adapter and type-2 IMS Universal drivers, which converts the calls to DL/I calls.




Figure 109. JMP or JBP application that is using the IMS Java dependent region resource adapter

Preparing to write a Java application with the IMS Java dependent region resource adapter

The IMS Java dependent region resource adapter is available as an SMP/E-installable driver (`imsutm.jar` through the Java on Demand Feature FMID).

Java application programs that use the IMS Java dependent region resource adapter require the Java Development Kit (JDK) 7.0 or later. They also require a way to generate the IMS database metadata, such as using the IMS Enterprise Suite Explorer for Development. The default segment encoding of the database metadata class produced by the IMS Explorer for Development is `cp1047`. To change the segment encoding, use the `com.ibm.ims.base.DLIBaseSegment.setDefaultEncoding` method.

Related concepts

[“Gathering requirements for message processing options” on page 91](#)

One of the tasks of application design is providing information about your application's requirements to the people in charge of designing and administering your IMS system.

Developing JMP applications with the IMS Java dependent region resource adapter

Java message processing (JMP) applications access the IMS message queue to receive messages to process and to send output messages.

Defining the input and output message classes

Before your JMP application can access the message queue, you must define input and output message classes by subclassing the `com.ibm.ims.application.IMSFieldMessage` class.

Recommendation: Use the IMS Enterprise Suite Explorer for Development to generate the necessary metadata class files from COBOL copybooks or PL/I resources when this support is available.

The IMS Java dependent region resource adapter provides the capability to process `IMSFieldMessage` objects.

Subclass IMSFieldMessage: input message sample code

This example code subclasses the `com.ibm.ims.application.IMSFieldMessage` class to make the fields in the message available to the program and creates an array of `com.ibm.ims.base.DLTypeInfo` objects for the fields in the message. For the `DLTypeInfo` class, the code identifies first the field name, then the data type, the position, and finally the length of the individual fields within the array. This allows the application to use the access functions within the `IMSFieldMessage` class hierarchy to automatically convert the data from its format in the message to a Java type that the application can process. In addition to the message-specific fields that it defines, the `IMSFieldMessage` class provides access functions that allow it to determine the transaction code and the length of the message.

This class defines an input message that accepts a 2-byte type code of a car model to query a car dealership database for available car models.

```
package dealership.application;
import com.ibm.ims.db.*;
import com.ibm.ims.base.*;
import com.ibm.ims.application.*;

/* Subclasses IMSFieldMessage to define application's input messages */
public class InputMessage extends IMSFieldMessage {

    /* Creates array of DLTypeInfo objects for the fields in message */
    final static DLTypeInfo[] fieldInfo={
        new DLTypeInfo("ModelTypeCode", DLTypeInfo.CHAR, 1, 2)
    };

    public InputMessage() {
        super(fieldInfo, 2, false);
    }
}
```

Subclass IMSFieldMessage: output message sample code

The following code example shows how to subclass the `com.ibm.ims.application.IMSFieldMessage` class to define an output message that displays the available car models from a type code query.

This sample code creates an array of `com.ibm.ims.base.DLTypeInfo` objects and then passes that array, the byte array length, and the Boolean value false, which indicates a non-SPA message, to the `IMSFieldMessage` constructor. For each `DLTypeInfo` object, you must first identify the field data type, then the field name, the field offset in the byte array, and finally the length of the byte array.

```
package dealership.application;
import com.ibm.ims.db.*;
import com.ibm.ims.base.*;
import com.ibm.ims.application.*;

/*Subclasses IMSFieldMessage to define application's output messages */
public class ModelOutput extends IMSFieldMessage {

    s    /* Creates array of DLTypeInfo objects for the fields in message */
    final static DLTypeInfo[] fieldInfo={
        new DLTypeInfo("Type", DLTypeInfo.CHAR, 1, 2),
        new DLTypeInfo("Make", DLTypeInfo.CHAR, 3, 10),
        new DLTypeInfo("Model", DLTypeInfo.CHAR, 13, 10),
        new DLTypeInfo("Year", DLTypeInfo.DOUBLE, 23, 4),
        new DLTypeInfo("CityMiles", DLTypeInfo.CHAR, 27, 4),
        new DLTypeInfo("HighwayMiles", DLTypeInfo.CHAR, 31, 4),
        new DLTypeInfo("Horsepower", DLTypeInfo.CHAR, 35, 4)
    };

    public ModelOutput() {
        super(fieldInfo, 38,false);
    }
}
```


JMP programming models

JMP applications can retrieve input messages from the IMS message queue, access IMS and Db2 for z/OS databases, commit or roll back transactions, and send output messages.

Creating the main method for a JMP application

The main method (`public static void main(String[] args)`) is the program entry point for all JMP and JBP applications.

A JMP application starts when IMS receives a message with a transaction code for the JMP application and schedules the message. A JMP application typically ends when there are no more messages with that transaction code to process.

JMP application main method code sample

The following code sample shows how to implement a JMP application to access the hospital database and send messages:

```
package hospital.ims;

import java.sql.*;
import com.ibm.ims.dli.tm.*;
import com.ibm.ims.dli.DLIException;

public static void main(String args[]) {
    try {
        Application app = null;
        MessageQueue messageQueue = null;
        IOMessage inputMessage = null;
        IOMessage outputMessage = null;
        Transaction tran = null;

        app = ApplicationFactory.createApplication();
        inputMessage = app.getIOMessage("class://hospital.ims.InMessage");
        outputMessage = app.getIOMessage("class://hospital.ims.OutMessage");
        messageQueue = app.getMessageQueue();
        tran = app.getTransaction();

        IMSDataSource dataSource = new IMSDataSource();
        dataSource.setMetadataURL("class://hospital.ims.HospitalDBView");
        dataSource.setDriverType(IMSDataSource.DRIVER_TYPE_2);
        dataSource.setDatastoreName("IMS1");

        Connection conn = dataSource.getConnection();
        conn.setAutoCommit(false);
        Statement st = conn.createStatement();

        String in = new String("");

        // Returns true if a message is read from the queue
        while (messageQueue.getUnique(inputMessage)) {

            in = inputMessage.getString("Message").trim();
            if (!in.equals("")) {

                // Query the database for all hospital names
                ResultSet rs
                = st.executeQuery("SELECT HOSPNAME FROM PCB01.HOSPital");

                while (rs.next()) {

                    // Return hospital name in output message
                    outputMessage.setString("Message", rs.getString("HOSPNAME"));
                    messageQueue.insert(outputMessage,
                        MessageQueue.DEFAULT_DESTINATION);

                    // Commit this transaction
                    tran.commit();
                }
                conn.close();
            }
        }
    } catch (Exception e) {
        e.printStackTrace();
    }
}
```

```
}  
}
```

Accessing Db2 for z/OS data from a JMP application

When a JMP application accesses only IMS data, it must open a database connection only once to process multiple transactions. However, a JMP application that accesses Db2 for z/OS data must open and close a database connection for each message that is processed.

Processing an input message in a JMP application

A transaction begins when the application receives an input message and ends when the application commits the results from processing the message. To get an input message, the application calls the `MessageQueue.getUnique` method.

Processing an input message sample code

The following code example shows how an input message is processed in a JMP application.

```
import com.ibm.ims.dli.tm.*;  
  
public static void main(String args[]) {  
    conn = dataSource.getConnection(...); //Establish DB connection  
    while(messageQueue.getUnique(...)){ //Get input message, which  
                                        //starts transaction  
        results=statement.executeQuery(...); //Perform DB processing  
        ...  
        messageQueue.insert(...); //Send output messages  
        ...  
    }  
    conn.close(); //Close DB connection  
    return;  
}
```

Rolling back IMS changes in a JMP application

A JMP application can roll back IMS changes any number of times during a transaction. A rollback call backs out all output messages to the most recent commit.

Use the `com.ibm.ims.dli.tm.Transaction` class to issue commit and rollback operations from your JMP application.

The following code example shows how a JMP application rolls back IMS changes.

```
import com.ibm.ims.dli.tm.*;  
import java.sql.*;  
  
public static void main(String args[]) {  
    conn = dataSource.getConnection(...); //Establish DB connection  
    Application app = ApplicationFactory.createApplication();  
    Transaction tran = app.getTransaction();  
    MessageQueue mq = app.getMessageQueue();  
  
    while(mq.getUnique(...)){ //Get input message, which  
                            //starts transaction  
        results=statement.executeQuery(...); //Perform DB processing  
        ...  
        mq.insertMessage(...); //Send output messages  
        ...  
        tran.rollback(); //Roll back output messages  
  
        results=statement.executeQuery(...); //Perform more DB processing  
        //(optional)  
        ...  
        mq.insert(...); //Send more output messages  
        //(optional)  
    }  
}
```

```

    }
    conn.close(); //Close DB connection
}

```

Additional message handling considerations for JMP applications

The following considerations apply to JMP applications that access the IMS message queue when handling conversational transactions, multi-segment messages, messages with repeating structures, and multiple input messages.

Conversational transactions

The IMS Java dependent region resource adapter supports access to IMS conversational transactions.

Conversational transactions

A conversational transaction does not process the entire transaction at the same time. A conversational program divides processing into a connected series of terminal-to-program-to-terminal interactions. Use conversational processing when one transaction contains several parts. In contrast, a nonconversational program receives a message from a terminal, processes the request, and sends a message back to the terminal.

A conversational program receives a message from a terminal and replies to the terminal, but it saves the data from the transaction in a scratchpad area (SPA). When the user at the terminal enters more data, the program has the data it saved from the last message in the SPA, so it can continue processing the request without the user at the terminal having to enter the data again.

Conversational transaction sample

The following code example shows how to write a JMP application to process a conversational transaction.

```

package mytest.jdbo;

import com.ibm.ims.dli.DLIException;
import com.ibm.ims.dli.tm.*;

public class MyConversationalSample {

    public static void main(String[] args) {
        Transaction tran = null;
        try {
            Application app
                = ApplicationFactory.createApplication();
            IOMessage spaMessage
                = app.getIOMessage("class://mytest.jdbo.SPAMessage");
            IOMessage inputMessage
                = app.getIOMessage("class://mytest.jdbo.InMessage");
            IOMessage outputMessage
                = app.getIOMessage("class://mytest.jdbo.OutMessage");
            MessageQueue msgQueue = app.getMessageQueue();
            tran = app.getTransaction();

            // Read the SPA message
            while (msgQueue.getUnique(spaMessage)) {
                // before reading the application messages.
                if (msgQueue.getNext(inputMessage)) {
                    String inField
                        = inputMessage.getString("Message").trim();
                    try {
                        int sum = (new Integer(inField)).intValue();
                        spaMessage.setString("Message", "" + sum);
                        msgQueue.insert(spaMessage,
                            MessageQueue.DEFAULT_DESTINATION);
                        outputMessage.setString("Message",
                            "The initial value is: " + sum);
                        msgQueue.insert(outputMessage,
                            MessageQueue.DEFAULT_DESTINATION);
                    } catch (NumberFormatException e) {
                        if (inField.equalsIgnoreCase("stop")) {
                            // End the conversation
                        }
                    }
                }
            }
        } catch (DLIException e) {
            // Handle exception
        }
    }
}

```

```

        spaMessage.setString("Message",
            "Exit requested, so I am exiting");
spaMessage.setTransactionName(IOMessage.END_CONVERSATION_BLANKS);
msgQueue.insert(spaMessage, MessageQueue.DEFAULT_DESTINATION);
    }
}
}
tran.commit();
} catch (DLIException e) {
    e.printStackTrace();
    try {
        // Roll back the transaction
        if (tran != null) {
            tran.rollback();
        }
    } catch (DLIException e1) {
        e1.printStackTrace();
    }
}
}
}
}

```

Conversational transaction sequence of events

When the message is a conversational transaction, the following sequence of events occurs:

1. 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. The transaction code 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.
2. After the conversational program prepares its reply, it inserts the SPA to IMS. The program then inserts the actual text of the reply as segments of an output message.
3. IMS saves the SPA and routes the message to the input LTERM (logical terminal).
4. If the SPA insert specifies 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.
5. 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.
6. To terminate a conversation, the program places blanks in the transaction code field of the SPA and inserts the SPA to IMS. To terminate a conversation when using the IMS Java dependent region resource adapter, set the SPA transaction code to the constant `IOMessage.END_CONVERSATION_BLANKS`.
7. The conversation can also be terminated if the transaction code in the SPA is replaced by any transaction code from a nonconversational program, and the SPA is inserted to IMS. After the next terminal input, IMS routes that message to the queue of the other program in the normal way.

Related concepts

[Conversational transactions \(Communications and Connections\)](#)

Handling multi-segment messages

Message-driven applications can have multi-segment input messages. That is, more than one message needs to be read from the message queue in order to retrieve the entire message.

The following code shows how the `IOMessage` and the `MessageQueue` classes are used to retrieve multi-segment messages:

```

//Create a message queue
MessageQueue messageQueue = app.getMessageQueue();

//Create the first input message
IOMessage input1
    = app.getIOMessage("class://InputMessage1");

```

```
//Create the second input message
IOMessage input2
= app.getIOMessage("class://InputMessage2");

try {
    //Read the first message from the queue
    messageQueue.getUnique(input1);
    ...
    //Read additional messages from the queue
    while(messageQueue.getNext(input2)) {
        ...
    } catch (DLIException e) {
        ...
    }
}
```

Coding and accessing messages with repeating structures

Messages with repeating structures can be defined by using the `DLTypeInfoList` class. With the `DLTypeInfoList` class, you can specify a repeating list of fields and the maximum number of times the list can be repeated. These repeating structures can contain repeating structures.

The following code example is a sample output message that contains a set of Make, Model, and Color fields:

Sample output message with repeating structures

```
public class ModelOutput extends IMSFieldMessage {
    static DLTypeInfo[] modelTypeInfo = {
        new DLTypeInfo("Make", DLTypeInfo.CHAR, 1, 20),
        new DLTypeInfo("Model", DLTypeInfo.CHAR, 21, 20),
        new DLTypeInfo("Color", DLTypeInfo.CHAR, 41, 20),
    };
    static DLTypeInfoList[] modelTypeInfoList = {
        new DLTypeInfoList("Models", modelTypeInfo, 1, 60, 100),
    };
    public ModelOutput() {
        super(modelTypeInfoList, 6004, false);
    }
}
```

To access the nested structures that are defined in a `DLTypeInfoList` object, use a dotted notation to specify the fields of the field within a repeating structure. For example, the "Color" field in the fourth "Models" definition in the output object is accessed as "Models.4.Color" within the output message. The following code sets the fourth "Color" in the output message to "Red."

```
IOMessage output = app.getIOMessage("class://ModelOutput");
output.setString("Models.4.Color", "Red");
```

Flexible reading of multiple input messages

JMP applications can process multiple input messages that require different input data types.

The following car dealership sample application supports requests to list models, show model details, find cars, cancel orders, and record sales. Each of these requests requires different input data.

The following steps explain how to define the messages to support these requests, and how to access the messages from the application.

1. Define the primary input message.

The primary input message is the message that you pass to the `MessageQueue.getUnique` method to retrieve all of your input messages. Your primary input message must have an I/O area that is large enough to contain any of the input requests that your application might receive. It must also contain at least one field in common with all of your input messages. This common field allows you to determine the input request. In the following code example, the common field is `CommandCode`, and the maximum length of each message is 64 (the number passed to the `IMSFieldMessage` constructor):

```
public class InputMessage extends IMSFieldMessage {
    final static DLTypeInfo[] fieldInfo =
```

```

    {
        new DLTypeInfo("CommandCode",
            DLTypeInfo.CHAR, 1, 20),
    };

    public InputMessage()
    {
        super(fieldInfo, 64, false);
    }
}

```

2. Define separate input messages for each request.

Each of these input messages contains the same CommandCode field as its first field. Each of these input messages also uses an IMSFieldMessage constructor that takes an IMSFieldMessage object and a DLTypeInfo array. The IMSFieldMessage constructor allows you to remap the contents of the primary input message using the same type of information with each request; therefore, you do not copy the I/O area of the message, only a reference to this area. The following code example illustrates how to create the input messages for the requests ShowModelDetails, FindACar, and CancelOrder.

```

public class ShowModelDetailsInput extends IMSFieldMessage {
    final static DLTypeInfo[] fieldInfo = {
        new DLTypeInfo("CommandCode", DLTypeInfo.CHAR, 1, 20),
        new DLTypeInfo("ModelTypeCode", DLTypeInfo.CHAR, 21, 2)
    };

    public ShowModelDetailsInput(InputMessage inputMessage) {
        super(inputMessage, fieldInfo);
    }
}

public class FindACarInput extends IMSFieldMessage {
    final static DLTypeInfo[] fieldInfo = {
        new DLTypeInfo("CommandCode", DLTypeInfo.CHAR, 1, 20),
        new DLTypeInfo("Make", DLTypeInfo.CHAR, 21, 10),
        new DLTypeInfo("Model", DLTypeInfo.CHAR, 31, 10),
        new DLTypeInfo("Year", DLTypeInfo.CHAR, 41, 4),
        new DLTypeInfo("LowPrice", DLTypeInfo.PACKEDDECIMAL, 45, 5),
        new DLTypeInfo("HighPrice", DLTypeInfo.PACKEDDECIMAL, 50, 5),
        new DLTypeInfo("Color", DLTypeInfo.CHAR, 55, 10),
    };

    public FindACarInput(InputMessage inputMessage) {
        super(inputMessage, fieldInfo);
    }
}

public class CancelOrderInput extends IMSFieldMessage {
    final static DLTypeInfo[] fieldInfo = {
        new DLTypeInfo("CommandCode", DLTypeInfo.CHAR, 1, 20),
        new DLTypeInfo("OrderNumber", DLTypeInfo.CHAR, 21, 6),
        new DLTypeInfo("DealerNumber", DLTypeInfo.CHAR, 21, 6),
    };

    public CancelOrderInput(InputMessage inputMessage)
    {
        super(inputMessage, fieldInfo);
    }
}

```

Note the following details about the previous code examples:

- The CommandCode field is defined in every message that reads the command code. If you do not define the field, you must adjust the offsets of the following fields to account for the existence of the CommandCode in the byte array. For example, you can delete the DLTypeInfo entry for CommandCode in the CancelOrderInput class, but the OrderNumber field must still start at offset 21.
- The length of the base class InputMessage must be large enough to contain any of its subclasses. In this example, the InputMessage class is 64 bytes because the fields of the FindACarInput method require it.
- Each InputMessage subclass must provide a constructor to create itself from an InputMessage object. This constructor uses a new constructor in the IMSFieldMessage class, called a *copy constructor*.

Given this design, an application can provide message-reading logic like in the following code example.

```
while (messageQueue.getUnique(inputMessage)) {
    string commandCode=inputMessage.getString("CommandCode").trim();
    if (commandCode.equals("ShowModelDetails")) {
        showModelDetails(new ShowModelDetailsInput(inputMessage));
    } else if(commandCode.equals("FindACar")) {
        findACar(new FindACarInput(inputMessage));
    } else {
        //process an error
    }
}
```

Developing JBP applications with the IMS Java dependent region resource adapter

JBP applications are similar to JMP applications, except that JBP applications do not receive input messages from the IMS message queue. Unlike batch message processing (BMP) applications, JBP applications must be non-message-driven applications.

Symbolic checkpoint and restart

Similarly to batch message processing (BMP) applications, JBP applications can use symbolic checkpoint and restart calls to restart the application after an abend. To issue a symbolic checkpoint and restart when using the IMS Java dependent region resource adapter, use these methods of the `com.ibm.ims.dli.tm.Transaction` interface:

- `Transaction.checkpoint()`
- `Transaction.restart()`

These methods perform functions that are analogous to the DL/I system service calls: (symbolic) CHKP and XRST.

A JBP application connects to a database, makes a restart call, performs database processing, periodically checkpoints, and disconnects from the database at the end of the program. The program must issue a final commit before ending. On an initial application start, the `Transaction.restart()` method notifies IMS that symbolic checkpoint and restart is to be enabled for the application. The application then issues periodic `Transaction.checkpoint()` calls to take checkpoints. The `Transaction.checkpoint()` method allows the application to provide a `com.ibm.ims.dli.tm.SaveArea` object that contains one or more other application Java objects whose state is to be saved with the checkpoint.

If a restart is required, it defaults to the last checkpoint ID. The `Transaction.restart()` method returns a `SaveArea` object that contains the application objects in the same order in which they were inserted at checkpoint time. If the `SaveArea` object returned is `null`, this means there were no objects stored in the `SaveArea` object at checkpoint time.

Symbolic checkpoint and restart calls may also be used with GSAM data, or z/OS data sets. To restart using a basic z/OS checkpoint, you must identify the restart checkpoint.

Code sample of JBP symbolic checkpoint and restart

The following symbolic checkpoint/restart sample JBP application demonstrates the use of the checkpoint and restart functionality support with the IMS Java dependent region resource adapter.

The two symbolic checkpoint methods `checkpoint()` and `checkpoint(SaveArea saveArea)` require the application to be restarted (in the case of any abnormal end of the program) using the 4-character constant "LAST".

```
package samples.dealership.chkp_xrst;

import java.sql.*;
import java.io.*;

import com.ibm.ims.dli.DLIException;
import com.ibm.ims.dli.tm.Application;
import com.ibm.ims.dli.tm.ApplicationFactory;
import com.ibm.ims.dli.tm.SaveArea;
import com.ibm.ims.dli.tm.Transaction;
import com.ibm.ims.jdbc.IMSDataSource;

public class CheckpointRestartAutoSample {

    private SaveArea saveAreaOut;
    private Connection connection;
    private Transaction transaction;

    // The entry point of the application
    public static void main(String[] args) throws Exception {
        CheckpointRestartAutoSample crSample
            = new CheckpointRestartAutoSample();

        crSample.setup();

        crSample.runSample();

        crSample.closeDown();
    }

    // Set up for the application:
    // 1. Enable trace
    // 2. Creates connection
    void setup() throws Exception {
        this.createConnection();
        Application app = ApplicationFactory.createApplication();
        this.transaction = app.getTransaction();
    }

    void closeDown() throws Exception {
        // close the connection
        connection.close();

        // Commit the IMS DB work
        this.transaction.commit();
    }

    // Creates a connection to the auto dealership database
    void createConnection() throws Exception {
        try {
            IMSDataSource ds = new IMSDataSource();
            ds.setDriverType(IMSDatasource.DRIVER_TYPE_2);
            ds.setMetadataURL("class://samples.dealership.AUTPSB11DatabaseView");

            connection = ds.getConnection();

        } catch (SQLException e) {
            String errorMessage
                = new String("During connection creation: "
                    + e.toString());
            throw new Exception(errorMessage);
        }
    }

    void runSample() throws Exception {
        // the restart call is always the first call in a
        // checkpoint/restart application
        saveAreaOut = this.transaction.restart();

        // if the SaveArea object returned is null it
        // is a normal program start, otherwise it is a restart
        if (saveAreaOut != null) {
            // Check the SaveArea object to determine
            // where to restart from
            if (saveAreaOut.isEmpty()) {

```



```

        sqlMethod(true);
    } else {
        String str =
            (String)saveAreaOut.getObject(1);
        System.out.println("Retrieved string = "+str);
    }
} else {
    sqlMethod(false);
}
}

void sqlMethod(boolean isRestart)
    throws DLIException, SQLException {

    String sql
        = new String("SELECT * FROM Dealer.DealerSegment");
    Statement statement = connection.createStatement();
    ResultSet results = statement.executeQuery(sql);

    // this part of the code will be executed only during a normal
    // program start
    if (!isRestart && results.next()) {
        System.out.println("At first GetSegment call to the DealerDB: ");
        System.out.println("Dealer Number = "
            + results.getString("DealerNo"));
        System.out.println("Dealer Name = "
            + results.getString("DealerName"));
        System.out.println("DealerCity = "
            + results.getString("DealerCity"));
        System.out.println("DealerZip = "
            + results.getString("DealerZip"));
        System.out.println("DealerPhone = "
            + results.getString("DealerPhone"));
    }

    //String ckptid = null;
    for (int i=1; results.next(); i++) {
        System.out.println("GetSegment call to the DealerDB:");
        System.out.println("Dealer Number = "
            + results.getString("DealerNo"));
        System.out.println("Dealer Name = "
            + results.getString("DealerName"));
        System.out.println("DealerCity = "
            + results.getString("DealerCity"));
        System.out.println("DealerZip = "
            + results.getString("DealerZip"));
        System.out.println("DealerPhone = "
            + results.getString("DealerPhone"));

        // The checkpoint call, apart from storing program information,
        // causes the program to lose its position in the database
        this.transaction.checkpoint();
    }
}
}

```

Rolling back changes in a JBP application

Similar to JMP applications, a JBP application can roll back database processing and output messages any number of times during a transaction. A rollback call backs out all database processing and output messages to the most recent commit.

Use the `com.ibm.ims.dli.tm.Transaction` class to issue commit and rollback operations from your JMP application.

Related concepts

[“Restarting your program” on page 65](#)

If you issue symbolic checkpoint calls (for batch and BMP programs), you can use the Extended Restart system service request (XRST) to restart your program after an abnormal termination.

Related reference

[“Recovering databases and maintaining database integrity” on page 263](#)

You can issue checkpoints, restart programs, and maintain database integrity in your application programs.

Accessing GSAM data from a JBP application

GSAM data are frequently referred to as z/OS data sets or as flat files. This kind of data is non-hierarchical in structure. You can access data from GSAM databases from a JBP application.

The JMP application connects to a GSAM database, performs database processing, periodically commits, and disconnects from the database at the end of the application. To access the GSAM data, you will need to supply your JBP application with the Java database metadata class for that database.

If your IMS system includes an activate IMS catalog database, you can connect to the catalog instead of using a database metadata class file.

Sample metadata class for a car dealership database

The following Java code sample provides an example of the Java database metadata class.

```
package samples.dealership.gsam;

import com.ibm.ims.db.*;
import com.ibm.ims.base.*;

public class AUTOGSAMDatabaseView extends DLIDatabaseView {

    // This class describes the data view of PSB: AUTOGSAM
    // PSB AUTOGSAM has database PCBs with 8-char PCBNAME or label:
    // AUTOLPCB
    // PCBGSGAMG
    // PCBGSGAML

    // The following describes Segment:
    // DEALER ("DEALER") in PCB: AUTOLPCB ("AUTOLPCB")
    static DLTypeInfo[] AUTOLPCBDEALERArray= {
        new DLTypeInfo("DLRNO", DLTypeInfo.CHAR, 1, 4,
            "DLRNO", DLTypeInfo.UNIQUE_KEY),
        new DLTypeInfo("DLRNAME", DLTypeInfo.CHAR,
            5, 30, "DLRNAME"),
        new DLTypeInfo("CITY", DLTypeInfo.CHAR,
            35, 10, "CITY"),
        new DLTypeInfo("ZIP", DLTypeInfo.CHAR,
            45, 10, "ZIP"),
        new DLTypeInfo("PHONE", DLTypeInfo.CHAR,
            55, 7, "PHONE")
    };
    static DLISegment AUTOLPCBDEALERSegment= new DLISegment
        ("DEALER","DEALER",AUTOLPCBDEALERArray,61);

    // The following describes Segment: MODEL ("MODEL")
    // in PCB: AUTOLPCB ("AUTOLPCB")
    static DLTypeInfo[] AUTOLPCBMODELArray= {
        new DLTypeInfo("MODKEY", DLTypeInfo.CHAR, 3, 24,
            "MODKEY", DLTypeInfo.UNIQUE_KEY),
        new DLTypeInfo("MODTYPE", DLTypeInfo.CHAR, 1, 2, "MODTYPE"),
        new DLTypeInfo("MAKE", DLTypeInfo.CHAR, 3, 10, "MAKE"),
        new DLTypeInfo("MODEL", DLTypeInfo.CHAR, 13, 10, "MODEL"),
        new DLTypeInfo("YEAR", DLTypeInfo.CHAR, 23, 4, "YEAR"),
        new DLTypeInfo("MSRP", DLTypeInfo.CHAR, 27, 5, "MSRP"),
        new DLTypeInfo("COUNT1", DLTypeInfo.CHAR, 32, 2, "COUNT")
    };
    static DLISegment AUTOLPCBMODELSegment= new DLISegment
        ("MODEL","MODEL",AUTOLPCBMODELArray,37);

    // An array of DLISegmentInfo objects follows
    // to describe the view for PCB: AUTOLPCB ("AUTOLPCB")
    static DLISegmentInfo[] AUTOLPCBArray = {
        new DLISegmentInfo(AUTOLPCBDEALERSegment,DLIDatabaseView.ROOT),
        new DLISegmentInfo(AUTOLPCBMODELSegment,0),
    };

    // Warning: PCB: PCBGSGAMG has no SENSEGS
    // The following describes GSAM Record:
    // JAVGSAM1 ("JAVGSAM1") in PCB: PCBGSGAMG ("GSAMRead")
    static DLTypeInfo[] PCBGSGAMGJAVGSAM1Array= {
        new DLTypeInfo("DealerNo", DLTypeInfo.INTEGER, 1, 4),
```

```

new DLTypeInfo("DealerName", DLTypeInfo.CHAR, 5, 30),
new DLTypeInfo("ModelType", DLTypeInfo.CHAR, 35, 2),
new DLTypeInfo("ModelKey", DLTypeInfo.CHAR, 37, 24),
new DLTypeInfo("Make", DLTypeInfo.CHAR, 37, 10),
new DLTypeInfo("Model", DLTypeInfo.CHAR, 47, 10),
new DLTypeInfo("Year", "yyyy", DLTypeInfo.DATE, 57, 4),
new DLTypeInfo("MSRP", "S999999V99", DLTypeInfo.PACKEDDECIMAL, 61, 5),
new DLTypeInfo("Counter", DLTypeInfo.SMALLINT, 66, 2)
};
static GSAMRecord PCBGSAMGRecord= new GSAMRecord
("PCBGSAMGRecord",PCBGSAMGJAVGSAM1Array,80);

// An array of DLISegmentInfo objects follows
// to describe the view for PCB: PCBGSAMG ("GSAMRead")
static DLISegmentInfo[] PCBGSAMGarray = {
new DLISegmentInfo(PCBGSAMGRecord, DLIDatabaseView.ROOT)
};

// Warning: PCB: PCBGSAML has no SENSEGS
// The following describes GSAM Record:
// JAVGSAM1 ("JAVGSAM1") in PCB: PCBGSAML ("GSAMLoad")
static DLTypeInfo[] PCBGSAMLJAVGSAM1Array= {
new DLTypeInfo("DealerNo", DLTypeInfo.INTEGER, 1, 4),
new DLTypeInfo("DealerName", DLTypeInfo.CHAR, 5, 30),
new DLTypeInfo("ModelType", DLTypeInfo.CHAR, 35, 2),
new DLTypeInfo("ModelKey", DLTypeInfo.CHAR, 37, 24),
new DLTypeInfo("Make", DLTypeInfo.CHAR, 37, 10),
new DLTypeInfo("Model", DLTypeInfo.CHAR, 47, 10),
new DLTypeInfo("Year", "yyyy", DLTypeInfo.DATE, 57, 4),
new DLTypeInfo("MSRP", "S999999V99",
DLTypeInfo.PACKEDDECIMAL, 61, 5),
new DLTypeInfo("Counter", DLTypeInfo.SMALLINT, 66, 2)
};
static GSAMRecord PCBGSAMLRecord= new GSAMRecord
("PCBGSAMLRecord",PCBGSAMLJAVGSAM1Array,80);

// An array of DLISegmentInfo objects follows
// to describe the view for PCB: PCBGSAML ("GSAMLoad")
static DLISegmentInfo[] PCBGSAMLarray = {
new DLISegmentInfo(PCBGSAMLRecord, DLIDatabaseView.ROOT)
};

// Constructor
public AUTOGSAMDatabaseView() {
super("2.0","AUTOGSAM", "AUTOLPCB", "AUTOLPCB",
AUTOLPCBarray);
addDatabase("GSAMRead", "PCBGSAMG", PCBGSAMGarray);
addDatabase("GSAMLoad", "PCBGSAML", PCBGSAMLarray);
} // end AUTOGSAMDatabaseView constructor

} // end AUTOGSAMDatabaseView class definition

```

Sample JBP application for accessing a GSAM database

The following code example is a JBP application that relies on the previous code sample to access GSAM data.

```

package samples.dealership.gsam;

import java.io.*;
import java.util.Properties;
import java.math.BigDecimal;

import com.ibm.ims.dli.*;
import com.ibm.ims.dli.tm.*;

/**
 * This is an auto dealership sample application
 * demonstrating the use of the
 * GSAM database functionality support in
 * the IMS Java dependent region resource adapter.
 */
public class GSAMAuto {
private final String readOnlyGSAMPCB
= new String("GSAMRead");
private final String writeOnlyGSAMPCB

```

```

    = new String("GSAMLoad");

private PSB psb;

/**
 * The entry point of the application
 */
public static void main(String[] args) {
    GSAMAuto gsamLoadSample = new GSAMAuto();
    if (System.getProperty("com.ibm.ims.jdbcenvironment")
        == null) {
        Properties properties = System.getProperties();
        properties.put("com.ibm.ims.jdbcenvironment", "IMS");
    }

    try {
        gsamLoadSample.setup();
    } catch (Exception e) {
        e.printStackTrace();
    }

    try {
        gsamLoadSample.runSample();

        gsamLoadSample.closeDown();
    } catch (Throwable e) {
        e.printStackTrace();
    }

}

/**
 * This method does the set up for the application:
 * 1. Enable trace
 * 2. Creates dbConnection
 * 3. Creates GSAMConnection object
 * @throws IOException
 * @throws SecurityException
 * @throws DLIException
 */
void setup() throws SecurityException,
    IOException, DLIException {
    IMSConnectionSpec cSpec
        = IMSConnectionSpecFactory.createIMSConnectionSpec();
    cSpec.setDatastoreName("IMS1");
    cSpec.setDriverType(IMSConnectionSpec.DRIVER_TYPE_2);
    cSpec.setMetadataURL("class://samples.dealership.gsam.AUTOGSAMDatabaseView");
    psb = PSBFactory.createPSB(cSpec);
}

/**
 * This method does the clean up before application exit.
 * 1. Commits the database work done. IMS Java dependent
 * regions require all applications to commit before exiting.
 * @throws DLIException
 */
void closeDown() throws DLIException {
    try {
        Application app = ApplicationFactory.createApplication();
        Transaction transaction = app.getTransaction();

        // Always commit any work before exiting
        transaction.commit();
    } catch (DLIException e) {
        System.out.println("IMS commit failed. Reason: "
            + e.toString());
        throw e;
    }
}

/**
 * Demonstrates how to write to and read from a
 * GSAM database. Also shows different data types
 * being stored into the GSAM database using the
 * internal data conversion methods.
 */
void runSample() {

    final int dealerNo = 1171;
    final String dealerName = "ABC Autos";

```

```

final String modelType = "LX";
final String make = "Santro";
final String model = "Zen";
final java.sql.Date year
    = java.sql.Date.valueOf("2011-05-18");
final BigDecimal msrp = new BigDecimal(17750.00);
final short count = (short) 8;

try {
    GSAMPCB pcb1 = psb.getGSAMPCB(this.writeOnlyGSAMPCB);

    Path myGSAMRecord = pcb1.getPathForInsert();

    // Set values to individual fields in a GSAM record
    myGSAMRecord.setInt("DealerNo", dealerNo);
    myGSAMRecord.setString("DealerName", dealerName);
    myGSAMRecord.setString("ModelType", modelType);
    myGSAMRecord.setString("Make", make);
    myGSAMRecord.setString("Model", model);
    myGSAMRecord.setDate("Year", year);
    myGSAMRecord.setBigDecimal("MSRP", msrp);
    myGSAMRecord.setShort("Counter", count);

    // Insert the GSAM record data
    // and save the RSA of the record
    RSA rsa = pcb1.insert(myGSAMRecord);

    // Close the GSAM database explicitly
    // for writing/loading data
    pcb1.close();

    // Open a GSAM Connection to write the GSAM dataset
    GSAMPCB pcb2 = psb.getGSAMPCB(this.readOnlyGSAMPCB);

    // Read the GSAM record data using
    // the RSA stored earlier
    Path gsamRecord = pcb2.getUnique(rsa);

    // Print the GSAM data
    if (gsamRecord != null) {
        System.out.println("Dealer Number: "
            + gsamRecord.getInt("DealerNo"));
        System.out.println("Dealer Name: "
            + gsamRecord.getString("DealerName"));
        System.out.println("Model Type: "
            + gsamRecord.getString("ModelType"));
        System.out.println("Make: "
            + gsamRecord.getString("Make"));
        System.out.println("Model: "
            + gsamRecord.getString("Model"));
        System.out.println("Year: "
            + gsamRecord.getDate("Year"));
        System.out.println("MSRP: "
            + gsamRecord.getBigDecimal("MSRP"));
        System.out.println("Counter: "
            + gsamRecord.getShort("Counter"));

        System.out.println
            ("\nSuccessful completion of GSAM sample application");
    } else {
        System.out.println("GSAM DB is empty");
    }

} catch (DLIException e) {
    System.out.println
        ("GSAM sample failed. Reason: " + e.toString());
}
}

```

Related concepts

[“Processing GSAM databases” on page 283](#)

GSAM databases are available to application programs that can run online in IMS batch message processing (BMP) regions (message-driven or non-message-driven) or Java batch processing (JBP) regions or standalone in DLIBATCH regions.

Related reference

[“GSAM coding considerations” on page 289](#)

The calls your program uses to access GSAM databases are not the same as the DL/I calls. The system service calls that you use with GSAM are symbolic CHKP and XRST.

IMS Java dependent region resource adapter support for ICAL callout with control data

IMS transactions that are written in Java and leverage the IMS Java dependent region resource adapter can issue ICAL calls that include control data.

By using control data, you can include any information in an ICAL callout message such as the endpoint information or other routing specifications.

The support for ICAL callout with control data uses the IMS implementation of the Java Message Service (JMS) API in the IMS Java dependent region resource adapter (imsutm.jar). The IMS Java dependent region resource adapter invokes the Universal Drivers C library (DFSCLIBU, in the SDFSJLIB data set) through JNI calls to issue the calls from C to the AIBTDLI interface with the ICAL information.

The following example illustrates an IMS IMS Java dependent region resource adapter JMP application that issues a synchronous callout message with control data:

```
public static void main(String args[]) {
    // Create an IMS JMP application
    app = ApplicationFactory.createApplication();

    // Get the IMS JMS queue connection factory
    IMSQueueConnectionFactory jmsConnectionFactory = app
        .getIMSQueueConnectionFactory();
    QueueConnection jmsConnection = null;
    QueueSession jmsQueueSession = null;
    javax.jms.Queue jmsQueue = null;
    QueueRequestor jmsQueueRequestor = null;

    try {
        // Get a reference to the IMS message queue
        msgQueue = app.getMessageQueue();

        // Create an input message object
        inputMessage = app.getIOMessage("class://MyInputMessage");

        // Create an output message object
        outputMessage = app.getIOMessage("class://MyOutputMessage");

        // Retrieve messages off the queue
        while (msgQueue.getUnique(inputMessage)) {

            // Setting the JMS settings to issue an ICAL call

            // Specify the amount of time to wait for a response from an
            // ICAL call. This value
            // corresponds to the RSFLD value in the AIB
            jmsConnectionFactory.setTimeout(9999999);

            // Specify the expected size of the response message from the
            // ICAL call. This value
            // corresponds to the OAUSE value in the AIB
            jmsConnectionFactory.setResponseAreaLength(0x00000033);

            // Create the JMS queue connection
            jmsConnection = jmsConnectionFactory.createQueueConnection();

            // Create the JMS queue session
            jmsQueueSession = jmsConnection.createQueueSession(false, 1);

            // Specify the OTMA Routing descriptor which describes the
            // target that the ICAL call
            // will be sent to. This value correponse with the RSNM1 value
            // in the AIB
            jmsQueue = jmsQueueSession.createQueue("DEST0001");

            // Create the JMS queue requestor
            jmsQueueRequestor = new QueueRequestor(jmsQueueSession,
                jmsQueue);

            // Build the request area for the ICAL call
            // For synchronous program switch a BytesMessage object must be
```

```

// used
BytesMessage sendMsg = jmsQueueSession.createBytesMessage();

// The content of the request area must follow the existing
// format for synchronous program switch:
// LL + ZZ + SWITCH-TO-TRAN + TRAN-INPUT
short ll = 50;
short zz = 0;
sendMsg.writeShort(ll); // Specify the LL value
sendMsg.writeShort(zz); // Specify the ZZ value

// The name of the SWITCH-TO-TRAN is 8 bytes long and encoded in
// CP1047
// This value must be converted to bytes to be written into the
// BytesMessage object
String trancode = new String("SWTCHTRN");
sendMsg.writeBytes(trancode.getBytes("Cp1047"));

// Specify the input data for the switch to transaction
sendMsg.writeUTF(inputMessage.getString("MYINPUT"));

// Build the control data object
IMSControlArea controlArea = new IMSControlArea();
byte controlData[] = "Richard".getBytes();
controlArea.addControlDataItem("name", controlData);

// Attach the control data object to the message object
((BytesMessageImpl) sendMsg).addControlArea(controlArea);

// The length of the request area can be retrieved by calling
// the BytesMessage.getBodyLength() method
// This value corresponds to the OALEN value in the AIB
System.out.println("Request Message Length (AIBOALEN): "
    + sendMsg.getBodyLength());

// Submit the ICAL call
// >>-ICAL--aib--request_area--responseArea---control_area---<<
// For synchronous program switch, the reply message will be a
// BytesMessage object
BytesMessage replyMsg = (BytesMessage) jmsQueueRequestor
    .request(sendMsg);

// The response message will have the following format
// LL + ZZ + TRAN-OUTPUT

// Retrieve the LL field
replyMsg.readShort();

// Retrieve the ZZ field
replyMsg.readShort();

// Retrieve the output data from the switch to transaction and
// place it in the output message
outputMessage.setString("MYOUTPUT", replyMsg.readUTF());

// Send the output message back to IMS
msgQueue
    .insert(outputMessage, MessageQueue.DEFAULT_DESTINATION);

}

// Terminate the application and free up any associated resources
app.end();

} catch (Exception e) {
    // Error scenario, free up resources
    app.end();
    e.printStackTrace();
}

}

```

Programming with the Callout API

Use the `IMSCallout` API to issue IMS synchronous callout requests from JMP or JBP applications that are running in a Java dependent region. Another option is to use the Java Message Service (JMS) API, which provides limited support.

IMSCallout API

Use the `IMSCallout` class to issue Synchronous Callout (ICAL) DL/I calls directly from a Java-enabled dependent region with the IMS Universal JDBC driver `IMSCallout` API.

The following code shows a simple application using the `IMSCallout` to issue a callout request:

```
ByteBuffer icalRequestArea = null;
ByteBuffer icalResponseArea = null;
icalRequestArea = app.get31BitDirectByteBuffer(requestSize);
icalRequestArea.position(0);
icalRequestArea.put("CALLOUT FROM ICAL PGM".getBytes("CP1047"));
icalResponseArea = app.get31BitDirectByteBuffer(actualResponseBufferSize);

IMSCallout imsCallout = app.createImsCallout();

imsCallout.setOtmaDescriptor(otmaDescriptor)
.setRequestArea(icalRequestArea)
.setResponseArea(icalResponseArea)
.setSubFunction(IMSCallout.SubFunction.SENDRECV)
.setTimeout(10000)
.execute();

System.out.println(imsCallout.getReturnCode());
System.out.println(imsCallout.getReasonCode());
System.out.println(imsCallout.getErrorCodeExtension());
System.out.println(imsCallout.getOutputAreaUsed());

app.free31BitDirectByteBuffer(icalResponseArea);
app.free31BitDirectByteBuffer(icalRequestArea);
app.end();
```

ByteBuffers and Type-2 Java applications

If your JDBC application is running in Type-2 mode, you can optimize JVM storage by obtaining Byte Buffers directly. When your specified work is complete, you must release the buffers.

Obtain direct storage in the following ways:

- If your application is running in a 31-bit JVM, you can obtain Byte Buffers directly by using the **`java.nio.ByteBuffer`** class. `ByteBuffer.allocateDirect()` works only in a 31-bit JVM.

```
ByteBuffer directBuffer = ByteBuffer.allocateDirect(10);
```

- Use the Universal Drivers Application object to obtain direct 31-bit storage. If your application is running in the 64-bit JVM, you must use this method.

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);
```

Release the buffers that were obtained with the `Application` class in the following ways:

- Specifically free the buffers with the `free31BitDirectByteBuffer()` method.

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);
// Doing work

app.free31BitDirectByteBuffer(directBuffer);
```


- Use the `end()` method. IMS tracks all buffers that are allocated with the **Application** class and automatically releases them as part of the `end()` procedure.

```
Application app = new ApplicationFactory.createApplication();
ByteBuffer directBuffer = app.get31BitDirectByteBuffer(20);
// Doing work
app.end();
```

Java Message Service API

The IMSCallout API is the newer, IMS-specific way to issue synchronous callout requests from a Java dependent region. However, an IMS implementation of the Java Message Service (JMS) that uses JMS classes can provide limited support for synchronous callout requests from JMP or JBP applications.

Support is limited to the Point-to-Point (PTP) messaging domain, and is provided only for nontransacted QueueSession objects with Session.AUTO_ACKNOWLEDGE mode. Calling unsupported JMS methods or passing unsupported JMS arguments throws a JMSEException exception.

The JMS API is included in the Java SDK. To use JMP and JBP support for synchronous callout, the JMS API (`jms.jar`) file must be on your classpath. Alternatively, IMS Enterprise Suite V3.2 is available for purchase, which includes the JMS API.

Use the following steps to send a message and synchronously receive a response:

1. Create a `com.ibm.ims.jms.IMSQueueConnectionFactory` object.
2. Create a JMS QueueConnection instance by calling the `createQueueConnection` method on the `IMSQueueConnectionFactory` object.
3. Create a JMS QueueSession instance by calling the `createQueueSession` method on the QueueConnection instance. In the method call, set the input parameter values to **false** and **Session.AUTO_ACKNOWLEDGE** to specify that the generated QueueSession instance is nontransacted and runs in AUTO_ACKNOWLEDGE mode.
4. Create a queue identity by calling the `createQueue` method on the QueueSession instance. In the method call, set the input parameter value to the OTMA descriptor name for the synchronous callout operation.
5. Create a JMS QueueRequestor instance and pass in the QueueSession instance from step 3 and the Queue instance from step 4 as input parameters to the QueueRequestor constructor method.
6. Create a TextMessage instance by calling the `createTextMessage` method on the QueueSession instance from step 3. Set the string that contains the message data.
7. To send the message and retrieve a response, call the `request` method on the QueueRequestor object from step 5. In the method call, pass in the TextMessage instance from step 6. Cast the return value from the request method call to a TextMessage instance. If the call is successful, the return value is the response to the synchronous callout request.

The following example shows a simple JMP or JBP application that sends a message to an external application and synchronously receives a response message. In the example, an `IMSQueueConnectionFactory` instance has a timeout value of 10 seconds and allocates 128 KB of space to hold response messages.

```
import javax.jms.JMSEException;
import javax.jms.Queue;
import javax.jms.QueueConnection;
import javax.jms.QueueRequestor;
import javax.jms.QueueSession;
import javax.jms.Session;
import javax.jms.TextMessage;
import com.ibm.ims.jms.IMSQueueConnectionFactory;

public class IMS_Sample
{
    public static void main(String argv[])
    {
        IMSQueueConnectionFactory jmsConnectionFactory
            = new IMSQueueConnectionFactory();
```

```

QueueConnection jmsConnection = null;
QueueSession jmsQueueSession = null;
Queue jmsQueue = null;
QueueRequestor jmsQueueRequestor = null;

try {
    jmsConnectionFactory.setTimeout(1000);
    // set the timeout to 10 seconds
    jmsConnectionFactory.setResponseAreaLength(128000);
    // allocate 128k to hold the response message
    jmsConnection = jmsConnectionFactory.createQueueConnection();
    jmsQueueSession
    = jmsConnection.createQueueSession(false, Session.AUTO_ACKNOWLEDGE);
    // set session to be non-transacted and in AUTO_ACKNOWLEDGE mode
    jmsQueue = jmsQueueSession.createQueue("OTMDEST1");
    // pass in the OTMA descriptor name

    jmsQueueRequestor
    = new QueueRequestor(jmsQueueSession, jmsQueue);
    TextMessage sendMsg = jmsQueueSession.createTextMessage();
    sendMsg.setText("MyMessage");
    System.out.println("Sending message: "+sendMsg.getText());
    TextMessage replyMsg
    = (TextMessage)jmsQueueRequestor.request(sendMsg);

    System.out.println("\nReceived message: "+replyMsg.getText());
} catch (JMSException e) {
    e.printStackTrace();
}
}
}

```

Related concepts

[“Using Byte Buffers with the IMS Universal DL/I driver” on page 663](#)

The IMS Universal DL/I driver includes several methods to handle complex data types that are not fully represented by the IMS Catalog.

Related tasks

[“Implementing the synchronous callout function” on page 453](#)

To issue a synchronous callout request from your IMS application, issue the ICAL call and specify the OTMA descriptor name.

Related reference

[“Using the expandArrayResultSet property” on page 611](#)

The IMS Universal JDBC driver allows for easier access to Array datatype fields using the **expandArrayResultSet** property.

Program switching in JMP and JBP applications

IMS allows you to switch programs in JMP and JBP applications. You can perform immediate program switches in JMP and JBP applications, and you can also make a deferred program switch in a conversational JMP application.

Related concepts

[“Sending messages to other IMS application programs” on page 393](#)

A program-to-program switch occurs when an IMS application running in an IMS dependent region sends a message to another IMS application running in an IMS dependent region.

Immediate program switching for JMP and JBP applications

The IMS Java dependent region resource adapter supports immediate program switching in JMP and JBP applications. An immediate program switch passes the conversation directly to another conversational program that is specified by an alternate PCB.

When an application makes an immediate program switch, the first `MessageQueue.insert` call sends the SPA to the other conversational program, but subsequent `MessageQueue.insert` calls will send messages to the new program. The program does not return or respond to the original terminal.

The `setAlternatePCBName` method of the `com.ibm.ims.dli.tm.MessageDestinationSpec` class sets the name of the alternate PCB for the program switch. The `setAlternatePCBName` method issues the DL/I CHNG call.

To make an immediate program switch in a JMP or JBP application:

1. Call the `MessageDestinationSpec.setAlternatePCBName` method to set the name of the alternate PCB.
2. Call the `MessageQueue.insert` method to send the message to the alternate PCB.

Code sample of immediate program switching

The following code sample demonstrates how immediate program switching is performed in a JMP application.

```
package sample.jmp;

import com.ibm.ims.dli.tm.Application;
import com.ibm.ims.dli.tm.ApplicationFactory;
import com.ibm.ims.dli.tm.IOMessage;
import com.ibm.ims.dli.tm.MessageDestinationSpec;
import com.ibm.ims.dli.tm.MessageQueue;
import com.ibm.ims.dli.tm.Transaction;

public class SampleJMPIImmediatePgmSwitch {
    private static IOMessage outputMessage = null;
    private static MessageQueue msgQueue = null;
    private static Application app = null;
    private static IOMessage inputMessage = null;

    public static void main(String[] args) {
        try {
            app = (Application) ApplicationFactory.createApplication();
            msgQueue = (MessageQueue) app.getMessageQueue();
            inputMessage
                = app.getIOMessage("class://sample.jmp.InMessage");
            outputMessage
                = app.getIOMessage("class://sample.jmp.OutMessage");

            //Define Message Destinations Specs
            MessageDestinationSpec mds2
                = new MessageDestinationSpec();
            mds2.setAlternatePCBName("TPPCB1");
            mds2.setDestination("JAVTRANJ");

            String in = new String("");
            while (msgQueue.getUnique(inputMessage)){
                in = inputMessage.getString("Message").trim();
                if(in.equalsIgnoreCase("ImmediatePGMSwitch1")){
                    outputMessage.setString("Message",
                        "Running ImmediatePGMSwitch1 Call");
                    msgQueue.insert(outputMessage,
                        MessageQueue.DEFAULT_DESTINATION);

                    // Insert Message to JAVTRANJ TPPCB1: DLIWithCommit
                    outputMessage.setString("Message",
                        "Insert Message to JAVTRANJ TPPCB1");
                    msgQueue.insert(outputMessage,
                        MessageQueue.DEFAULT_DESTINATION);

                    outputMessage.setString("Message", "DLIWithCommit");
                    outputMessage.setTransactionName("JAVTRANJ");

                    // Insert message to JAVTRANJ
                    msgQueue.insert(outputMessage, mds2);

                    // Commit transaction
                    Transaction tran = app.getTransaction();
                    tran.commit();
                } else {
                    outputMessage.setString("Message",
                        "Invalid input - valid input is 'ImmediatePGMSwitch1'");
                    msgQueue.insert(outputMessage,
                        MessageQueue.DEFAULT_DESTINATION);
                    Transaction tran = app.getTransaction();
                    tran.commit();
                }
            }
        }
    }
}
```

```

    }
  } catch (Exception e) {
    e.printStackTrace();
  }
}
}

```

To perform immediate program switching in a JBP application, the steps are similar to a JMP application. In the main module, setup a `MessageDestinationSpec` instance then issue an `insert` call to another transaction. For example:

```

MessageDestinationSpec mds2 = new MessageDestinationSpec();
mds2.setAlternatePCBName("TPPCB1");
mds2.setDestination("JAVTRANJ");
...
outputMessage.setString("Message", "Some Message");
outputMessage.setTransactionName("JAVTRANJ");
msgQueue.insert(outputMessage, mds2);

```

Related concepts

[“Passing the conversation to another conversational program” on page 404](#)

A conversational program can pass the conversation to another conversational program in by performing a deferred switch or a immediate switch.

Deferred program switching for conversational JMP applications

You can make a deferred program switch in a conversational JMP application. A deferred program switch changes the transaction code in the scratchpad area (SPA) before the SPA is returned to IMS. When an application makes a deferred program switch, the application replies to the terminal and passes the conversation to another conversational application.

Use the `setTransactionName(String)` method of the `com.ibm.ims.dli.tm.IOMessage` class to specify the transaction code in the SPA.

To make a deferred program switch in a conversational JMP application:

1. Call the `insert(IOMessage)` method to send the output message to the terminal.
2. Call the `setTransactionName(String)` method to set the name of the transaction code in the SPA.
3. Call the `insert(IOMessage)` method to send the SPA to IMS.

Code sample of deferred program switching

The following code sample demonstrates how deferred program switching is performed in a JMP application.

```

package sample.jmp;

import com.ibm.ims.dli.tm.Application;
import com.ibm.ims.dli.tm.ApplicationFactory;
import com.ibm.ims.dli.tm.IOMessage;
import com.ibm.ims.dli.tm.MessageDestinationSpec;
import com.ibm.ims.dli.tm.MessageQueue;
import com.ibm.ims.dli.tm.Transaction;

public class SampleJMPDeferredPGM {
    private static IOMessage spaMessage = null;
    private static MessageQueue msgQueue = null;
    private static Application app = null;
    private static IOMessage inputMessage = null;
    private static Transaction tran = null;

    public static void main(String[] args) {
        try {
            app = ApplicationFactory.createApplication();
            spaMessage
                = app.getIOMessage("class://sample.jmp.SPAMessage");
            inputMessage
                = app.getIOMessage("class://sample.jmp.InMessage");
            msgQueue = app.getMessageQueue();

```

```

        tran = app.getTransaction();

        MessageDestinationSpec mds
            = new MessageDestinationSpec();
        mds.setAlternatePCBName("TPPCB1");
        mds.setDestination("IVTCM");

        String in = new String("");
        while (msgQueue.getUnique(spaMessage)) {
            if (msgQueue.getNext(inputMessage)) {
                in = inputMessage.getString("Message").trim();
                if (in.equalsIgnoreCase("DeferredPGMSwitch2")) {
                    inputMessage.setString("Message", spaMessage.getString("Message"));
                    msgQueue.insert(inputMessage, MessageQueue.DEFAULT_DESTINATION);

                    // Setting Deferred Program Switch
                    inputMessage.setString("Message", "Setting Deferred Program Switch");
                    msgQueue.insert(inputMessage, MessageQueue.DEFAULT_DESTINATION);

                    spaMessage.setString("Message", "SampleJMPDeferredPGM");
                    spaMessage.setTransactionName("IVTCM");
                    msgQueue.insert(spaMessage, mds);

                    inputMessage.setString("Message", "SampleJMPDeferredPGM Completed");
                    msgQueue.insert(inputMessage, MessageQueue.DEFAULT_DESTINATION);

                    tran.commit();
                } else {
                    inputMessage.setString("Message", spaMessage.getString("Message"));
                    msgQueue.insert(inputMessage, MessageQueue.DEFAULT_DESTINATION);

                    inputMessage.setString("Message",
                        "Input Message was not 'DeferredPGMSwitch2'");
                    msgQueue.insert(inputMessage, MessageQueue.DEFAULT_DESTINATION);

                    tran.commit();
                }
            }
        }
    } catch (Exception e) {
        e.printStackTrace();
    }
}
}
}

```

Related concepts

[“Passing the conversation to another conversational program” on page 404](#)

A conversational program can pass the conversation to another conversational program in by performing a deferred switch or a immediate switch.

Issuing synchronous program switch requests from a Java dependent region

IMS provides support for synchronous program switch functionality from Java message processing (JMP) or Java batch processing (JBP) applications through an IMS implementation of the Java Message Service (JMS).

To use the JMP support for synchronous program switch, the JMS API (`jms.jar`) file must be on your classpath. The JMS API is included in the Java SDK. Alternatively, you can order IMS Enterprise Suite V3.2, which includes the JMS API.

You must create an OTMA destination descriptor with the type IMSTRAN to use this function.

Restrictions:

- The IMS implementation of JMS is limited to supporting the Point-to-Point (PTP) messaging domain only. In addition, support is only provided for non-transacted `QueueSession` objects with `Session.AUTO_ACKNOWLEDGE` mode.
- The DFSYIOEO and DFSMSCEO exit routines are not called for synchronous program switch requests.
- The target transaction is not part of the RRS commit scope for the initiating application program.
- JBP applications cannot make synchronous program switch requests in a DBCTL environment.
- The target transaction has read-only access to Fast Path Main Storage Databases (MSDBs).

- The target transaction cannot be a conversational transaction.
- Synchronous program switch requests can be used in a shared queues environment only if all of the participating IMS systems have a DBRC MINVERS value of 13.1 or greater.

If the application attempts to call a JMS method that is not supported by IMS or pass an unsupported argument in a JMS method call, the API throws a `JMSEException`.

The following procedure shows the high-level programming flow to implement a synchronous program switch operation from a JMP application.

1. Create an IMS JMP application object.
2. Create an `IMSQueueConnectionFactory`.
The connection factory and application objects are reusable.
3. Get an IMS message queue reference.
4. Create input and output message objects.
5. Configure the timeout and response area size of your connection factory and create a connection.
6. From the connection, create a JMS queue session.
7. Configure the queue session object with the name of your IMSTRAN OTMA destination descriptor and create a JMS queue requestor.
8. Create a `BytesMessage` object for the ICAL call input message area.
For synchronous program switch requests, only `BytesMessage` is supported.
9. Configure the message object with the target tran code, message length information, and input data.
10. Submit the ICAL call.
11. Retrieve the output message, parse the response message length information, and get the output message data.
12. Clean up the output message and terminate the application.

This example demonstrates how to make a synchronous program switch request from a JMP application.

```
import javax.jms.BytesMessage;
import javax.jms.QueueConnection;
import javax.jms.QueueRequestor;
import javax.jms.QueueSession;
import com.ibm.ims.jms.IMSQueueConnectionFactory;
import com.ibm.ims.dli.tm.Application;
import com.ibm.ims.dli.tm.ApplicationFactory;
import com.ibm.ims.dli.tm.IOMessage;
import com.ibm.ims.dli.tm.MessageQueue;

public class SyncCalloutSample
{
    private static IOMessage inputMessage = null;
    private static IOMessage outputMessage = null;
    private static MessageQueue msgQueue = null;
    private static Application app = null;

    public static void main(String args[])
    {
        // Create an IMS JMP application
        app = ApplicationFactory.createApplication();

        // Get the IMS JMS queue connection factory
        IMSQueueConnectionFactory jmsConnectionFactory = app.getIMSQueueConnectionFactory();
        QueueConnection jmsConnection = null;
        QueueSession jmsQueueSession = null;
        javax.jms.Queue jmsQueue = null;
        QueueRequestor jmsQueueRequestor = null;

        try {
            // Get a reference to the IMS message queue
            msgQueue = app.getMessageQueue();
```

```

// Create an input message object
inputMessage = app.getIOMessage("class://MyInputMessage");

// Create an output message object
outputMessage = app.getIOMessage("class://MyOutputMessage");

// Retrieve messages off the queue
while(msgQueue.getUnique(inputMessage)) {

    // Setting the JMS settings to issue an ICAL call

    // Specify the amount of time to wait for a response
    // from an ICAL call. This value corresponds to the
    // RSFLD value in the AIB
    jmsConnectionFactory.setTimeout(999999);

    // Specify the expected size of the response message
    // from the ICAL call. This value corresponds to the
    // OAUSE value in the AIB
    jmsConnectionFactory.setResponseAreaLength(0x00000033);

    // Create the JMS queue connection
    jmsConnection = jmsConnectionFactory.createQueueConnection();

    // Create the JMS queue session
    jmsQueueSession = jmsConnection.createQueueSession(false, 1);

    // Specify the OTMA Routing descriptor which describes the
    // target that the ICAL call will be sent to. This value
    // correponse with the RSNM1 value in the AIB
    jmsQueue = jmsQueueSession.createQueue("DEST0001");

    // Create the JMS queue requestor
    jmsQueueRequestor = new QueueRequestor(jmsQueueSession, jmsQueue);

    // Build the request area for the ICAL call
    // For synchronous program switch a BytesMessage object must be used
    BytesMessage sendMsg = jmsQueueSession.createBytesMessage();

    // The content of the request area must follow the existing
    // format for synchronous program switch:
    // LL + ZZ + SWITCH-TO-TRAN + TRAN-INPUT
    short ll = 50;
    short zz = 0;
    sendMsg.writeShort(ll); // Specify the LL value
    sendMsg.writeShort(zz); // Specify the ZZ value

    // The name of the SWITCH-TO-TRAN is 8 bytes long and encoded in CP1047
    // This value must be converted to bytes to be written into the BytesMessage
    String trancode = new String("SWTCHTRN");
    sendMsg.writeBytes(trancode.getBytes("Cp1047"));

    // Specify the input data for the switch to transaction
    sendMsg.writeUTF(inputMessage.getString("MYINPUT"));

    // The length of the request area can be retrieved by calling the
    // BytesMessage.getBodyLength() method
    // This value corresponds to the OALEN value in the AIB
    System.out.println("Request Message Length (AIBOALEN): " +
sendMsg.getBodyLength());

    // Submit the ICAL call
    // For synchronous program switch, the reply message will be a BytesMessage
    BytesMessage replyMsg = (BytesMessage)jmsQueueRequestor.request(sendMsg);

    // The response message will have the following format
    // LL + ZZ + TRAN-OUTPUT

    // Retrieve the LL field
    replyMsg.readShort();
}

```

```

        // Retrieve the ZZ field
        replyMsg.readShort();

        // Retrieve the output data from the switch to transaction and place it in
        // the output message
        outputMessage.setString("MYOUTPUT", replyMsg.readUTF());

        // Send the output message back to IMS
        msgQueue.insert(outputMessage, MessageQueue.DEFAULT_DESTINATION);

    }

    // Terminate the application and free up any associated resources
    app.end();

}

catch(Exception e) {
    // Error scenario, free up resources
    app.end();
    e.printStackTrace();
}

}

}

```

This example demonstrates how to make a synchronous program switch request from a JBP application.

```

package testcases.udb.opendb.t2;

import javax.jms.BytesMessage;
import javax.jms.QueueConnection;
import javax.jms.QueueRequestor;
import javax.jms.QueueSession;

import com.ibm.ims.jms.IMSQueueConnectionFactory;
import com.ibm.ims.dli.tm.Application;
import com.ibm.ims.dli.tm.ApplicationFactory;

public class SyncPgmSwitchFromJBPSample {
    private static Application app = null;

    public static void main(String args[]) {
        // Create an IMS JBP application
        app = ApplicationFactory.createApplication();

        // Get the IMS JMS queue connection factory
        IMSQueueConnectionFactory jmsConnectionFactory = app.getIMSQueueConnectionFactory();
        QueueConnection jmsConnection = null;
        QueueSession jmsQueueSession = null;
        javax.jms.Queue jmsQueue = null;
        QueueRequestor jmsQueueRequestor = null;

        try {

            // Setting the JMS settings to issue an ICAL call

            // Specify the amount of time to wait for a response from an
            // ICAL call. This value corresponds to the RSFLD value in the AIB
            jmsConnectionFactory.setTimeout(999999);

            // Specify the expected size of the response message from the
            // ICAL call. This value corresponds to the OAUSE value in the AIB
            int expectedResponseLength = 50;
            jmsConnectionFactory.setResponseAreaLength(expectedResponseLength);

            // Create the JMS queue connection
            jmsConnection = jmsConnectionFactory.createQueueConnection();

            // Create the JMS queue session
            jmsQueueSession = jmsConnection.createQueueSession(false, 1);

            // Specify the OTMA Routing descriptor which describes the
            // target that the ICAL call will be sent to. This value
            // corresponds with the RSNM1 value in the AIB
            jmsQueue = jmsQueueSession.createQueue("MYDEST");

```



```

// Create the JMS queue requester
jmsQueueRequestor = new QueueRequestor(jmsQueueSession,
    jmsQueue);

// Build the request area for the ICAL call
// For synchronous program switch a BytesMessage object must be used
BytesMessage sendMsg = jmsQueueSession.createBytesMessage();

// The content of the request area must follow the existing
// format for synchronous program switch:
// LL + ZZ + SWITCH-TO-TRAN + TRAN-INPUT
short ll = 50;
short zz = 0;
sendMsg.writeShort(ll); // Specify the LL value
sendMsg.writeShort(zz); // Specify the ZZ value

// The name of the SWITCH-TO-TRAN is 8 bytes long and encoded in CP1047
// This value must be converted to bytes to be written into the
// BytesMessage object
String trancode = new String("SWTCHTRN");
sendMsg.writeBytes(trancode.getBytes("Cp1047"));

// Specify the input data for the switch to transaction
sendMsg.writeUTF("MYINPUT");

// The length of the request area can be retrieved by calling
// the BytesMessage.getBodyLength() method
// This value corresponds to the OALEN value in the AIB
System.out.println("Request Message Length (AIBOALEN): "
    + sendMsg.getBodyLength());

// Submit the ICAL call
// >>-ICAL--aib--request_area--responseArea-----<<
// For synchronous program switch, the reply message will be a
// BytesMessage object
BytesMessage replyMsg = (BytesMessage) jmsQueueRequestor
    .request(sendMsg);

// The response message will have the following format
// LL + ZZ + TRAN-OUTPUT

// Retrieve the LL field
replyMsg.readShort();

// Retrieve the ZZ field
replyMsg.readShort();

// Retrieve the Data field
byte[] messageBody = new byte[(int) replyMsg.getBodyLength()];
replyMsg.readBytes(messageBody);

// Terminate the application and free up any associated resources
app.end();

} catch (Exception e) {
    // Error scenario, free up resources
    app.end();
    e.printStackTrace();
}

}

}

```

Related concepts

[Synchronous program switch requests \(Communications and Connections\)](#)

IBM Enterprise PL/I for z/OS and Java language interoperability

With IBM Enterprise PL/I for z/OS support for Java and PL/I language interoperability, you can write PL/I and Java applications that run in a Java dependent region and call PL/I programs already in production. .

You can use this support to write and run applications specific to your workloads and business environment.

For example:

- Write front-end message processing applications in Java that call back-end PL/I applications to process databases.
- Build PL/I applications with a main routine that can call Java routines.

You can access PL/I code in a JMP or JBP region by using PL/I language syntax features. For example:

- Define classes with methods and data that are implemented in PL/I.
- Create instances of Java and PL/I classes.
- Run methods on Java and PL/I objects.
- Write classes that inherit from Java classes or other PL/I classes.
- Define and run overloaded methods.
- From PL/I, call services that are provided by the JNI to obtain Java-oriented capabilities.

To learn more about building IBM Enterprise PL/I for z/OS applications that run in an IMS dependent region, refer to [Enterprise PL/I for z/OS](#).

IBM Enterprise COBOL for z/OS interoperability with JMP and JBP applications

With the IBM Enterprise COBOL for z/OS support for COBOL and Java language interoperability, you can write Java and Object-Oriented (OO) COBOL applications that execute in a Java dependent region and invoke existing COBOL programs.

With this support, you can:

- Call an object-oriented (OO) COBOL application from a Java application by building the frontend application, which processes messages, in Java, and the back end, which processes databases, in OO COBOL.
- Build an OO COBOL application containing a main routine that can invoke Java routines.

You can access COBOL code in a JMP or JBP region because Enterprise COBOL provides object-oriented language syntax that enables you to:

- Define classes with methods and data implemented in COBOL
- Create instances of Java and COBOL classes
- Invoke methods on Java and COBOL objects
- Write classes that inherit from Java classes or other COBOL classes
- Define and invoke overloaded methods

In IBM Enterprise COBOL for z/OS programs, you can call the services provided by the JNI to obtain Java-oriented capabilities in addition to the basic OO capabilities available directly in the COBOL language.

In IBM Enterprise COBOL for z/OS classes, you can code CALL statements that interface with procedural COBOL programs. Therefore, COBOL class definition syntax can be especially useful for writing wrapper classes for procedural COBOL logic, enabling existing COBOL code to be accessed from Java.

Java code can create instances of COBOL classes, invoke methods of these classes, and can extend COBOL classes.

Related Reading: For details building applications that use IBM Enterprise COBOL for z/OS and that run in an IMS dependent region, see *Enterprise COBOL for z/OS Programming Guide*.

Related concepts

[“Overview of the IMS Java dependent regions” on page 673](#)

The IMS Java dependent regions are two types of IMS dependent regions that provide a Java Virtual Machine (JVM) environment for Java applications: Java message processing (JMP) regions and Java batch processing (JBP) regions.

IBM Enterprise COBOL for z/OS backend applications in a JMP or JBP region

When you define an object-oriented (OO) COBOL class and compile it with the IBM Enterprise COBOL for z/OS compiler, the compiler generates a Java class definition with native methods and the object code to implement the native methods. After compiling the class, you can create an instance and invoke the methods of the compiled class from a Java program that runs in a JMP or JBP region.

For example, you can define an OO COBOL class with the appropriate DL/I call in COBOL to access an IMS database.

To make the implementation of this class available to a Java application running with IMS:

1. Compile the COBOL class with the IBM Enterprise COBOL for z/OS compiler to generate a Java source file, which contains the class definition, and an object module, which contains the implementation of the native methods.
2. Compile the generated Java source file with the Java compiler to create the application class file.
3. Link the object module into a dynamic link library (DLL) in the z/OS file system (zFS) file (.so).
4. Update the application class path (ibm.jvm.application.class.path) for the JMP or JBP region to allow access to the Java class file.
5. Update the library path for the JMP or JBP region to allow access to the DLL.

IBM Enterprise COBOL for z/OS frontend applications in a JMP or JBP region

The object-oriented syntax of IBM Enterprise COBOL for z/OS enables you to build COBOL applications with a main method, which can be run directly in a JMP or JBP region.

The JMP or JBP region locates, instantiates, and invokes the main method of an OO COBOL application in the same way it does for the main method of a Java application.

You can write an application for an JMP or JBP region entirely with OO COBOL, but a more likely use for a frontend COBOL application is to call a Java routine from a COBOL application.

When running within the JVM of an JMP or JBP region, the IBM Enterprise COBOL for z/OS runtime support automatically locates and uses the JVM to invoke methods on Java classes.

A frontend OO COBOL application with a main routine that runs in a JMP or JBP region has the same requirements as a Java program that runs in a JMP or JBP region.

Accessing Db2 for z/OS databases from JMP or JBP applications

A JMP or JBP application can access Db2 for z/OS using the latest JDBC driver for Db2 for z/OS.



Attention: If you access a Db2 for z/OS database using both Java and COBOL in the same application, you might experience unexpected behavior, but only if the commit or rollback processing is done in COBOL while active cursors are in the Java portion.

The JMP or JBP region that the application is running in can also be defined with Db2 for z/OS attached by the DB2 Recoverable Resource Manager Services attachment facility (RRSAF).

Accessing Db2 for z/OS data from a JMP or JBP application is like accessing IMS data. When writing a JMP or JBP application that accesses Db2 for z/OS data, consider both the differences from IMS database access and the differences from accessing Db2 for z/OS data in other environments:

- You must create a DB2 plan for each PSB (typically each Java application) that is used to access Db2 for z/OS.
- You can have only one active Db2 for z/OS connection open at any time.
- If you are using the type-2 JDBC drivers for Db2 for z/OS, you must use the default connection URL in the application program. For example, `jdbc:db2os390:` or `db2:default:connection`.

- If you are using the type-4 DB2 JDBC drivers, you can use a specific connection URL in the application program.
- To commit or roll back work, use the `Transaction.commit` method or the `Transaction.rollback` method.
 - For JMP applications, the `Transaction.commit` method commits all work, including SQL calls. Calling the `Transaction.commit` and `Transaction.rollback` methods does not automatically reset the connection to Db2 for z/OS. The connection to Db2 for z/OS is reset when you issue a `MessageQueue.getUnique` call.
 - For JBP applications, the `Transaction.commit` method commits SQL calls.
- Because RRSAF is the coordinator, you cannot use the `Connection.setAutoCommit` or `Connection.commit` method of the JDBC driver for Db2 for z/OS.

Related concepts

[“Overview of the IMS Java dependent regions” on page 673](#)

The IMS Java dependent regions are two types of IMS dependent regions that provide a Java Virtual Machine (JVM) environment for Java applications: Java message processing (JMP) regions and Java batch processing (JBP) regions.

Related tasks

[Preparing your system to use the DB2 Attach Facility \(Communications and Connections\)](#)

[DB2: Installing the IBM Data Server Driver for JDBC and SQLJ](#)

Chapter 42. 31-bit COBOL and 64-bit Java interoperability

You can enable interoperability between 31-bit COBOL code and 64-bit Java code by adding value **JVM=3164** to an IMS dependent region's parameter list, installing required software, and configuring the dependent region. Supported dependent region types include MPR, BMP, IFP, JBP, and JMP.

Adding parameter **JVM=3164** indicates to an IMS dependent region that separate Language Environment instances are to be established for each addressing mode: a primary 31-bit Language Environment enclave, and a secondary 64-bit Language Environment enclave that hosts the Java virtual machine (JVM). When JVM=3164, IMS uses module DFSJVM36.

Software requirements

This interoperability requires specific functions from the COBOL compiler, the Language Environment, and the 64-bit IBM Java Software Development Kit (SDK). Required software includes the following items:

- IMS APAR PH37140.
- Language Environment APAR PH28966, which provides the base support for interoperability.
- Language Environment APAR PH40444, which provides a fix for S422 ABEND (BMP region) or U4083 ABEND out of CELQLIB during IMS dependent region termination.
- COBOL APAR PH37101.
- Java IBM SDK, Java Technology Edition V8.0.6.36 (FixPack 36) or newer.

Additionally, consider the following important information regarding Java, COBOL (including Object-Oriented COBOL), and IMS.

Java considerations

The following Java considerations for IMS applications are based on *IBM SDK, Java Technology Edition*. To learn more, see [Using 31-bit native C or C++ code with the 64-bit Java VM \(z/OS only\)](#).

Limitations

- The 31-bit native application must be a single-thread. The 64-bit Java application can be multi-threaded but only a single Java application thread is allowed to interoperate across the address mode (AMODE) boundary. If a second application thread tries to load a 31-bit library or call a 31-bit C or C++ function, a `java.lang.UnsatisfiedLinkError` is usually thrown.
- You must compile the 31-bit native application with standard CALL linkage; XPLINK is not supported.
- Language Environment Condition handling support is not available across the AMODE boundary. Specifically, `-XCEEHDLR`, `-Xsignal:userConditionHandler=percolate`, and related options are not supported. Language Environment conditions are not converted into Java exceptions.
- Signal chaining is not supported across the AMODE boundary.
- Environment variables are copied across the AMODE boundary on the first transition only. Subsequent updates in one AMODE environment are not propagated back to the other; instead, the variables are tracked separately in each AMODE environment. To learn more, see *Environment variables propagation to secondary Language Environment* in [Introduction to AMODE 31 and AMODE 64 programs interoperability](#).

COBOL considerations

The following COBOL considerations are based on COBOL technical note *An enhancement in COBOL/JNI interface*. To learn more, see [COBOL programs in AMODE 31 interacting with Java programs in AMODE 64](#).

Object-Oriented (OO) COBOL

Important: This enhancement cannot be used by COBOL programs that contain object-oriented features such as INVOKE statements, OBJECT-REFERENCE data items, or class definitions. An exception is JNIEnvPtr, a COBOL special register that can be used by OO COBOL and non-OO COBOL programs to make Java Native Interface (JNI) calls.

COBOL programming considerations

In a mixed 31/64-bit environment, references to Java types must be declared as PIC 9(18) COMP-5 in COBOL (that is, as an 8-byte data item). The OBJECT-REFERENCE type cannot be used because this data item maps to an incorrect 4-byte size in AMODE 31 COBOL programs. This applies to COBOL data items that interact with JNI types such as jobject, jclass, jfieldID, jmethodID, jarray, and jarray<type> when they are used as parameters in JNI calls.

Note that accesses to EXTERNAL data items and EXTERNAL files are not serialized.

Building the mixed AMODE COBOL/Java interoperability application

It is recommended that you build and run COBOL JNI applications in z/OS UNIX.

Compiling

To compile the COBOL program, use the cob2 command in a z/OS UNIX shell. Specify the RENT, DLL and PGMNAME (LONGMIXED) compiler options. The THREAD option is not mandatory if only one COBOL program is active in the run-unit at any time.

If the COBOL program includes the file JNI.cpy by using a COPY statement, specify the include subdirectory of the COBOL install directory (typically /usr/lpp/cobol/include) in the search order for copybooks. You can specify the include subdirectory by using the -I option of the cob2 command or by setting the SYSLIB environment variable.

Linking

To prepare the COBOL program for execution, link the object file with the following two DLL side files to create an executable module:

- libjvm31.x, which is provided with your 64-bit IBM Java Software Development Kit: \$JAVA_HOME/lib/s390x/j9vm.
- igzxjni2.x, which is provided in the lib subdirectory of the cobol directory in the z/OS UNIX file system (typically /usr/lpp/cobol/lib). This DLL side file is also available as the member IGZXJNI2 in the SCEELIB PDS.

Do not link with the sidedeck igzcjava.x, which is used only by object-oriented features and is not supported in the mixed AMODE environment. Linking both igzxjni2.x and igzcjava.x in the same program can lead to ABENDs.

Note that \$JAVA_HOME is the home directory of the 64-bit IBM Java Software Development Kit.

IMS considerations

The following list includes known limitations, things to consider, and changes in behavior between a 31-bit only environment and a mixed address mode environment running 31-bit COBOL and 64-bit Java.

- The UMASK environment option is not supported on the 64-bit Java environment.
- Unhandled exceptions, such as divide by zero, in a COBOL first application will cause the Language Environment to terminate and trigger the IMS dependent region controller to go through termination without issuing a U0101 ABEND with reason code 'B'x (11 decimal).
- For Java application programs, only the main thread can properly interact with the IMS control region. This is an existing restriction for Java application programs.

Dependent region setup

Dependent region setup for 31-bit COBOL and 64-bit Java interoperability is slightly different than for interoperability between 31-bit COBOL and 32-bit Java. Differences include a different positional startup option for the dependent region's JCL (JVM=3164), a different IMS module (DFSJVM36), a different Java

DLL (libjvm31.so), and a different JVM option (-XX:+Enable3164Interoperability), which is automatically included by DFSJVM36 when starting the JVM. IMS module DFSJVM36 will print messages that begin with "DFSJVM36:" to the dependent region's job log.

To set up a dependent region that supports 31-bit COBOL code and 64-bit Java code, make the following changes:

- Update the dependent region's JCL to include the new option **JVM=3164**. You can tailor the JCL examples in the following Samples section to start the different dependent region types.
- Update the dependent region's JCL to include the IMS SDFSJLIB as part of the STEPLIB concatenation. This data set contains new IMS code that is needed for proper execution of the dependent region.
- Update the LIBPATH and CLASSPATH to point to the 64-bit Java installation directories to be able to start the 64-bit JVM. See the samples of the ENVIRON= member and the JVMOPMAS= member or the STDENV DD card contents in the following section.

Samples

IMS dependent region procedure samples are available in the following topics:

- [DFSMPR procedure](#)
- [IMSBATCH procedure](#)
- [IMSFP procedure](#)
- [DFSJBP procedure](#)
- [DFSJMP procedure](#)

The following samples show JVM and other configuration information specified in the ENVIRON=DFSJVM36 and JVMOPMAS=DFSJVM36 members:

```
*****
HLQ.PROCLIB(DFSJVM36)
*****
LIBPATH=/tmp:>
javahome/J8.0_64/lib/s390x/j9vm/>
javahome/J8.0_64/lib/s390x/compressedrefs:>
javahome/J8.0_64/lib/>
/usr/lpp/ims/ims15/imsjava/lib/:
```

```
*****
HLQ.PROCLIB(DFSJVM36)
*****

-Xms256m
-Xmx1024m
-Djzos.logging=E
-Djzos.merge.sysout=false
-Djava.class.path=/tmp:>
javahome/J8.0_64/>
javahome/J8.0_64/lib/>
javahome/J8.0_64/lib/ext:>
/tmp/udbtest-0.0.1.jar:>
/usr/lpp/ims/ims15/imsjava/imsudb.jar
```

The alternative to using ENVIRON= and JVMOPMAS= is to use the STDENV DD card on the JCL for starting the dependent region. The following is a sample of the STDENV member that can be specified on the STDENV DD card:

```
*****
HLQ.PDS(STDENV01)
*****
LP=/tmp
LP="$LP":javahome/J8.0_64/lib/s390x/j9vm/
LP="$LP":javahome/J8.0_64/lib/s390x/compressedrefs/
LP="$LP":javahome/J8.0_64/lib/
LP="$LP":/usr/lpp/ims/ims15/imsjava/lib/
export LIBPATH="$LP":
CP=/tmp
```

```

CP="$CP":javahome/J8.0_64/
CP="$CP":javahome/J8.0_64/lib/
CP="$CP":javahome/J8.0_64/lib/ext
CP="$CP":/tmp/udbtest-0.0.1.jar
CP="$CP":/usr/lpp/ims/ims15/imsjava/imsudb.jar
export CLASSPATH="$CP":
IJO="-Xms256m"
IJO="$IJO -Xmx512m"
IJO="$IJO -Djzos.logging=E"
IJO="$IJO -Djzos.merge.sysout=false"
export IBM_JAVA_OPTIONS="$IJO "

```

Troubleshooting common issues

You see message DFS529E indicating an invalid value on the JVM exec parameter. For example:

```
DFS529E MPP00001.REGION.MPP00001 - AN INVALID VALUE IS SPECIFIED ON THE JVM EXEC PARAMETER
```

- This message indicates that the job MPP00001 to start dependent region MPP00001 contains an invalid value for JVM=. This can indicate that IMS APAR PH37140 is not on the system, or that the value specified for parameter JVM= is not 31 (default) or 3164.

You see multiple messages that begin with 'DFSJVM00: Option' or 'DFSJVM64: Option' in the dependent region's job output.

- The DFSJVM00 messages indicate that the dependent region was started without the parameter JVM=3164, or with the parameter in the wrong position. The DFSJVM64 messages indicate that a JMP or JBP dependent region type was started with JVM=64.

You see message DFSJVM00 in the dependent region's job output, indicating a failure to open libjvm.so:

```
DFSJVM00: dlopen(libjvm.so) failed; function(JNI_CreateJavaVM); EDC5254S An AMODE31 application is attempting to load an AMODE64 DLL load module.
```

- This indicates that you are starting a dependent region without the parameter JVM=3164 or with the parameter in the wrong position, and that you properly configured the LIBPATH to point to the 64-bit Java SDK.

You see message DFSJVM36 in the dependent region's job output, indicating a failure to open libjvm31.so:

```
DFSJVM36: dlopen(libjvm31.so) failed; function(JNI_CreateJavaVM);
```

- This indicates that the LIBPATH is not pointing to the 64-bit Java SDK, or that the 64-bit Java SDK is not at V8.0.6.36 or newer. To resolve, ensure that you installed Java SDK V8.0.6.36 or newer and that the LIBPATH points to it.

You see messages DFSJVM36 and CEL4R064 in the dependent region's job output, indicating a failure to start the JVM and no support for CEL4R064:

```

DFSJVM36: JNI_CreateJavaVM() failed
CEL4R064: CEL4R064 runtime support not found.
JNI_CreateJavaVM failed. CEL4R064 rc: 6 - CEL4R064 runtime support not found.

```

- The CEL4R064 message produced by the Java Native Interface (JNI) indicates that the Language Environment CEL4R064 runtime support function is not present on the system. To resolve, ensure that the required Language Environment APARs are installed on the system.

The dependent region terminates with U0101 ABEND and reason code 'E'x (14 decimal).

- Reason code 'E'x (14 decimal) for U0101 ABEND indicates that the dependent region was started without the SDFSJLIB in the STEPLIB concatenation. The SDFSJLIB contains files that are needed by the IMS dependent region to support 31-bit COBOL and 64-bit Java interoperability.

Related concepts

[Language environments \(Application Programming\)](#)

Related reference

[APARM= parameter for procedures \(System Definition\)](#)

[DFSJBP procedure \(System Definition\)](#)

[DFSJMP procedure \(System Definition\)](#)

[DFSJVMEV \(JVM environment settings member\) \(System Definition\)](#)

[DFSMPR procedure \(System Definition\)](#)

[IMSBATCH procedure \(System Definition\)](#)

[IMSFP procedure \(System Definition\)](#)

[JVM= parameter for procedures \(System Definition\)](#)

Part 7. PL/I top-down development for IMS Enterprise Suite SOAP Gateway web services

You can generate a PL/I application template from a web service description language (WSDL) document that describes the operations and messages of a web service by using IBM Developer for System z. Then use the PL/I segmentation APIs to add your business logic to the generated application before enabling this application as a web service running on IMS Enterprise Suite SOAP Gateway.

IBM Developer for System z has a batch processor that can generate an IMS web service application template in Enterprise PL/I from a WSDL file. The approach is known as the *top-down* approach, where the application to serve as a web service is generated from a web service description file. The batch processor generates the data structure, XML converters, and PL/I application template file from the WSDL file that describes the web service.

Chapter 43. WSDL-to-PL/I segmentation APIs for adding business logic in generated PL/I templates

IBM Developer for System z generates metadata to record the high-level relationships between the WSDL file that you supply and the generated artifacts based on the PL/I segmentation APIs. Use this set of APIs to add your business logic before enabling it as a web service on IMS Enterprise Suite SOAP Gateway.

Important: The IRZPWSIO segmentation APIs in IBM Developer for System z V9.0.1 or older versions is renamed to DFSPWSIO in IMS. Starting with IMS Enterprise Suite V3.1 SOAP Gateway and IBM Developer for System z V9.0.1.1, you must use the DFSPWSIO segmentation APIs in IMS.

Language structures are written to a single include file that begins with an operation-to-language-structure dictionary comment. The metadata file is in XML format and is used by the batch processor to generate XML converters, deployment metadata, and template programs. Annotations are added to the generated source code to describe the relationships between the generated language structures and the XML schemas from which they are derived.

The annotations appear as language comments immediately preceding the definitions of the language structures or language structure members to which they apply. The WSDL2PLI component in IBM Developer for System z uses a set of segmentation APIs in the DFSPWSH include file that is required during compilation of the PL/I program. These APIs define how to consume and produce IMS messages. The DFSPWSH include file in the SDFSSMPL data set provides the PL/I binding and offers pointers to the data structures.


For each operation on the specified service and port, the following is generated:

- The PL/I structure(s) for operation input message
- The XSD to PL/I mapping session for operation input message
- The PL/I to XSD mapping session for operation output message
- The PL/I structure(s) for operation output message

For each operation in the WSDL, an `operationNameHandler` procedure and an `operationNameImpl` procedure are created in the generated template. The `operationNameHandler` procedure contains protocol logic while the `operationNameImpl` procedure is ready to be filled out and customized with your business logic.

A DFSPWSHK user exit is provided to demonstrate how you can inspect, modify, or replace the buffer which contains the current data structure that is being processed by the DFSPWSIO segmentation APIs. Customize the DFSPWSHK user exit if you need to inspect, modify, or replace data structures as they are processed by the DFSPWSIO APIs. The DFSPWSHK user exit allows you to keep track of whether or not a data structure has been modified or replaced via the `dfs_in_struct_state` and `dfs_out_struct_state` parameters. These and other parameters are documented in detail in the sample procedure and in *IMS Version 15.3 Exit Routines*.

The DFSPWSHK user exit is invoked by DFSPWSIO segmentation APIs when the APIs are called by XML converters in IMS Connect or by message processing programs. Therefore the DFSPWSHK exit must be compiled and linked into a data set such that IMS Connect and Message Processing Regions could find it by using the standard MVS search order (for example, STEPLIB and LINKLIST). The following diagram demonstrates how the DFSPWSHK exit is called by both the XML converters in IMS Connect and by message processing programs.



- **DFSXSETS:** This API is used by the PL/I XML converter in IMS Connect to set a language structure that contains either the SOAP header, SOAP body, or SOAP fault. This API does not copy language structures into the IMS Connect output buffer until it is instructed to do so through the parameter @dfs_commit_structs. It is an error to deallocate or otherwise invalidate structure pointers passed to the API via parameter @dfs_struct_ptr before instructing the API to commit (copy) all structures to the IMS Connect output buffer.
- **DFSQGETS:** This API is used by the message processing PL/I program to get that language structure from the IMS Message Queue by using the CEETDLI in-interface. The language structure contains either a SOAP header, SOAP body, or SOAP fault. All language structures must be retrieved from the IMS Message Queue prior to setting the language structures by using the DFSQSETS API.
- **DFSQSETS:** This API sets a language structure that contains either the SOAP header, SOAP body, or SOAP fault. This API does not insert language structures into the IMS Message Queue until it is instructed to do so via parameter @dfs_commit_structs. Therefore it is an error to deallocate or otherwise invalidate structure pointers that are passed to the API through the @dfs_struct_ptr parameter before instructing the API to commit (insert) all structures to the IMS Message Queue.
- **DFSXGETS:** This API is used by the PL/I XML converter in IMS Connect to get a language structure that contains either the SOAP header, SOAP body, or SOAP fault. Since the IMS Message Queue is not available to XML Conversion in IMS Connect, language structures are retrieved from the IMS Connect input buffer. The expected format of the IMS Connect input buffer is an LLZZDATA byte stream.

For more information about the generated DFSPWSH include file, see *IMS Version 15.3 Application Programming APIs*.

Related reference

[Include file DFSPWSH \(Application Programming APIs\)](#)

[WSDL-to-PL/I segmentation APIs exit routine \(DFSPWSHK\) \(Exit Routines\)](#)

Chapter 44. Sample of a generated PL/I application template

IMS PL/I application requires a message protocol and segmentation API (DFSPWSIO) by the service provider message processing programs (MPPs) and XML converters.

Here is an example operation defined in the FAST247.wsd1 included in this sample. The operation here is a check balance operation that takes in a request and returns a response. A SOAP fault element is also defined.

```
<wsdl:operation name="CheckBalanceOperation">
  <soap:operation soapAction="CheckBalanceOperation" style="document" />
  <wsdl:input name="CheckBalanceRequest">
    <soap:body parts="CheckBalancePart" use="literal" />
  </wsdl:input>
  <wsdl:output name="CheckBalanceResponse">
    <soap:body parts="CheckBalancePart" use="literal" />
  </wsdl:output>
  <wsdl:fault name="ServiceExceptionFault">
    <soap:fault use="literal" name="ServiceExceptionFault" />
  </wsdl:fault>
</wsdl:operation>
```

The top-down application development tooling creates the corresponding operation in the PL/I application template:

```
CheckBalanceOperationImpl: procedure(iopcb_mask_ptr, checkBalanceReques
t_ptr, checkBalanceResponse_ptr, ServiceException_ptr) internal;
  dcl iopcb_mask_ptr pointer byvalue;
  dcl checkBalanceRequest_ptr pointer byvalue;
  dcl checkBalanceResponse_ptr pointer byaddr;
  dcl ServiceException_ptr pointer byaddr;

  return;
end CheckBalanceOperationImpl;
```


Chapter 45. Trace output for WSDL-to-PL/I segmentation APIs

Trace information for the segmentation APIs is typically written to standard out and therefore can be found in the job log of the Message Processing Region.

Each API has a trace mode that when enabled writes out information from the message header and a two-column hex dump of the language structure. These dumps can be helpful, because an IMS Connect Recorder Trace shows only the first 670 bytes of an IMS message.

The following is an example of source code:

```
01: /* Invoke API DFSQSETS to set the SOAP body language
02: * structure and commit it to the IMS Message Queue.
03: */
04: @dfs_struct_name      = 'getteam_1_0Response';
05: @dfs_struct_ptr       = getteam_1_0Response_ptr;
06: @dfs_struct_size      = storage(getteam_1_0Response);
07: @dfs_commit_structs   = '1'b;
08: @dfs_cee_feedback_ptr = addr(@dfs_cee_feedback);
09: @dfs_debug            = '1'b;
10:
11: @return_code =
12:   DFSQSETS(@dfs_async_msg_header_ptr,
13:   @dfs_iopcb_mask_ptr, @dfs_soap_body_struct,
14:   @dfs_struct_name, @dfs_struct_ptr,
15:   @dfs_struct_size, @dfs_commit_structs,
16:   @dfs_cee_feedback_ptr, @dfs_debug);
17:
18: if (@return_code != @dfs_success) then do;
19:   display('MYMPP#handle_getteam():
20:   || 'ERROR, DFSQSETS @dfs_soap_body_struct, '
21:   || '@return_code: ' || trim(@return_code) || '.');
22:   return;
23: end;
```

The corresponding trace output for this source code is as follows:

```
...: DFSPWSIO#DFSQSETS() @20140415152643909 :...
o @dfs_async_msg_header_ptr: 877656904.
o @dfs_iopcb_ptr: 110672.
o @dfs_struct_type: 2.
o @dfs_struct_name: getteam_1_0Response.
o @dfs_struct_ptr: 878837800.
o @dfs_struct_size: 150274.
o @dfs_commit_struct: 1.
o @dfs_cee_feedback_ptr: 875679616.
o DFSQSETS#setBodyStruct()
o body_struct_ptr: 878837800.
o body_struct_size: 150274.
o body_struct_ptr(1:body_struct_size):
00000000: 000001F4 00000001 00000002 000DE296 |...4.....So|
00000010: 86A3A681 998540E3 85A2A300 00000000 |ftware Test....|
00000020: 00000000 00000000 00000000 00000000 |.....|
...
```

Chapter 46. Limitations and restrictions of the segmentation APIs

The APIs are designed to support SOAP header, body, and fault structures, but currently the APIs implement only the SOAP body and fault structures.

Part 8. IMS Transaction Manager Resource Adapter

You can use the IMS Transaction Manager Resource Adapter (also known as the IMS TM resource adapter) to create Java Platform, Enterprise Edition (Java EE, previously known as J2EE) applications to access IMS transactions over the Internet, as well as to make callout requests to external Java EE applications from IMS applications that run in IMS dependent regions.

Using this resource adapter within a WebSphere or Rational development environment, you can:

- Develop components of business processes in support of service-oriented architecture
- Create Java EE applications from Java beans
- Develop service-based applications

The applications can then be deployed on application servers such as WebSphere Application Server, WebSphere Process Server, WebSphere Transformation Extender, or IBM Integration Bus.

Chapter 47. IMS Transaction Manager Resource Adapter overview

The IBM IMS Transaction Manager Resource Adapter (also known as the IMS TM resource adapter) is used by Java applications, Java Platform, Enterprise Edition (Java EE, previously known as J2EE) applications or web services to access IMS transactions that are running on host IMS systems.

The IMS TM resource adapter implements the Java EE Connector Architecture (JCA), which connects enterprise information systems (EISs) such as IMS to the Java EE platform. JCA provides your applications with the qualities of service that can be provided by a Java EE application server, such as connection management, transaction management, and security management. In addition, the IMS TM resource adapter implements the JCA Common Client Interface (CCI), a programming interface that you use in your application to communicate with IMS Transaction Manager.

The IMS TM resource adapter can be used with any generic Java EE 1.4 or later application server, such as IBM WebSphere Application Server and IBM WebSphere Application Server Liberty Profile. The IMS TM resource adapter can also enable a Java application to access IMS transactions that are running on a host IMS system. In addition, the IMS TM resource adapter enables an IMS application to act as a client to invoke applications in a Java EE server.

Although the IMS TM resource adapter is intended for use primarily by Java applications or web services that submit transactions to IMS, the IMS TM resource adapter can also be used by services that submit IMS commands to IMS or for IMS applications to invoke external Java EE applications.

Related information

J2EE Connector Architecture Specification See the J2EE Connector Architecture Specification at <http://java.sun.com/j2ee/connector/index.jsp>.

WebSphere Application Server request metrics tool information For more information about request metrics, see IBM Documentation for WebSphere Application Server.

IMS TM Resource Adapter V15 release notes IMS TM Resource Adapter V15 release notes

Components of the IMS TM resource adapter

The IMS TM resource adapter includes a development-time component and a runtime component.

The IMS TM resource adapter runtime component must be deployed on IBM WebSphere Application Server (z/OS or distributed), IBM WebSphere Integration Developer, IBM Integration Bus, or IBM WebSphere Process Server.

The IMS TM resource adapter also includes a development component that lets you create your application by using an integrated development environment (IDE), such as IBM Rational Application Developer for WebSphere Software, IBM Rational Software Architect, WebSphere Integration Developer, IBM WebSphere Transformation Extender, and IBM Integration Bus. The development component is included with the optional Java EE Connector Architecture (J2C) feature in these IDEs.

You can use the J2C wizard in these IDEs to create J2C applications, EJB components, and Web services, either as stand-alone programs, or as added functionality to existing applications. You can also use this wizard to dynamically import a specific version of the resource adapter if it is not already deployed in your workspace. Even though you can code your own application by using the Common Client Interface (CCI) of the Java EE Connector Architecture, use of an IDE significantly simplifies the development process.

Runtime process of the IMS TM resource adapter

The communication between the IMS TM resource adapter and IMS are through IMS Connect, which communicates with IMS Open Transaction Manager Access (OTMA).

When an initiating client (a Java application or a web service) issues a request to access an IMS transaction, the IMS TM resource adapter communicates with IMS Connect through TCP/IP or Local Option connections. IMS Connect then sends the transaction request to IMS OTMA by using the cross-system coupling facility (XCF), 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, message-driven bean (MDB), or web service (also known as the IMS callout function), the callout requests from IMS applications are placed in the IMS OTMA hold queue. The Java application in WebSphere Application Server is set up to start and obtain a connection to IMS Connect through the IMS TM resource adapter. The IMS TM resource adapter polls IMS Connect to retrieve the callout requests from the hold queue. The Java application processes the request, and returns any response data to IMS by issuing a normal IMS transaction request. The response might be returned in the same transaction or a different transaction, depending on whether IMS application issued a synchronous or asynchronous callout request.

The following figure illustrates the IMS TM resource adapter runtime process for the first scenario, where an initiating client issues a request to access an IMS transaction.




Figure 110. IMS TM resource adapter runtime process

IMS TM resource adapter features

IMS TM resource adapter Version 15 provide the key functionality for Java applications to access various types of IMS transactions with secure socket connections.

The base functionality available in IMS TM resource adapter includes:

- Component-managed and container-managed security.
- Pooling and reuse of connections.
- Both commit mode 1 and commit mode 0 IMS transactions.
- Secure Sockets Layer (SSL) communication between the IMS TM resource adapter and IMS Connect.
- Support for SSL null encryption
- Use of RACF keyrings as SSL keystores and truststores.

- Retrieval of output messages queued as the result of a failed commit mode 0 interaction or by insertion to an alternate Program Control Block (PCB).
- Retrieval of asynchronous output with an alternate client ID when using shareable persistent socket connections
- Control of whether undelivered output for commit mode 0 interactions on shareable persistent socket connections is queued or discarded. This function is controlled by the `purgeAsyncOutput` property.
- Specification of the name of a destination for undelivered output for commit mode 0 interactions on shareable persistent socket connections. This function is controlled by the `reRoute` flag and `reRouteName` properties.
- Conversational processing.
- Support for commit mode 1, sync level CONFIRM applications
- Support for PL/I IMS applications

With Rational Application Developer (or other WebSphere and Rational development environments that include the required version of Rational Application Developer), you can generate Java EE and Web services applications to use the IMS TM resource adapter to invoke your PL/I IMS applications.

- Global transaction and two-phase-commit support.
- Run-as-thread identity support.

Related information

[IMS TM Resource Adapter V15 release notes](#)[IMS TM Resource Adapter V15 release notes](#)

New features in IMS TM Resource Adapter Version 15

V15.1 adds the support for distributed network security credentials in IMS 15.

Support for IMS 15 distributed network security credentials

IMS TM Resource Adapter can pass the original, distributed user credentials to be audited in IMS log records with both inbound and outbound (callout) requests. This feature requires IMS 15 and IMS Connect 15, and the application server, WebSphere Application Server or WebSphere Liberty, must be configured to enable the network security credentials.

An extendable Java Authentication and Authorization Service (JAAS) login module is provided with IMS TM resource adapter to enable network security credentials to be passed from a Java EE application to IMS. You must install and link the JAAS module to your application before network security credentials can be passed to IMS and audited in IMS log records.

The activation specification is enhanced with the `resumeTpipeNsc` property to enable IMS TM resource adapter to support network security credentials in IMS synchronous callout messages. To enable IMS TM resource adapter to support network security credentials in asynchronous callout messages, the IMS interaction specification is enhanced with the `setResumeTpipeNSC` property.

Distributed network security credentials

[“IMS TM resource adapter security” on page 797](#)

The Java EE Connector Architecture (JCA) specifies that the application server and the Enterprise Information System (EIS) must collaborate to ensure that only authenticated users are able to access the EIS.

[“Enabling support for distributed network security credentials” on page 810](#)

IMS TM resource adapter can pass the original, distributed network security credentials, including the network session ID and the network user ID, between Java EE applications and IMS. The network security credentials that are passed by IMS TM resource adapter are written to IMS log records.

[“Configuring WebSphere Application Server for distributed network security credentials” on page 810](#)

You can enable IMS to audit the network security credentials that are passed to and from a Java EE application that runs on WebSphere Application Server.

[“Configuring WebSphere Liberty for distributed network security credentials” on page 812](#)

You can enable IMS to audit the network security credentials that are passed to and from a Java EE application that runs on WebSphere Liberty.

Supported platforms

The IMS TM resource adapter runtime component supports WebSphere Application Server on various distributed platforms and on z/OS.

The IMS TM resource adapter runtime component can run on WebSphere Application Server on the following platforms:

- z/OS
- Windows
- AIX®
- HP-UX
- Linux®
- Linux on System z
- Solaris

For other generic Java EE application servers, see the documentation for the specific application server for its supported platforms.

Supported software configurations

The IMS TM resource adapter is based on Java EE Connector Architecture (JCA) Version 1.5.

The following table lists the supported software configurations for each version of TM resource adapter. Supported configurations are paired 1:1 with each version of IMS. To view the supported software configurations for IMS TM resource adapter Version 15, see IBM Documentation.

Table 110. Supported software configurations by IMS TM resource adapter Version 15

Software	Supported versions
IMS with integrated IMS Connect ^{“1” on page 726}	• IMS Version 15 • IMS Version 14
IBM WebSphere Application Server	Version 8.5 or later ^{“2” on page 726???}
IBM WebSphere Liberty servers	Version 19 or later ^{“2” on page 726???}
IBM Integration Bus	Version 10
Non-IBM Java EE application servers	Any generic Java EE 1.7 or later certified application servers ^{“2” on page 726}

1. The integrated IMS Connect function can coexist with earlier or newer versions of IMS. See the *Release Planning Guide* for the version of IMS that you are using for IMS Connect coexistence considerations and coexistence APARs.
2. Generic Java EE 1.7 or later certified application servers, are supported if the IMS TM resource adapter installation verification program (IVP) can run successfully. For these application servers, only a subset of IMS TM resource adapter functions are supported unless additional conditions are met. See the [restrictions topic](#) for usage restrictions. For generic Java EE 1.7 or later certified application servers, IBM support is provided for specific functions if:
 - The problem can be recreated in the supported versions of WebSphere Application Server or WebSphere Application Server Liberty Profile.
 - The problem is confirmed to be caused by the IMS TM resource adapter through diagnostic traces.

See the requirements and restrictions topics for more information about support for specific functionality on these supported software configurations.

For a more detailed description of the compatibility of supported architecture, development environments, and runtime environments, including the use of WebSphere Process Server and IBM Integration Bus, see the IMS TM resource adapter supported development environments and runtime environments technote.

Related tasks

[Verifying installation by using the installation verification program](#)

Use the IMS TM resource adapter installation verification program (IVP) to verify that an enterprise application can successfully access the target IMS system from the application server by using the IMS TM resource adapter.

Related reference

[Requirements for the IMS TM resource adapter](#)

Some IMS TM Resource Adapter features require TCP/IP connections (no Local Option).

[Restrictions for the IMS TM resource adapter](#)

IMS TM Resource Adapter has some restrictions with their support of JCA 1.5 implementations.

Related information

[Technote on supported development environments and runtime environments](#) For supported development environments and runtime environments, see this technote.

Requirements for the IMS TM resource adapter

Some IMS TM Resource Adapter features require TCP/IP connections (no Local Option).

The following list describes other requirements:

- IMS TM Resource Adapter Version 15 supports Java SDK version 8, Java 11, Java 17, and Java 21.
- The support for complex data formats by using the WebSphere Transformation Extender requires WebSphere Transformation Extender Version 8.2.0.2 and later.
- If the synchronous callout request is issued from an IMS application in the IMS Java dependent region, the Java Message Service (JMS) API in the Java SDK is required.

See the IMS TM Resource Adapter download website or the [IMS TM Resource Adapter release notes page](#) for more information on the new features and enhancements available in each release level.

Related concepts

[IMS TM resource adapter features](#)

IMS TM resource adapter Version 15 provide the key functionality for Java applications to access various types of IMS transactions with secure socket connections.

Restrictions for the IMS TM resource adapter

IMS TM Resource Adapter has some restrictions with their support of JCA 1.5 implementations.

- IMS TM Resource Adapter does not provide the JCA 1.5 Kerberos support.
- For a complete description of the restrictions for building composite business applications that invoke IMS conversational transactions, see the IMS Connect conversational support topic in *IMS Communications and Connections* information.
- New features that are added after IMS TM Resource Adapter Version 10 require TCP/IP connections (no Local Option).
- For generic Java EE application servers and WebSphere Liberty servers, two-phase commit (2PC), MFS functions, and conversational and global transactions (XA) are not supported. These features are supported only when the IMS TM resource adapter runs in IBM WebSphere Application Server.

Important: The information provided in this set of IMS TM resource adapter documentation is mainly based on the use of WebSphere Application Server. IMS TM resource adapter functions that are not

supported in WebSphere Application Server Liberty Profile are not supported in other generic Java EE application servers.

WebSphere Application Server platform configurations and communication protocol considerations

Use TCP/IP as your communication protocol between IBM WebSphere Application Server and IMS Connect.

The IMS TM resource adapter can be deployed to WebSphere Application Server for distributed platforms (AIX, HP-UX, Linux, Linux on System z, Solaris, or Windows) and to WebSphere Application Server for z/OS. The IMS TM resource adapter communicates with IMS Connect by using the TCP/IP communication protocol. .

The following diagram shows several possible deployment scenarios for the IMS TM resource adapter, and the supported communication protocols.




Figure 111. IMS TM resource adapter connections

The following table describes the communication protocols that are supported for the different platform configurations. All platform configurations support global transactions (two-phase commit).

WebSphere Application Server platform used with the IMS TM resource adapter	Supported communication protocols
AIX	TCP/IP
HP-UX	TCP/IP
Linux	TCP/IP
Linux on System z	TCP/IP
Solaris	TCP/IP
Windows	TCP/IP
z/OS	TCP/IP

To use global transaction support for your IMS and WebSphere Application Server, see the topic on global transaction and two-phase commit support process for more information.

Related concepts

[Global transaction and two-phase commit support process](#)

A Java EE-compliant application server uses a Java transaction manager to communicate and coordinate among the application components and the resource managers.

Related information

[Configuring IMS connection factories](#)

You use the IMS connection factory in your application to interact with the IMS TM resource adapter. Through the IMS connection factories, you create pre-configured connections to the IMS Transaction Manager (IMS TM).

Chapter 48. Installing the IMS TM resource adapter runtime component

Installing the IMS TM resource adapter runtime component consists of extracting the contents of the installation archive file to a target installation directory for later deployment to an application server, such as WebSphere Application Server.

Prerequisite:

Follow the preparation steps in [“Preparing to use the IMS TM resource adapter”](#) on page 731 to determine the version of the IMS TM resource adapter that you need, and [download the runtime component](#).

The IMS TM resource adapter is available as both a .tar file and a .zip file as follows:

Table 111. Installation files for the IMS TM resource adapter in different formats for different platforms

File name	Description
icoxxxx.zip	IMS TM resource adapter runtime component for distributed in a compressed format, where xxxx is the version of IMS TM resource adapter that you are installing.
icoxxxx.tar	IMS TM resource adapter runtime component for distributed in a compressed format, where xxxx is the version of IMS TM resource adapter that you are installing.
icoxxxxzos.tar	IMS TM resource adapter runtime component for z/OS in a compressed format, where xxxx is the version of IMS TM resource adapter that you are installing.

To install the IMS TM resource adapter:

1. Extract and install the archive file for your platform:
 - [“Extracting the compressed file for installation on distributed platforms”](#) on page 733.
 - [“Extracting the compressed file for installation on z/OS”](#) on page 734.
2. Verify the extracted content.
3. Install the resource adapter RAR file:
 - [“Installing the resource adapter on WebSphere Application Server”](#) on page 735
 - [“Installing the resource adapter on WebSphere Liberty servers”](#) on page 737
4. Create a connection factory:
 - [“Creating a connection factory in WebSphere Application Server”](#) on page 736
 - [“Configuring the connection factory for WebSphere Liberty servers”](#) on page 738
5. [“Verifying installation by using the installation verification program”](#) on page 739.

Preparing to use the IMS TM resource adapter

Before you use the IMS TM resource adapter, you must first determine the version of the IMS TM resource adapter and the development environment you need to use.

1. Determine which version of the IMS TM resource adapter you want to use.

The version of the IMS TM resource adapter that you use depends on whether you need the new features that are added to IMS TM Resource Adapter and the version of your IMS.

 - [“Supported software configurations”](#) on page 726
2. Determine which integrated development environment (IDE) you will use.
3. If you plan to deploy the IMS TM resource adapter to an application server other than IBM WebSphere Application Server, including WebSphere Liberty servers, International Components for Unicode (ICU)

Version 4.4.2 or later is required. This component is included in a special runtime package for WebSphere Liberty servers on the [IMS TM resource adapter download site](#).

4. Review the [potential migration issues](#) and decide if the issues apply to you.
5. Download the runtime component for the version that you choose.
The runtime component is available for downloads from the [IMS TM resource adapter download site](#). It is also available for installation on z/OS by using SMP/E.
6. If you are upgrading from an older version of the IMS TM resource adapter, follow the steps in [“Installing IMS TM resource adapter service and updates”](#) on page 748.

Potential migration issues for the IMS TM resource adapter

You might need to modify existing applications that use older connector frameworks, architectures, or APIs when you upgrade to newer versions of the IMS TM resource adapter

Deprecated functions and support

- The SYNC_RECEIVE_ASYNCOUTPUT interaction verb was deprecated in IMS TM Resource Adapter Version 10, and is replaced by SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT. No migration step is required.

See the topic on "Asynchronous output programming model" for the differences between the new interaction verbs SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT. These two verbs provide finer control of the interactions when checking for the output on the IMS OTMA asynchronous hold queue.

Removed functions and support

- The Common Connector Framework (CCF) is not supported. Applications that use CCF must be migrated to the JCA specification.
- You might encounter warnings if you run applications that use non-managed connections or were generated by using a previous version of JCA 1.0-based IMS TM Resource Adapter with the latest versions of IMS TM Resource Adapter. The TransactionResourceRegistration property available in IMS TM Resource Adapter Version 9 was removed from the custom properties of an IMS TM resource adapter connection factory. This property supported dynamic transaction enlistment for JCA 1.0-based resource adapters. This function was incorporated into the JCA 1.5-based resource adapters, so the TransactionResourceRegistration property is no longer needed.
- MFS web services applications that were generated by WebSphere Studio Application Developer Integration Edition are not supported by IMS TM Resource Adapter.

Related concepts

[Asynchronous output programming model](#)

Use this programming model to retrieve output that has been queued by IMS.

Updating the IMS TM resource adapter

Use the **Update RAR** function in the administrative console in WebSphere Application Server Version 7 and later to install the new version of the IMS TM resource adapter.

To update the IMS TM resource adapter version without deleting and reconfiguring all existing connection factories, use the **Update RAR** function in the WebSphere Application Server administrative console and update the class path accordingly.

1. In the WebSphere Application Server administrative console, click to expand **ResourcesResource Adapters**.
2. Click **Resource adapters**.
3. Click the check box to select the IMS TM resource adapter that you want to update and click **Update RAR**.




Figure 112. Updating the RAR for new versions of the resource adapter

4. In Step 1, Select either **Local file system** or **Remote file system**, and click **Browse** to navigate to where the new IMS TM resource adapter RAR file is.
For example, on a local file system, the RAR file is located in `install_path/IBM/IMS/IC0xx/Vxxxx/JCA15/imsxxxx.rar`
5. Click **Next**.
6. In Step 2, verify the version numbers of the existing and the new resource adapters and click **Next**.
7. In Step 3, configure any new properties introduced by the new version of resource adapter as prompted. Click **Next**.
8. Confirm the changes in the summary step, and click **Finish**.
You are returned to the Resource adapter page.
9. Click the name of the resource adapter you updated.
10. In the class path field under General Properties, update the RAR file name.
For example, if you are upgrading to Version 14.1 of the IMS TM resource adapter, change the class path to: `${CONNECTOR_INSTALL_ROOT}/imsico1410.RAR`
11. Click **OK**.
You are returned to the Resource adapter page.
12. Click **Save** at the top of the page.

The IMS TM resource adapter RAR file is updated. All existing connection factories now use the new version of the RAR.

Extracting the compressed file for installation on distributed platforms

On distributed platforms, you install the IMS TM resource adapter by extracting the installation files from a compressed file to the correct location for later deployment to the application server.

Download, and then extract the compressed file contents to a target installation directory.

One of the many options, some platform-dependent, for extracting the contents of the installation archive is as follows:

```
tar -xvf ic0xxx.tar install_path
```

The tar program is available as part of the operating system in UNIX and Linux variants. Versions of the tar program are also available from a variety of sources as free software program downloads. Note that some extraction tools are able to handle .tar files as well as .zip files.

Successful completion of the extraction places the contents of the compressed file in the directories `install_path/IBM/IMS/IC0xxx/Vxxxx`, where `install_path` is the target installation path of your choice and `Vxxxx` is the version number of the IMS TM resource adapter installation compressed file that you downloaded. For example, `c:\Program Files\IBM\IMS\IC0xxx\Vxxxx` on Windows or `/opt/IBM/IMS/IC0xxx/Vxxxx` for AIX.

Proceed to verify the extracted directory structure and file contents.

Extracting the compressed file for installation on z/OS

For z/OS platform, install the IMS Java On Demand features that are included in IMS, or extract the z/OS TAR file that you download from the IMS TM Resource Adapter website.

If you install the IMS Java On Demand features that are included in IMS, SMP/E installation instructions are in the *Program Directory for Information Management System (IMS) Transaction and Database Servers*.

If you download the z/OS TAR file from the [IMS download page](#), to install the IMS TM resource adapter:

1. Transfer the downloaded TAR file to a UNIX System Services file system in binary mode.
2. Use the tar program to extract the contents of the TAR file to a target installation directory.

For example:

```
tar -xvf icoxxxx.tar
```

Successful extraction and installation, either by using the SMP/E process or by manually expanding the contents of the downloaded TAR file, places the files in *install_path/Vxxxx* where *install_path* is */usr/lpp/ims/icoxxxx*. Vxxxx is the IMS TM resource adapter version number of the TAR file that you extracted the files from, or the APAR or PTF that is used for the installation.

Proceed to [verify the extracted directory structure and file contents](#).

Verifying file contents for the IMS TM resource adapter runtime component

After you extracted the .tar or .zip file for the IMS TM resource adapter runtime component, verify that you have the required files in the correct directories.

File name	Description
/JCA15/imsxxxx.rar	The resource adapter archive (RAR) file for the IMS TM resource adapter.
/JCA15/imsicoivp.ear	The enterprise archive (EAR) file for the IMS TM resource adapter installation verification program (IVP). This IVP verifies the proper deployment of the IMS TM resource adapter in WebSphere Application Server.
/JCA15/ imsicocalloutivp.ear	The enterprise archive (EAR) file for the IMS TM resource adapter callout IVP. This IVP includes a message-driven bean (MDB) application that can be used to verify that, using the IMS TM resource adapter, the application can receive a callout request from the host IMS system and return a response for a synchronous callout request.
/JCA15/jboss/ imsxxxx.rar	The JBoss resource adapter archive (RAR) file for the IMS TM resource adapter. Must only be used with application servers using Java EE.
/JCA15/twas/ imsxxxx.rar	The traditional Websphere Application Server resource adapter archive (RAR) file for the IMS TM resource adapter. Must only be used with application servers using Java EE.
/Jakarta/imsxxxx.rar	The Jakarta resource adapter archive (RAR) file for the IMS TM resource adapter.
/Jakarta/ imsicoivp.ear	The enterprise archive (EAR) file for the IMS TM resource adapter installation verification program (IVP). This IVP verifies the proper deployment of the IMS TM resource adapter in WebSphere Application Server.

File name	Description
/Jakarta/ imsicocalloutivp.ear	The enterprise archive (EAR) file for the IMS TM resource adapter callout IVP. This IVP includes a message-driven bean (MDB) application that can be used to verify that, using the IMS TM resource adapter, the application can receive a callout request from the host IMS system and return a response for a synchronous callout request.
/Jakarta/jboss/ imsxxxx.rar	The JBoss resource adapter archive (RAR) file for the IMS TM resource adapter. Must only be used with application servers using JakartaEE.
/Jakarta/twas/ imsxxxx.rar	The traditional Websphere Application Server resource adapter archive (RAR) file for the IMS TM resource adapter. Must only be used with application servers using JakartaEE.
/licenses/ license_xx.txt	License agreements in all supported languages. Users must accept all the terms and conditions in these license agreements before installing or using the IMS TM resource adapter. For the z/OS platform, these license agreements are formatted as EBCDIC-encoded text files in the z/OS compressed files (imsxxxxzos.tar and imsxxxxzos.tar) for direct viewing in the z/OS file system (zFS).
/README.html	The readme file with information on new features and bug fixes.

Note: Additional files are provided in the .tar or .zip file containing Jakarta equivalent components.

If your extracted files do not have the same directory structure and file contents, you must extract the contents of the .tar or .zip file again into the correct directories.

Proceed to install the extracted RAR file for the IMS TM resource adapter:

- [“Installing the resource adapter on WebSphere Application Server” on page 735](#)
- [“Installing the resource adapter on WebSphere Liberty servers” on page 737](#)

Installing the resource adapter on WebSphere Application Server

Use the WebSphere Application Server administrative console to deploy the IMS TM resource adapter RAR file.

Important:

1. If you are installing the IMS TM resource adapter archive (RAR) on a generic application server, additional class libraries might be needed for use with the IMS TM resource adapter. For class libraries that are not included in the IMS TM resource adapter installation archive, contact your product representative for the development environment of your choice for information regarding the licensing and location of any needed JAR files.
2. When you use the **Install RAR** dialog to install a RAR file, the scope you define on the Resource Adapters page has no effect on where the RAR file is installed. You can install RAR files only at the node level, which you specify on the Install RAR page. To set the scope of an RAR file to a specific cluster, or server, after you install the RAR file at each node level, create a copy of the RAR file with the appropriate cluster or server scope.

Prerequisites:

- You must have installed the IMS TM resource adapter RAR to a file system that is accessible to your WebSphere Application Server.
- Start the WebSphere Application Server and log in to the administrative console.

To deploy the RAR on a WebSphere Application Server:

1. In the navigation pane in the administrative console (also known as the Integrated Solutions Console), click **Resources > Resource Adapters > Resource adapters**.




Figure 113. Resource adapter configuration in WebSphere Application Server

2. If a RAR that you no longer want to use is already installed, delete it first. Stop and restart the WebSphere Application Server. Repeat the previous step to return to the resource adapter configuration page.
3. In the content pane, click **Install RAR**.
You might need to scroll down in the pane to locate the button.
The Install RAR File page displays.
4. Click **Local file system** or **Remote file system**, depending on the location of the RAR file.
 - If you choose **Local file system**, click **Browser** to locate and specify the IMS TM resource adapter RAR file that is installed, for example, `install_path/IBM/IMS/IC0xx/Vxxxxx/JCA15/imsxxxxx.rar`.
5. Click **Next**.
6. On the **Configuration** page, type a name for the RAR, for example, `imstmravxxxx`. Click **OK**.
7. Optionally, create a copy of the RAR file with a different scope level. After you install the RAR file at each node level, you can create another copy of the file that has a specific server or cluster as the scope for that file. See [Installing a resource adapter archive](#).
8. To save your results in the master configuration, click **Save** at the top of the panel.

The IMS TM resource adapter runtime RAR file is deployed.

You can proceed to [create a connection factory for the IMS TM resource adapter](#).

Related tasks

[Running your application on a stand-alone WebSphere Application Server](#)

You can test your application by using a stand-alone WebSphere Application Server. You must first export your Java EE application as an Enterprise Application Archive (EAR) file and then install the EAR file on the server.

Creating a connection factory in WebSphere Application Server

After you install and deploy the IMS TM resource adapter on WebSphere Application Server, create a connection factory for the IMS TM resource adapter, if it does not already exist.

The connection factory is used by one or more applications to create connections between IMS and IMS Connect.

Prerequisite: Start the WebSphere Application Server, and log in to the WebSphere Application Server administrative console.

To create a connection factory for the IMS TM resource adapter:

1. In the WebSphere Application Server administrative console, click to expand **Resources > Resource Adapters**.
2. Click **Resource adapters**.
3. Click the name of the IMS TM resource adapter that you deployed and want to create a connection factory for.
4. In the Additional Properties section, select **J2C connection factories**.
5. Click **New**, and in the Name field, type a name for the connection factory.
For example, myIMSTMRA.
6. Specify the JNDI name if necessary.
7. Click **Apply**.
8. Click **Save** at the top of the page.
9. Click the name of the connection factory that you just created, myIMSTMRA.
The Configuration page opens.
10. In the Configuration page, in the column under Additional Properties, click **Custom properties**.
11. Click the fields to fill in the values to configure your connection factory. Host name, port number, and data store name must be configured.
12. When you are done configuring the properties, click **Save** at the top of the page.
13. Stop the server and restart it to see the new connection factory.

You have created a connection factory for your IMS TM resource adapter.

You can proceed to [verify the installation by using the installation verification program \(IVP\)](#).

Installing the resource adapter on WebSphere Liberty servers

Download the IMS TM resource adapter RAR file and set up the WebSphere Application Server Liberty Profile environment before you install the resource adapter.

The following software is required:

- IMS TM Resource Adapter Version 13 or later runtime package for WebSphere Liberty Profile
- WebSphere Application Server Version 8.5.5.4 Liberty Profile or later, with the JCA feature installed

Prerequisites:

1. Your WebSphere Liberty server runtime environment must be properly installed and set up.
Download the JCA Connection Architecture 1.6 feature. For details and download options, see [WebSphere Liberty for developers](#).
2. Download the IMS TM resource adapter runtime component archive file for WebSphere Application Server Liberty Profile from the [IMS TM resource adapter download site](#). Select the version of the runtime offering that you want to use, and then select the IMS TM Resource Adapter vxxxx runtime component for WebSphere Liberty Profile.

To install the IMS TM resource adapter, you need to edit the `server.xml` file. WebSphere Application Server Liberty Profile does not provide a graphical administrative console. In general, you can use WebSphere Application Server Developer Tools for Eclipse or a text editor to edit the `server.xml` file.

To install the resource adapter, take the following steps:

1. Open the `server.xml` file of your WebSphere Liberty server instance in your editor.
2. Enable the JCA 1.6 and JNDI 1.0 features by adding these features to the `<featureManager>` section:

```
<featureManager>
  ...
  <feature>jca-1.6</feature>
  <feature>jndi-1.0</feature>
  ...
</featureManager>
```

3. Install the IMS TM resource adapter by adding a `<resourceAdapter>` entry:

```
<resourceAdapter id="IMS_adapter_id" location="IMS_adapter_RAR_file_location">
```

IMS_adapter_id is used to differentiate the resource adapter from other resource adapters that are installed. This value is used later when you configure the connection factory to connect to IMS.

IMS_adapter_RAR_file_location is the absolute path, including the resource adapter filename, to the IMS TM resource adapter RAR file.

The `server.xml` might look as follows:

```
<server description="myLibertyServer">
  <!-- Enable features -->
  <featureManager>
    <feature>jsp-2.2</feature>
    <!-- Added jca-1.6 and jndi-1.0 features required by the IMS
         TM resource adapter -->
    <feature>jca-1.6</feature>
    <feature>jndi-1.0</feature>
  </featureManager>

  <!-- To access this server from a remote client add a host attribute
       to the following element, e.g. host="*" -->
  <httpEndpoint id="defaultHttpEndpoint" httpPort="9080"
    httpsPort="9443" />

  <resourceAdapter id="IMSTMRA"
    location="C:\IBM\IMS\ico1410\IBM\IMS\IC014\V1410\JCA15\imsico1410.rar" />

</server>
```

In this example, IMSTMRA is the resource adapter ID, and it points to the `imsico1410.rar` file.

4. Save your changes.
5. Start your WebSphere Libertyserver.

The following J2CA7001I message is issued in the console, indicating a successful installation.

```
J2CA7001I: Resource adapter imsicovxxx installed in 0.735 seconds.
```

Next, you need to [configure a connection factory](#) to connect to your IMS host system.

Related information

[Configuring and deploying a basic JCA resource adapter \(WebSphere Application Server Liberty Core V8.5.5\)](#)

Configuring the connection factory for WebSphere Liberty servers

Configure a connection factory in your `server.xml` file for connection to your IMS TM host system.

For the WebSphere Liberty server, the connection factory is configured in the `server.xml` file.

1. Open the `server.xml` file.
2. Configure a connection factory for connection to your IMS host system by adding a `<connectionFactory>` entry:

```
<connectionFactory jndiName="your_JNDI_name"
  type="javax.resource.cci.ConnectionFactory">
  <properties.IMS_adapter_id hostName="your_host_name"
    portNumber="xxxx" dataStoreName="data_store_name" />
</connectionFactory>
```

IMS_adapter_id is the ID that is specified in the <resourceAdapter> entry when you installed and deployed the resource adapter.

The server.xml might look as follows:

```
<server description="myLibertyServer">
  <!-- Enable features -->
  <featureManager>
    <feature>jsp-2.2</feature>
    <!-- Added jca-1.6 and jndi-1.0 features required by the IMS
         TM resource adapter -->
    <feature>jca-1.6</feature>
    <feature>jndi-1.0</feature>
  </featureManager>

  <!-- Added connection factory for the IMS TM resource adapter -->
  <connectionFactory jndiName="myJNDIName"
    type="javax.resource.cci.ConnectionFactory">
    <properties.<b>IMSTMRA</b> hostName="MY.IMS.IBM.COM"
      portNumber="9999" dataStoreName="IMS1" />
  </connectionFactory>

  <resourceAdapter id="IMSTMRA"
    location="C:\IBM\IMS\ico1410\IBM\IMS\IC014\V1410\JCA15\imsico1410.rar" />
</server>
```

3. Save your changes.

You are ready to proceed to test the connection by [deploying and testing the provided installation verification program \(IVP\)](#).

Related information

[Configuring and deploying a basic JCA resource adapter \(WebSphere Application Server Liberty Core V8.5.5\)](#)

Verifying installation by using the installation verification program

Use the IMS TM resource adapter installation verification program (IVP) to verify that an enterprise application can successfully access the target IMS system from the application server by using the IMS TM resource adapter.

The IMS TM resource adapter IVP is a simple Java EE application that is installed when the IMS TM resource adapter runtime component is installed. The objective of the IVP is to receive a message from IMS in response to the **/STA OTMA** command. The IVP does not run an IMS application program, so running an IMS application program is not required.

Tip: The IVP does not verify that IMS transactions can run successfully; it performs only the following verifications:

- Communications among the client, IMS Connect, and IMS are working.
- IMS is running, and is able to execute an IMS command and return the output of that command.
- The client is able to submit an IMS command request and receive the output from that command.

Therefore, the IVP does not require a host IMS application.

Prerequisites for running the IVP

Before you run the IVP, you must ensure that several components are running on your target host system.

- For IMS Connect:
 - Ensure that the outstanding IMS Connect reply is displayed on the system console of target system. For example:

```
HWSC0000I *IMS CONNECT READY* IMS_Connect_Name
```


For message HWSC0000I to display on the system console when IMS Connect is active, ensure either that **WTORCMD=Y** is specified in the HWS statement of the IMS Connect HWSCFGxx configuration member or that the **WTORCMD** parameter is omitted from the HWS statement.

- Verify that the target data store and port are active by entering the IMS Connect command **VIEWHWS** at the IMS Connect outstanding reply. If necessary, use the IMS Connect commands **OPENDS** *datastore_name* and **OPENPORT** *port_number*, respectively, to activate the data store and port.
- Ensure that the IMS outstanding reply is displayed on the system console of the target system. For example:

```
DFS996I *IMS READY*
```

- Verify that the XCF status of both the IMS and IMS Connect members is active by entering the IMS command **/DISPLAY OTMA** at the outstanding IMS reply:

DFS000I	GROUP/MEMBER	XCF-STATUS	USER-STATUS	SECURITY
DFS000I	GROUPNM			
DFS000I	-IMSMEM	ACTIVE	SERVER	NONE
DFS000I	-ICONNMEM	ACTIVE	ACCEPT	TRAFFIC

- Obtain the latest update of IMS TM resource adapter for fixes and enhancements.
- Ensure that you have created a connection factory.

For WebSphere Application Server, see [“Creating a connection factory in WebSphere Application Server”](#) on page 736.

After these prerequisites are met, proceed to [deploying the IVP EAR file to WebSphere Application Server](#).

Deploying the IVP EAR file in the Java EE application server

The installation verification program (IVP) for the IMS TM resource adapter is packaged as an enterprise archive (EAR) file that must be deployed to the application server to test the proper installation of the IMS TM resource adapter (a RAR file).

Prerequisites:

- A connection factory must be created to connect to the IMS host system. If you have not yet done so, [create a connection factory for use by the IVP](#).
- The application server must be started and running properly.
- For WebSphere Application Server, you must be logged in to the administrative console of the server.

Deploying the IVP EAR file to WebSphere Application Server

To deploy the IVP EAR file on WebSphere Application Server, you first create a new application by uploading the IVP EAR file, and then start the application.

In the WebSphere Application Server administrative console:

1. Install the enterprise application:
 - a) In the administrative console, expand **Applications** in the navigation tree, and click **New Application**.
 - b) In the New Application page, click **New Enterprise Application**.
 - c) Click **Local file system** or **Remote file system**, depending on where the EAR file for the IVP (imsicoivp.ear) is.
 - d) Click **Browse** and navigate to where the EAR file is.
The file `imsicoivp.ear` is in `target_dir/IBM/IMSICO/ICOxx/xx/JCAxx/`, where `target_dir` is a directory of your choice and `xx` is the corresponding version number for the IMS TM resource adapter installation.
 - e) Click **Next**.
 - f) In the Preparing for application installation page, click **Next**.

- g) Click **Fast Path**.
The Fast Path installation option displays only options that require your attention, based on the contents of your application or module.
 - h) Click **Next**.
You are prompted to specify installation options and map modules to servers.
 - i) In Step 1, specify or change any installation options as needed, and click **Next**.
 - j) In Step 2, change the server that this module is mapped to if needed, and click **Next**.
 - k) In the Summary page, verify the information that is displayed, and click **Finish**.
 - l) Click the **Save directly to the master configuration** link.
The IVP EAR file is installed.
2. Start the enterprise application.
 - a) In the administrative console, expand **Applications > Application Types > WebSphere enterprise applications**.
IMSICOIVPSERVICEEAR is displayed on the Enterprise Applications page.
 - b) Select the application IMSICOIVPSERVICEEAR by selecting the check box next to it. Click **Start** to start the application.
The status changes to a green arrow, which indicates that the application is started.

Tip: If the status does not change to started, try restarting the server.

You have deployed the EAR file for the IMS TM resource adapter IVP.

You are ready to [run the IMS TM resource adapter IVP](#).

Deploying the IVP EAR file to WebSphere Liberty servers

V13.2 and later To deploy the IVP EAR file in WebSphere Liberty servers, edit the server.xml file.

1. Open the server.xml file.
2. Add an <enterpriseApplication> entry as follows:

```
<enterpriseApplication location="file_location" type="ear">
  <classloader classProviderRef="IMS_adapter_id" />
</enterpriseApplication>
```

IMS_adapter_id should be set to the value of resource adapter ID that you set in the <resourceAdapter> tag. For example, the server.xml might look as follows:

```
<server description="myLibertyServer">
  <!-- Enable features -->
  <featureManager>
    <feature>jsp-2.2</feature>
    <!-- Added jca-1.6 and jndi-1.0 features required by the IMS
         TM resource adapter -->
    <feature>jca-1.6</feature>
    <feature>jndi-1.0</feature>
  </featureManager>

  <resourceAdapter id="IMSTMRA"
    location="C:\IBM\IMS\ico1410\IBM\IMS\IC014\V1410\JCA15\imsico1410.rar" />

  <enterpriseApplication location="imsicoivp.ear" type="ear">
    <classloader classProviderRef="IMSTMRA" />
  </enterpriseApplication>
```

You have deployed the EAR file for the IMS TM resource adapter IVP.

You are ready to [run the IMS TM resource adapter IVP](#).

Running the IMS TM resource adapter IVP

You can run the IMS TM resource adapter IVP from any web browser.

To run the IMS TM resource adapter IVP:

1. From a web browser, enter the following URL to invoke the IVP:
 - If the application server is running on your local system: `http://localhost:9080/IMSIC0IVPServiceWeb/IMSIC0IVPInputForm.html`
 - If the application server is running on a remote system: `http://remote_server:port/IMSIC0IVPServiceWeb/IMSIC0IVPInputForm.html`

remote_server:port is the hostname and port number where your application server is running.

2. Click **Submit**.

You will receive one of the following messages from IMS, depending on your security configuration:

- DFS1292E SECURITY VIOLATION
- DFS058I *hh:mm:ss* START COMMAND COMPLETED

If you do not receive one of these messages, the IVP did not run successfully. Restart the server. If restarting the server does not solve the problem, see the topic on [“Diagnosing IVP failures”](#) on page 835.

If you plan to use the IMS TM resource adapter for processing IMS callout requests, [use the IMS TM resource adapter callout IVP sample](#) to test the proper setup of your environment.

Running the IMS TM resource adapter callout IVP sample

The callout IVP sample contains a message-driven bean that is designed to process both synchronous and asynchronous callout requests from sample IMS applications. The callout IVP sample also includes sample IMS applications and associated setup files on the IMS host system to enable the sample IMS applications to issue a callout request for the message-driven bean to process.

Prerequisite: The installation verification program (IVP) provides a set of jobs and tasks that you can modify to set up the IMS OTMA destination descriptors and the transaction pipes (tpipes). The OTMA destination descriptors, one for synchronous and the other for asynchronous callout requests, describe where the callout requests are queued (tpipe name).

Two jobs are provided for you to run the provided IMS applications to issue callout requests.

Table 112. The IVP jobs, TPIPE name, OTMA destination descriptor, and COBOL application part name for the asynchronous and synchronous callout samples

Type of callout requests	Job	TPIPE	OTMA destination descriptor	Part name for the COBOL application
Asynchronous	IV_S227J	IVPPIPE3	IVPDTOR3	DFSASCBL
Synchronous	IV_S228J	IVPPIPE4	IVPDTOR4	DFSSSCBL

Two simple IMS BMP applications are provided.

- One application uses the IMS DL/I testing program, DFSDDLTO, to issue a DL/I ICAL call for sending a synchronous callout request. This ICAL call specifies the tpipe for holding the callout request, the OTMA destination descriptor for routing of the callout message, and the request and response areas in the callout message.
- The other application uses the IMS DL/I testing program, DFSDDLTO, to issue an insert (ISRT) to an alternate PCB (program communication block) call to place the asynchronous callout request on the IVPPIPE3 tpipe.

The IMS TM resource adapter callout IVP sample is a simple message-driven bean that is installed with the IMS TM Resource Adapter runtime installation. The bean can be deployed to your WebSphere

Application Server environment to verify that the bean can receive either an asynchronous or synchronous callout request from an IMS application, and, in the case of a synchronous callout request, send a response. This sample is designed to work with the IMS callout IVP samples, with the pre-defined tpipes and OTMA destination descriptors.

The general tasks involved to implement a callout solution for an IMS application are as follows:

Table 113. Steps to implementing a callout solution for an IMS application

Step	Description	Environment in which the task is performed	How the tasks are handled by the callout sample
1	Create or modify an IMS application to issue a callout request by using the DL/I ICAL call (synchronous) or ISRT altpcb call (asynchronous)	IMS host system	The IMS callout sample IVP provides an IMS application. See the related IVP jobs in the IMS installation information. Run the IVP jobs to compile and bind the samples:
2	Define an OTMA destination descriptor		
3	Restart IMS for the newly defined OTMA destination descriptor		• IMS V14 installation information
4	Create a message-driven bean or J2C application to process the IMS inbound transaction.	IBM Rational Application Developer for WebSphere Software. or other application development environments	The IMS TM resource adapter callout IVP sample provides this message-driven bean.

Table 113. Steps to implementing a callout solution for an IMS application (continued)

Step	Description	Environment in which the task is performed	How the tasks are handled by the callout sample
5	Deploy the IMS TM resource adapter runtime in the Java EE application server.	Your Java EE application server.	See: • “Installing the resource adapter on WebSphere Application Server” on page 735 • “Installing the resource adapter on WebSphere Liberty servers” on page 737
6	Configure a J2C activation specification in the Java EE application server and specify the connection information, such as the IMS host name, port number, data store name, and queue name (tpipe name) for the IMS TM resource adapter to pull callout messages.		For WebSphere Application Server, see “Configuring a J2C activation specification for the callout IVP” on page 745. For WebSphere Liberty servers, the J2C activation specification information is specified as part of the deployment topics in Step 7.
7	Deploy the message-driven bean or J2C application in the Java EE application server, and start the application.		See: • “Deploying the sample application on WebSphere Application Server to process callout requests” on page 744 • “Deploying the sample application on WebSphere Liberty servers to process callout requests” on page 746
8	Run the IMS application to issue the synchronous callout request or asynchronous callout request	IMS host system	See “Running the IMS host callout IVP application” on page 747.

Deploying the sample application on WebSphere Application Server to process callout requests

Deploy the callout IVP application on WebSphere Application Server to listen for IMS synchronous callout requests.

Prerequisite: Complete [“Configuring a J2C activation specification for the callout IVP”](#) on page 745.

1. Install the callout application:

- In the administrative console, expand **Applications** in the navigation tree, and click **New Application**.
- In the New Application page, click **New Enterprise Application**.
- Click **Local file system** or **Remote file system**, depending on where the EAR file for the IVP (imsicocalloutivp.ear) is.
- Click **Browse** and navigate to where the EAR file is.

The file `imsicocalloutivp.ear` is in `target_dir/IBM/IMSICO/ICOxx/xx/JCAxx/`, where `target_dir` is a directory of your choice and `xx` is the corresponding version number for the IMS TM resource adapter installation.

- e) Click **Next**.
 - f) In the Preparing for application installation page, click **Next**.
 - g) Click **Fast Path**.
The Fast Path installation option displays only options that require your attention, based on the contents of your application or module.
 - h) Click **Next**.
You are prompted to specify installation options and map modules to servers.
 - i) In Step 1, specify or change any installation options as needed, and click **Next**.
 - j) In Step 2, change the server that this module is mapped to if needed, and click **Next**.
 - k) In the Summary page, verify the information that is displayed, and click **Finish**.
 - l) Click the **Save directly to the master configuration** link.
The callout IVP EAR file is installed.
2. Start the application.
 - a) In the administrative console, expand **Applications > Application Types > WebSphere enterprise applications**.
IMSICOIVPSERVICEEAR is displayed on the Enterprise Applications page.
 - b) Select the application IMSICOCalloutIVPMDBEAR by selecting the check box next to it. Click **Start** to start the application.
The status changes to a green arrow, which indicates that the application is started.

The application that listens for callout requests from IMS is ready.

Proceed to [“Running the IMS host callout IVP application” on page 747](#).

Configuring a J2C activation specification for the callout IVP

In the WebSphere Application Server administrative console, set up the connection information that is already defined in the IVP sample jobs and tasks for the resource adapter to pull callout messages and to reply with responses.

To use the IMS callout sample IVP, the queue name (or tppe) where synchronous callout requests are queued must be set to `IVPPPIPE4`. This `tppe` value is already defined in IMS OTMA destination descriptor. You also need to obtain the IMS host name and port number where IMS Connect is running, and the data store name.

To create and configure a J2C activation specification in WebSphere Application Server:

1. Create a J2C specification for the IMS TM resource adapter.
 - a) In the administrative console, select **Resources > Resource Adapters > J2C activation specifications** in the navigation pane.
 - b) Click **New** in the content pane.
A Configuration tab displays.
 - c) Select the IMS TM resource adapter from the list.
The Message listener type field at the bottom of this page is automatically populated based on the selection.
 - d) Specify a name for this J2C activation specification.
For example, `IMSICOCalloutIVP`.
 - e) Type the JNDI name.
For example, `eis/IMSICOCalloutIVP`.

- f) Click **Apply**.
You are back to the J2C activation specifications page.
- g) Click **Save** in the message box at the top to save the changes to the master configuration.
The J2C activation specification is created.
2. To configure the J2C activation specification for use with the IMS TM resource adapter:
 - a) Click the name of the J2C activation specification for the IMS TM resource adapter.
The Configuration tab displays.
 - b) Under the Additional Properties section on the side, click the **J2C activation specification custom properties** link.
Custom properties in this instance of J2C activation specification display.
 - c) Specify the values for the following required properties by clicking their names in the table and entering the appropriate values.
 - queueName:
 - Specify IVPPIPE3 , IVPPIPE4 if you plan to verify both asynchronous and synchronous callout messages
 - Specify IVPPIPE3 if you plan to verify only asynchronous callout messages
 - Specify IVPPIPE4 if you plan to verify only synchronous callout messages
 - portNumber: the port number for IMS Connect
 - hostName: the IMS Connect host name
 - dataStoreName: the IMS data store name.
 - d) Specify the values for the other optional properties based on your environment.
 - e) Click **Apply**.
 - f) Click **Save** in the message box at the top to save the changes to the master configuration.
3. Restart the server to see the new activation specification.

The activation specification for the sample application is configured.

Proceed to [“Deploying the sample application on WebSphere Application Server to process callout requests”](#) on page 744.

Deploying the sample application on WebSphere Liberty servers to process callout requests

To deploy the callout IVP EAR file in WebSphere Liberty servers, download the required extended content and the message-driven bean (MDB) feature. Then modify the server.xml file to configure the MDB feature and a JCA activation specification.

Prerequisite: The WebSphere Application Server Liberty Profile extended content and the MDB feature are required for callout request processing. If they are not yet downloaded and installed:

1. Review the instructions at [WebSphere Liberty for developers](#).
2. Follow the instructions to install the Liberty Profile Extended Programming Models that provide the WebSphere Application Server Liberty Profile Extended Content.
3. Follow the instructions to install the "Message-Driven Beans 3.1" feature.

Make the following changes in the server.xml file:

1. Configure the MDB feature.

```
<!-- Enable features -->
<featureManager>

  <!-- Added the mdb-3.1 feature required for callout support -->
  <feature>mdb-3.1</feature>
</featureManager>
```

2. Configure a JCA activation specification that uses the resource adapter so that the MDB acts as a listener.

A JCA activation specification is configured by adding an `<activationSpec>` tag:

```
<activationSpec id="imsicocalloutivp/IMSICOCalloutIVPMDB/IMSICOCalloutIVPMDB">
  <properties.IMS_adapter_id hostName="host_name" portNumber="port_number"
    dataStoreName="data_store_name" queueNames="tpipe_name" />
</activationSpec>
```

IMS_adapter_id is the ID for the IMS TM resource adapter that is specified in the `<resourceAdapter>` entry when you installed and deployed the resource adapter. In the earlier example, we set the ID to *IMSTMRA*, so the corresponding activation specification would look as follows:

```
<activationSpec id="imsicocalloutivp/IMSICOCalloutIVPMDB/IMSICOCalloutIVPMDB">
  <properties.IMSTMRA hostName="my.host.server.com" portNumber="9999"
    dataStoreName="myDataStoreName" queueNames="myTpipeName" />
</activationSpec>
```

3. Add an `<enterpriseApplication>` entry for the callout IVP.

```
<enterpriseApplication location="imsicocalloutivp.ear" type="ear">
  <classloader classProviderRef="IMS_adapter_id" />
</enterpriseApplication>
```

IMS_adapter_id, again, is the ID for the IMS TM resource adapter that is specified in the `<resourceAdapter>` entry, so in this example, it would be *IMSTMRA*:

```
<enterpriseApplication location="imsicocalloutivp.ear" type="ear">
  <classloader classProviderRef="IMSTMRA" />
</enterpriseApplication>
```

4. Save your changes.

You have configured the required MDB feature and the JCA activation specification, and deployed the EAR file for the IMS TM resource adapter callout IVP.

You are ready to [run the IMS callout application on the host system](#).

Running the IMS host callout IVP application

Run the IMS host callout IVP application to issue a synchronous or asynchronous callout message.

Prerequisites: The sample application must be deployed on your target Java EE application server.

For asynchronous callout requests:

1. Run the *IV_S227J* job as described in *IV_S001T*, which contains the introductory information asynchronous IMS Callout IVP application on the host.

The sample output shows the request message as "HELLO FROM IMS":

```
CALL=ISRT
      0100
SEGMENT =(0E00HELLO FROM IMS      )

COMP                                RET CODE=OK
ALTPCB SOURCE/DEST=IVPDTOR3 RET CODE=
```

You can locate the following message in your server log or console. The presence of the message confirms that the IMS TM resource adapter Callout IVP MDB receives the asynchronous callout request from the host IMS Callout IVP application.

```
SystemOut 0 Asynchronous callout request from IMS: HELLO FROM IMS
```

For synchronous callout requests:

2. Run the *IV_S228J* job as described in the IMS Callout IVP application on the host.

The output shows the request message "HELLO FROM IMS", and the response message "HELLO FROM WEBSPHERE MDB":

```
CALL=ICAL SENDRECV IVPDTOR4 001000 00050 00050
      0003      0003      0000
AIBOALEN = 0002, AIBOAUSE = 0002, AIBRSFLD = 003Y
-----
SEGMENT =(HELLO FROM IMS )
-----
CALL=ICAL SEGMENT = (HELLO FROM WEBSPHERE MDB )
COMP RET CODE=OK
```

You can locate the following messages in the WebSphere Application Server log or console. The presence of the messages confirm that the IMS TM resource adapter callout IVP message-driven bean has received the synchronous callout request from the host IMS callout IVP application, and send a response message to the host IMS application:

```
SystemOut 0 Synchronous callout request from IMS: HELLO FROM IMS
SystemOut 0 Synchronous callout response from WAS MDB: HELLO FROM WEBSPHERE MDB
```

Installing IMS TM resource adapter service and updates

Service and updates are available as newer versions of the IMS TM resource adapter. Use the **Update RAR** function in the WebSphere Application Server administrative console to install the new version of the IMS TM resource adapter.

Prerequisite:

1. Check the IMS TM resource adapter V15 release notes.
2. Download the new version (for all supported platforms). For z/OS, you can also obtain the updates as APARs or PTFs.

To update the IMS TM resource adapter version without deleting and reconfiguring all existing connection factories, use the **Update RAR** function in the WebSphere Application Server administrative console and update the class path accordingly.

To install the new version:

1. In the WebSphere Application Server administrative console, click to expand **Resources > Resource Adapters**.
2. Click **Resource adapters**.
3. Click the check box to select the IMS TM resource adapter that you want to update and click **Update RAR**.



Figure 114. Updating the RAR for new versions of the resource adapter

4. In Step 1, select either **Local file system** or **Remote file system**, and click **Browse** to navigate to where the new IMS TM resource adapter RAR file is.
For example, on a local file system, the RAR file is located in `install_path/IBM/IMS/IC0xx/Vxxxx/JCA15/imsxxxx.rar`
5. Click **Next**.

6. In Step 2, verify the version numbers of the existing and the new resource adapters and click **Next**.
7. In Step 3, configure any new properties introduced by the new version of resource adapter as prompted. Click **Next**.
8. Confirm the changes in the summary step, and click **Finish**.
You are returned to the Resource adapter page.
9. Click the name of the resource adapter you updated.
10. In the class path field under General Properties, update the RAR file name.
For example, if you are upgrading to Version 14.1 of the IMS TM resource adapter, change the class path to: `${CONNECTOR_INSTALL_ROOT}/imsico1410.RAR`
11. Click **OK**.
You are returned to the Resource adapter page.
12. Click **Save** at the top of the page.
13. Update the IMS TM resource adapter IVP EAR file by using the same approach if the EAR file is already deployed.
14. Optionally, remove the old IMS TM resource adapter RAR files.
 - For distributed platforms, remove the installation directory at `install_path/IBM/IMS/ICOxxx/Vxxxx`.
 - For z/OS, remove the directory at `install_path/Vxxxx`.

install_path is the target installation path and *Vxxxx* is the IMS TM resource adapter version that you want to remove.

The IMS TM resource adapter RAR file is updated. All existing connection factories now use the new version of the RAR.

Related information

[WebSphere Application Server Version 8 documentation](#) For more information, see the WebSphere Application Server Version 8 documentation.

Configuring for resource workload routing

Use the resource workload routing feature in WebSphere Application Server to provide data source and connection factory failover and fallback from a predefined alternate resource.

Prerequisite: You must have created at least two connection factories for the IMS TM resource adapter. For steps on how to create a connection factory, see [“Creating a connection factory in WebSphere Application Server”](#) on page 736.

WebSphere Application Server Version 8 or later provides a resource workload routing function that enables applications to recover from resource outages without having to embed alternate resource or configuration information.

A data source and connection factory can fail over and fail back automatically when a specified or default failure threshold value is reached. When fail over occurs, the application switches from using the primary resource to using the alternate resource. Fail back occurs when the application switches back from the alternate resource to the primary resource.

When an alternate resource is configured for the IMS TM resource adapter, WebSphere Application Server sends a ping message to IMS Connect to verify that it is up and running before re-establishing a connection to the currently available resource.

To enable the resource workload routing function, configure the `alternateResourceJNDIName` property as a custom connection pool property for your primary connection factory. Set the value of the `alternateResourceJNDIName` property to the JNDI name of the secondary connection factory.

1. In the WebSphere Application Server administrative console, expand **Resources > Resource Adapters**. Click **Resource adapters**.
2. Click the **Connection pool properties** link under the Additional Properties section.

3. In the Connection pools page, click **Connection pool custom properties**.
4. In the Custom properties page, click **New** to add the alternateResourceJNDIName property.
5. The property name must be set to alternateResourceJNDIName. The value is the JNDI name of the alternate connection factory.

[J2C connection factories](#) > [myIMSTMRA](#) > [Connection pools](#) > [Custom properties](#) > [New...](#)

Use this page to specify an arbitrary name and value pair. The value that is specified for the pair is a string that can set internal system configuration properties.

Configuration

General Properties

* Name
alternateResourceJNDIName

* Value
eis/myIMSTMRA2

Figure 115. The Custom properties page when you create a property.

6. Click **OK**.

An alternate resource is configured.

You have created an alternate connection factory for your IMS TM resource adapter.

You can repeat the same step to add other custom properties such as failureThreshold. For other properties that you can configure, see the WebSphere Application Server information.

Related information

[IBM Documentation for WebSphere Application Server](#)

Chapter 49. Developing an application for use with the IMS TM resource adapter

To interact with IMS Transaction Manager (IMS TM) through the IMS TM resource adapter, you can develop your application by using a Rational or WebSphere integrated development environment (IDE), or by coding your application outside of an IDE by using the Java EE Connector Architecture Common Client Interface (CCI).

An application uses the IMS TM resource adapter to run an IMS transaction in IMS Transaction Manager. The IMS TM resource adapter also supports other types of interactions with IMS TM, such as executing the IMS commands that are supported by IMS OTMA. Consider the type of interaction with IMS TM when you develop your application. In addition, the IMS TM resource adapter implements the Java EE Connector Architecture Common Client Interface (CCI) API, an application programming interface that you can use in your applications to communicate with an Enterprise Information System (EIS), IMS TM in this case, through the IMS TM resource adapter.

Recommendation: When you develop your application by using an IDE, you can use the provided wizards and file importers to generate a Java EE application that will run on WebSphere Application Server. When you use an IDE to generate your code, you do not need to write the quality of service (QOS) code or the CCI code that is needed by your Java application. Instead, you can concentrate on writing code to implement your business logic and function.

Related information

Submitting commands to IMS

Although the IMS TM resource adapter is intended primarily for you to run transactions on a host IMS system through a Java EE application, you can also issue IMS commands that are supported by IMS OTMA from your Java applications.

Interacting with the IMS Transaction Manager

A Java application interacts with the IMS Transaction Manager (IMS TM) to run an IMS transaction, to retrieve undelivered or asynchronous output messages, to retrieve and respond to an IMS callout request, or to invoke any of the IMS commands that are supported by IMS Open Transaction Manager Access (OTMA).

Java applications that run non-conversational IMS transactions are the most common type of Java applications that use the IMS TM resource adapter. To run a non-conversational transaction, the Java application interacts with IMS TM through IMS Connect, passing the input message of the transaction to IMS TM and accepting the output message from IMS TM in return.

In the Java EE Connector Architecture (JCA), this interaction is done using an instance of the `Interaction` object. The Java application interacts with IMS by invoking the `execute` method of an `Interaction` class, providing the input message of the IMS transaction and an `IMSInteractionSpec` object. The `IMSInteractionSpec` class in the `com.ibm.connector2.ims.ico` package describes the properties of the interaction.

The properties of the `IMSInteractionSpec` class determine the type of interaction that a Java application has with IMS TM. Interactions take place in a transaction pipe (tpipe), which is a logical connection between an OTMA client, such as IMS Connect, and the server, IMS OTMA.

You can configure an `IMSInteractionSpec` object in your Java application to invoke the `execute` method of an `Interaction` object in several ways:

- By using the wizards of an integrated development environment (IDE) to generate Common Client Interface (CCI) code for the Java application. In this case, the values for the properties of the `IMSInteractionSpec` class are provided to the wizards and the CCI code is generated for you.

Recommendation: Use this method to develop your application because the skeleton code is generated for you.

- By coding a Java application that directly calls the Common Client Interface of the IMS TM resource adapter. Use the `set` methods to configure an `IMSInteractionSpec` object and invoke the `execute` method of an `Interaction` object.
- By using a variation of the code-generating IDE model in which the properties of `IMSInteractionSpec` are dynamically provided at run time. This method is typically referred to as exposing the properties of the `IMSInteractionSpec` (or `IMSConnectionSpec`) object as data.

Related reference

[IMS interaction specification properties](#)

The properties of the `IMSInteractionSpec` object describe the interaction with IMS.

Related information

[Samples and tutorials](#)

Samples and tutorials for developing Java applications or Web services to access IMS transactions can be found in several WebSphere and Rational integrated development environments (IDEs).

Programming models

The programming model to use depends on the type of IMS transactions that the Java application interacts with.

The IMS TM resource adapter supports the following programming models:

Send/receive programming model

Use the send/receive programming model to run an IMS response mode transaction.

To run a transaction in IMS, a Java application executes a `SYNC_SEND_RECEIVE` interaction. In your application, provide the following values for the `IMSInteractionSpec` object that is used by the `execute` method of the `Interaction` object.

- A value of `SYNC_SEND_RECEIVE` for the `interactionVerb` property
- A value of 0 or 1 for the `commitMode` property

However, the `SYNC_SEND_RECEIVE` interaction processing is different for shareable and dedicated persistent socket connections. Depending on the type of socket connections, the processing model is different in either a normal processing of the transaction, or when an error or an execution timeout occurs.

If you convert a send-then-commit (CM1) application that expects a response to a commit-then-send (CM0) application that does not, set the `IMSInteractionSpec` `CM0Response` property to `true`. When this property is set for a CM0 transaction, if the IMS application does not reply to the IOPCB or complete a message switch to another transaction, IMS OTMA issues a DFS2082 message to the client, regardless of the transaction response mode.

Shareable persistent socket processing model

Shareable persistent socket connections are connections that can be used for commit mode 1 and commit mode 0 interactions. The following scenarios describe the `SYNC_SEND_RECEIVE` interaction on a shareable persistent socket during normal processing, error processing, and execution timeout.

Normal processing scenario

The IMS TM resource adapter, with the application server, obtains either an available connection from the connection pool or creates a new connection. The IMS TM resource adapter, as part of initializing a new connection, generates a client ID for the connection. The generated client ID identifies the socket connection, and for commit mode 0 interactions, the tpipe and associated OTMA asynchronous hold queue.

The IMS TM resource adapter ensures that a socket is associated with the connection and sends the request with input data to IMS Connect by using that socket. IMS Connect then sends the message to IMS, where IMS runs the transaction and returns the output message.

For commit mode 0 interactions, when it receives the output message, the IMS TM resource adapter sends an ACK message to IMS, which signals IMS to discard the output from the IMS queue. When the client application closes the connection or terminates, the connection is returned to the connection pool for reuse by other commit mode 0 or commit mode 1 interactions.

Error processing scenario

All errors result in a resource exception being thrown to the client application. In addition, some errors result in the socket being disconnected by IMS Connect. For commit mode 0 interactions, an exception means that the output message cannot be delivered to the client application. However, following the exceptions, undelivered output messages for commit mode 0 interactions on shareable persistent socket connections can be retrieved if the SYNC_SEND_RECEIVE interaction specifies that undelivered output is to be rerouted to a specific destination. To have an undelivered output message rerouted to a specific destination, the following additional properties must be specified in the IMSInteractionSpec object that is passed on the SYNC_SEND_RECEIVE interaction:

- Set the `purgeAsyncOutput` property to false so that undelivered output is not purged.
- Set the `reRoute` property to true, and specify a reroute destination in the `RouteName` property.

To retrieve undelivered output from a reroute destination, a separate client application issues a SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interaction on a dedicated persistent socket connection. The client application provides the reroute destination as the client ID of the interaction.

Alternatively, to retrieve undelivered output from a reroute destination, a separate client application issues a SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT, or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interaction on a shareable persistent socket connection by specifying the alternate client ID. Using the alternate client ID, a client application can retrieve undelivered asynchronous output messages from any tpipe.

The default value of the `purgeAsyncOutput` property is true. When the value of the `purgeAsyncOutput` property is true, the following output messages are purged:

- Undelivered output message inserted to the I/O Program Communications Block (I/O PCB) by the primary IMS application program
- Output messages inserted to the I/O PCB by secondary IMS application programs invoked by program-to-program switches

When the `purgeAsyncOutput` property is set to false, the reroute destination must be specified.

Execution timeout scenario

If an execution timeout occurs, the socket connection remains open but the output message is not delivered to the client application. However, following an execution timeout exception, undelivered output messages for commit mode 0 interactions on shareable persistent socket connections can be retrieved in either of the following two ways:

- The same client application that issued the SYNC_SEND_RECEIVE interaction can issue a SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interaction.
- The undelivered output message can be rerouted to a specific destination as described in the error processing scenario.

When the client application closes the connection or terminates, the connection is returned to the connection pool so it can be reused by other commit mode 0 or commit mode 1 interactions.

Related concepts

Shareable persistent sockets

A shareable persistent socket is a socket that can be shared (serially reused) by multiple applications that execute either commit mode 1 (CM1) or commit mode 0 (CM0) interactions.

Dedicated persistent socket processing model

Dedicated persistent socket connections are connections that can be used only for commit mode 0 interactions.

The following scenarios describe the commit mode 0 SYNC_SEND_RECEIVE interaction on a dedicated persistent socket during normal processing, error processing, and execution timeout.

Normal processing scenario

When a commit mode 0 SYNC_SEND_RECEIVE interaction is executed by a client application, the application server returns an existing connection with the user-specified client ID, or creates a new connection with the user-specified client ID. The user-specified client ID identifies the socket connection and the tpipes that are associated with the OTMA asynchronous hold queue.

The IMS TM resource adapter ensures that a socket is associated with the connection and sends the request with input data to IMS Connect using that socket. IMS Connect then sends the message to IMS, where IMS runs the transaction and returns the output message. When it receives the output message, the IMS TM resource adapter sends an ACK to IMS, which signals IMS to discard the output from the IMS queue. When the connection is closed or the application terminates, the connection is returned to the connection pool for reuse by another application that is running a commit mode 0 interaction with the same user-specified client ID.

Error processing scenario

All errors result in a resource exception being thrown to the client application. In addition, some errors result in the socket being disconnected by IMS Connect. For commit mode 0 interactions, an exception means the output message cannot be delivered to the client application. The undelivered output is queued to the tpipes that are associated with the user-specified client ID.

The properties, `purgeAsyncOutput` and `reRoute`, are not applicable to dedicated persistent sockets. You can not purge or reroute undelivered output messages on a dedicated persistent socket.

Execution timeout scenario

If an execution timeout occurs, the socket remains open and the output of the commit mode 0 interaction is queued to the tpipes associated with the user-specified client ID for later retrieval. When the connection is closed or the application terminates, the `IMSManagedConnection` object is returned to the connection pool for reuse by another application that is running a commit mode 0 interaction with the same user-specified client ID.

Related concepts

Dedicated persistent sockets

A dedicated persistent socket is a socket connection that is assigned to a specific client ID and remains dedicated to that particular client ID until it is disconnected.

Send-only programming model

Use the send-only programming model to run an IMS non-response mode transaction in the IMS Transaction Manager (IMS TM).

To run a non-response mode transaction in IMS TM, a Java application executes a SYNC_SEND interaction. In a SYNC_SEND interaction, the IMS TM resource adapter sends the request to IMS through IMS Connect and does not expect a response from IMS. Because the IMS TM resource adapter performs a send-only interaction with IMS, a SYNC_SEND interaction is typically used with a non-response mode transaction.

To use a SYNC_SEND interaction to run a transaction, your application must provide a value of SYNC_SEND for the `interactionVerb` property and a value of 0 for the `commitMode` property of the `IMSInteractionSpec` object that is used by the `execute` method. The SYNC_SEND interaction processing varies, depending on the type of persistent socket that is used (shareable or dedicated) and the type of IMS transaction that is run.

Important: The `IMSInteractionSpec` properties `purgeAsyncOutput`, `reRoute`, and `reRouteName` do not apply to SYNC_SEND interactions and are ignored by the IMS TM resource adapter.

Shareable persistent socket processing model

The following scenarios describe a SYNC_SEND interaction on a shareable persistent socket connection for different types of transactions.

- Non-response mode transaction

The IMS application program associated with a transaction that is defined to IMS as non-response mode typically does not insert an output message to the I/O PCB. Therefore, an output message is not created and nothing is queued on a tpipe.

- Response mode transaction

The IMS application program associated with a transaction that is defined to IMS as a response mode transaction typically inserts an output message to the I/O PCB. Because the IMS TM resource adapter does not expect a response from a SYNC_SEND interaction, the output message, if inserted, is queued on the tpipe with the name of the generated client ID. However, interactions SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT can be used to retrieve the response, if they are performed following the SYNC_SEND interaction in the same application and on the same connection.

- Non-response mode or response mode transactions that invoke an IMS application program that inserts a message to an alternate PCB

A message that is inserted to an alternate PCB can be retrieved by executing an interaction on a dedicated persistent socket connection. To insert this message, see the discussion on retrieving asynchronous output on dedicated persistent socket connections in the "Asynchronous output programming model" topic.

Dedicated persistent socket processing model

The following scenarios describe a SYNC_SEND interaction on a dedicated persistent socket connection for different types of transactions. SYNC_SEND interactions use commit mode 0 and dedicated persistent socket connections can be used only for commit mode 0 interactions.

- Non-response mode transaction

The IMS application program associated with a transaction that is defined to IMS as non-response mode typically does not insert an output message to the I/O PCB. Therefore, an output message is not created and nothing is queued on a tpipe.

- Response mode transaction

The IMS application program associated with a transaction that is defined to IMS as response mode typically inserts an output message to the I/O PCB. Because the IMS TM resource adapter does not expect a response from a SYNC_SEND interaction, the output message, if inserted, is queued on the tpipe with the name that is provided for the client ID of the interaction. Messages queued to this type of tpipe can be retrieved by issuing a SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions. The tpipe name is the client ID that is specified for the SYNC_SEND interaction. A client ID is required for interactions that use a dedicated persistent socket connection.

- Non-response mode or response mode transactions that invoke an IMS application that inserts to an alternate PCB

A message inserted to an alternate PCB can be retrieved by executing an interaction on a dedicated persistent socket connection. For more information, see the "Asynchronous output programming model" topic.

Related concepts

[Asynchronous output programming model](#)

Use this programming model to retrieve output that has been queued by IMS.

Asynchronous output programming model

Use this programming model to retrieve output that has been queued by IMS.

Your application might need to retrieve output that was queued at some earlier time. Such output is called *asynchronous output*. Asynchronous output can result from several situations: undelivered output from commit mode 0 interactions, output from program-to-program switches, output from ISRT calls to an alternate PCB, and output from an IMS application that returns multiple output messages.

To retrieve asynchronous output, a Java application executes a SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interaction. Your application must provide the value for the interactionVerb property and a value of 0 for the commitMode property of the IMSInteractionSpec object that is used by the execute method of the interaction.

Two types of socket connections that can be used to retrieve asynchronous output are shareable persistent socket and dedicated persistent socket. The way to retrieve asynchronous output messages varies, depending on the type of socket connection that is used. The interactionVerb property values that can be used to retrieve asynchronous output are:

- SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT
- SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT

Tip: SYNC_RECEIVE_ASYNCOUTPUT was deprecated in IMS TM Resource Adapter Version 10, and is replaced by SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT.

SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT is different from SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT in how IMS Connect checks for the output on the IMS OTMA asynchronous hold queue.

- For SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT interactions, if no asynchronous output is in the IMS OTMA asynchronous hold queue when the retrieve request is made, IMS Connect returns an execution timeout notification as soon as the execution timeout value that is specified by the client application has passed.

Recommendation: Specify the minimum execution timeout value of 10 for SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT interactions.

- For SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions, when the execution timeout value has passed, if no asynchronous output is in the hold queue, IMS Connect returns an execution timeout error. Otherwise, IMS Connect returns the first output message when it is placed in the queue.

Recommendation: For SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions, select an appropriate execution timeout value, rather than the minimum value.

The asynchronous output programming model requires commit mode 0 and can be used on both shareable persistent socket and dedicated persistent socket connections. In addition, the IMSInteractionSpec properties purgeAsyncOutput, reRoute, and reRouteName do not apply and are ignored when the interactionVerb property is set to either of these values. The interactionVerb property is invoked differently on dedicated and shareable persistent socket connections.

Related concepts

[Java exceptions that involve output messages](#)

If any errors occur when the IMS TM resource adapter passes the message through IMS Connect to IMS, and as IMS processes this transaction and attempts to return the output, the Java client receives an exception.

Securing message retrieval from IMS hold queues

You can specify the user ID and password information to ensure that only authorized users are allowed to retrieve asynchronous output or callout request messages from the IMS hold queue.

Retrieving asynchronous output on dedicated persistent socket connections

To retrieve asynchronous output on a dedicated persistent socket connection, use commit mode 0 interaction and specify a client ID.

1. Execute a commit mode 0 interaction with the `interactionVerb` property of `IMSInteractionSpec` set to `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT`, or `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT`.
2. Provide a value for the `clientId` property. The `clientId` property is required because it determines the tpipe from which the asynchronous output is retrieved.
 - To retrieve output messages from a commit mode 0 interaction on a dedicated persistent socket, the client ID that is specified on the `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT` or `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT` interaction must match the value that is specified for the original commit mode 0 interaction.
 - To retrieve output messages sent to an alternate PCB, the client ID that is specified on the `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT` or `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT` interaction must match the name of the alternate PCB.
 - To retrieve output messages that were rerouted to a `reRouteName` destination, the client ID that is specified on the `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT` or `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT` interaction must be set to the destination that is specified by the `reRouteName` property.

Related concepts

Dedicated persistent sockets

A dedicated persistent socket is a socket connection that is assigned to a specific client ID and remains dedicated to that particular client ID until it is disconnected.

Retrieving asynchronous output with an alternate client ID on shareable persistent socket connections

A client application can retrieve asynchronous output messages that are either rerouted or non-rerouted from any OTMA tpipe by using an alternate client ID.

This alternate client ID is specified by the `altClientId` property of the `IMSInteractionSpec` object. The alternate client ID corresponds with the name of the OTMA tpipe that has the asynchronous output messages. A client application cannot specify a reroute name (the `reRouteName` property) and an alternate client ID (the `altClientId` property) in the same interaction.

Alternate client IDs are supported when a client application meets all the following requirements:

1. Use a TCP/IP connection with a shareable persistent socket.
2. Use a valid value for the `interactionVerb` property for retrieval of asynchronous output messages (resume tpipe): `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT` and `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT`.
3. Set the `commitMode` property value to 0.

Related concepts

Shareable persistent sockets

A shareable persistent socket is a socket that can be shared (serially reused) by multiple applications that execute either commit mode 1 (CM1) or commit mode 0 (CM0) interactions.

Retrieving asynchronous output without an alternate client ID on shareable persistent socket connections

Asynchronous output must be retrieved on the same connection as the previous interaction if an alternate client ID is not specified.

Without using alternate client ID, a client application can retrieve only non-rerouted asynchronous output messages from an interaction on a shareable persistent socket connection by executing a SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interaction. This interaction must be on the same shareable persistent socket connection where the interaction from the same application led to the asynchronous output being queued.

The following table shows the behavior of the retrieval of asynchronous output message (resume tpipe) interaction when different property values are specified.

Table 114. Behavior of resume tpipe interactions with different IMSConnectionSpec property values

Interaction	Socket	Client ID	Reroute name	Alternate client ID	Comments
resume tpipe, without specifying alternate client ID	Shareable	null	null	null	Asynchronous output must be retrieved on the same connection as the previous interaction. Otherwise, the asynchronous output cannot be retrieved. If the retrieval failed, the message is lost.
resume tpipe, without specifying alternate client ID	Shareable	null	myRR	null	Asynchronous messages must be retrieved on the same connection as the previous interaction. If the retrieval failed, the message is rerouted to myRR.
resume tpipe, specifying alternate client ID	Shareable	null	N/A	myTpipe	Asynchronous messages are retrieved from tpipe myTpipe. the reRouteName and altClientID properties are mutually exclusive and cannot be used at the same time.
resume tpipe	Dedicated	myCID	N/A	N/A	Asynchronous messages must be retrieved by using the client ID myCID.

For shareable persistent socket connections, the commit mode for a SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT interaction is automatically set to 0 by the IMS TM resource adapter, regardless of the value that is set in the IMSInteractionSpec object for that interaction. This behavior is different from the behavior for SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT interactions on dedicated persistent socket connections. For dedicated socket connections, the commit mode must be explicitly set to 0.

Related concepts

[Shareable persistent sockets](#)

A shareable persistent socket is a socket that can be shared (serially reused) by multiple applications that execute either commit mode 1 (CM1) or commit mode 0 (CM0) interactions.

Displaying and interpreting output message counts

You can use IMS Connect commands to display output message counts when you want to monitor or troubleshoot the proper retrieval of asynchronous output messages.

For commit mode 0 interactions, the tpipe name is the client ID that is used for the interaction. For commit mode 0 interactions, the IMS OTMA asynchronous hold queue that is associated with the tpipe has the same name as the client ID.

For commit mode 1 interactions, the tpipe name is the IMS Connect port number that is used for the interaction. Each port has a tpipe that is used for all clients that run commit mode 1 interactions on that port.

1. To view counts of the output messages that are sent to the IMS TM resource adapter, as well as messages that are inserted to Alternate Program Communication Blocks (Alternate PCBs), use the IMS Connect command **/DISPLAY TMEBER IMSConnect_Name TPIPE ALL**.

The following sample output is from a **/DISPLAY TMEBER HWS1 TPIPE ALL** command. The types of tpipes and counts for the command output are also described.

DFS000I	MEMBER/TPIPE	ENQCT	DEQCT	QCT	STATUS	IMS1
DFS000I	HWS1					IMS1
DFS000I	-9999	0	0	0		IMS1
DFS000I	-HWSMIJRC	2	2	0		IMS1
DFS000I	-CLIENT01	3	2	1		IMS1
DFS000I	-ALTPCB1	2	1	1		IMS1
DFS000I	-HWS\$DEF	1	0	1		IMS1
DFS000I	-RRNAME	1	0	1		IMS1

2. To interpret the command output, determine which tpipe is the queue of interest, and the corresponding QCT column contain the message count.

The tpipe name is determined by the type of interactions and connections.

- For **commit mode 1 interactions on a shareable persistent socket**:
 - The tpipe name is the port number used for the interaction. In this sample, tpipe 9999 indicates that this queue is for a commit mode 1 interaction on a shareable persistent socket.
 - The enqueue count (ENQCT) and dequeue count (DEQCT) are the same, and the queue count (QCT) is 0, because undelivered output messages are not recoverable for commit mode 1 transactions.
- For **commit mode 0 interactions on a shareable persistent socket**:
 - The tpipe name is generated by the IMS TM resource adapter and has a prefix of HWS. In this example, tpipe name HWSMIJRC is generated by the IMS TM resource adapter.
 - The enqueue count (ENQCT) and dequeue count (DEQCT) are the same, and the queue count (QCT) is 0 if all messages are delivered to the IMS TM resource adapter.
 - The enqueue count (ENQCT) and dequeue count (DEQCT) are the same, the queue count (QCT) is 0, and all undelivered output messages are discarded if both the following conditions are met:
 - Output messages are not delivered to the IMS TM resource adapter on SYNC_SEND_RECEIVE interactions.
 - The default values for the reRoute property (false) and the purgeAsyncOutput property (true) are used.
 - The enqueue count (ENQCT) is greater than the dequeue count (DEQCT), and the queue count (QCT) is the number of messages that were not delivered to the IMS TM resource adapter if both the following conditions are met:
 - Output messages are not delivered to the IMS TM resource adapter on SYNC_SEND_RECEIVE interactions.
 - The reRoute property is set to true and the purgeAsyncOutput property is set to false

The tpipe name is the value specified for the reRouteName property. For example, RRNAME, or a default value, for example, HWS\$DEF.

- For SYNC_SEND interactions, output is not expected, and undelivered output does not apply. If SYNC_RECEIVE_ASYNCPUT_SINGLE_NOWAIT and SYNC_RECEIVE_ASYNCPUT_SINGLE_WAIT interactions are unsuccessful, the queue count does not change.
- For **commit mode 0 interactions on a dedicated persistent socket:**
 - Typically, the tpipe name is provided by the Java application and does not include a prefix of HWS. For example, CLIENT01. However, you might occasionally see a tpipe name of HWS\$DEF, which is the default value for the reRouteName property.
 - The enqueue count (ENQCT) and dequeue count (DEQCT) are the same, and the queue count (QCT) is 0 if all messages are delivered to the IMS TM resource adapter, and no undelivered messages were rerouted from interactions on shareable persistent socket connections.
 - If output messages are not delivered to the IMS TM resource adapter, or rerouted from interactions on shareable persistent socket connections, the enqueue count (ENQCT) is greater than the dequeue count (DEQCT), and the queue count (QCT) is the number of messages that were not delivered. The tpipe name is the client ID name that the user specified. For example, CLIENT01.
- For **output messages inserted to Alternate PCBs:**
 - The tpipe name is the name of the Alternate PCB, for example, ALTPCB1.

Callout programming models

Use this programming model to send outbound messages from an IMS application to request services or data from message-driven beans (MDBs), Enterprise JavaBeans (EJB) components, or web services and, optionally, receive response data.

In the IMS callout programming models, IMS acts as a client to interact with business data or logic that resides in a MDB, an EJB component, or a Web service. IMS applications invoke external Java applications by issuing callout requests. From the perspective of WebSphere Application Server, on the contrary, these messages from IMS applications are inbound messages.

Depending on whether a response from the external Java application is expected, and whether the response is expected by the initiating IMS application in the same transaction, different programming models can be employed.

Synchronous callout

A synchronous callout request is a request that expects a response from the external Java application or Web service to return in the same transaction to the same IMS application.

When an IMS application issues a synchronous callout request by using the IMS DL/I ICAL call, an OTMA transaction pipe (tpipe) is used to hold the request. The IMS application in the dependent region remains scheduled and waits for the response. After the external application pulls the request message from the tpipe, the tpipe is in a wait status until a response is returned from the external application.

For message-driven beans (MDBs), you can take advantage of the Java EE Connector Architecture 1.5 inbound specification, and the tooling support in various Rational or WebSphere integrated development environments (IDEs). The IMS TM resource adapter listens for callout messages from IMS Connect and handles the correlation of the request and its response.

For non-MDB applications, it is the application's responsibility to poll IMS Connect for any callout requests on the hold queue. The application also needs to handle the correlator token that is passed as a property of the IMSInteractionSpec class.

The key benefit of having the callout request and the response in the same transaction instance is that the programming logic for issuing the request and for processing the response can be contained in the same

IMS application. This approach does, however, require that the IMS application wait for the response, and the dependent regions would be blocked.

Asynchronous callout

An asynchronous callout request is a request that does not expect a response, or expects the response to return in a different transaction.

After the IMS application issues the asynchronous callout request by using the ISRT ALPCB call, an OTMA transaction pipe (tpipe) or an alternative destination (as coded in an exit routine), is used to hold the request. After the external application pulls the request message from the hold queue, the IMS application is no longer scheduled in the IMS dependent region, and the dependent region is freed up.

If a response is expected, the response message returns as a new IMS transaction request to either the same or a different IMS application with the output data, depending on your application design.

The key benefit of having the response return asynchronously is to avoid blocking the dependent regions and therefore locking IMS resources for an excessive period of time. However, this approach does require that, if a response needs to be processed, the IMS application must be designed to be capable of handling the output in a separate transaction instance. It is also the responsibility of the application to correlate that response with the request.

Related concepts

Execution timeouts

The value for the execution timeout is the maximum amount of time that is allowed for IMS Connect to send a message to IMS and receive a response to that message back from IMS.

Securing message retrieval from IMS hold queues

You can specify the user ID and password information to ensure that only authorized users are allowed to retrieve asynchronous output or callout request messages from the IMS hold queue.

Related reference

Diagnosing problems with callout requests

Callout request errors are often related to the execution timeout setting, incorrect or corrupted callout messages, or network failures.

OTMA destination descriptors

You can instruct an IMS application to route callout requests to services accessible to the IMS TM resource adapter without having to code assembler routing exits by using the OTMA destination descriptors.

You can define an IMS Connect destination (a transaction pipe, or tpipe) to where the IMS callout request messages are routed by using the OTMA destination descriptor. If no special routing is required, the tpipe name for the callout request can be specified directly in the ICAL call (for synchronous callout) or ISRT ALTPCB call (for asynchronous callout) in the IMS application.

To use the OTMA destination descriptor, your IMS system programmer must configure the descriptor in the DFSYDTx PROCLIB member. The following example demonstrates a descriptor for a destination named IMSTMRA that will route messages to the target member named HWS1 and a tpipe named HWS1TP01.

```
D IMSTMRA TYPE=IMSCON T MEMBER=HWS1 TPIPE=HWS1TP01
```

An OTMA destination descriptor is required for routing synchronous callout requests, but not required for routing asynchronous callout requests. However, the use of the OTMA destination descriptor for asynchronous callout requests is preferred because the descriptor approach is easier to use than the OTMA exit routines approach.

For more information, see the topic "OTMA descriptors" in *IMS Communications and Connections* information, and the DFSYDTx PROCLIB member in *IMS System Definition* information.

Security for retrieving callout requests

With the security feature that is provided by the IMS OTMA resume tpipe, you can protect callout messages from unauthorized use, with either RACF, the OTMA resume tpipe security exit routine (DFSYRTUX), or both.

IMS callout requests are retrieved from IMS Connect by using the RESUME TPIPE call. When security is enabled, 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 messages are sent to an OTMA client. If security is enabled and the tpipe does not exist at the time the RESUME TPIPE call is issued, the call is rejected.

For message-driven beans (MDBs), Secure Sockets Layer (SSL) authentication is supported for communication with IMS. Security information is specified in the J2C activation specification (IMSActivationSpec) that is configured in WebSphere Application Server.

For non-MDB applications, you can optionally specify a user ID in your Java application if you have configured your IMS security so that only authorized users can retrieve the callout request message from the hold queue. This user ID must be specified in the connection specification in your application or the connection factory that is used by your application.

Synchronous callout message flow

A synchronous callout message is placed on the tpipe hold queue and waits to be pulled by an external application or service, and the tpipe remains in a wait state until the response message from the external application returns before the next callout message can be delivered.

A typical message flow is as follows:




Figure 116. Synchronous callout message flow

1. A Java EE application in WebSphere Application Server starts and WebSphere Application Server connects to IMS Connect through the IMS TM resource adapter. The IMS TM resource adapter, in turn, issues a RESUME TPIPE request to the tpipe and waits for the callout request from IMS Connect.
2. An initiating client, such as a terminal or an IMS Connect or OTMA client starts an IMS application.

3. The IMS application issues an IMS DL/I ICAL call and specifies the OTMA destination descriptor that contains the destination tpipe name where the callout request message is to be queued. When the request is placed on the tpipe hold queue, a correlation token is attached with the request.
 4. If a callout request is not available at the time of the request, the IMS TM resource adapter is blocked and the Java application will wait for the next available callout message. In the case of an EJB application where the application itself handles the polling of IMS Connect for the callout messages, the IMS TM resource adapter waits until timeout occurs. As soon as the callout request is available in the tpipe, IMS Connect delivers the callout message to the IMS TM resource adapter.
 5. The IMS TM resource adapter receives the callout request message and sends the callout request to the Java application. The application processes the callout request.
 6. The Java application sends the response, with the correlation token attached, to the IMS TM resource adapter.
 7. The IMS TM resource adapter forwards the response to IMS Connect.
- If the send-only-with-acknowledgment protocol is enabled, IMS TM Resource Adapter receives an acknowledgment from IMS when IMS receives the synchronous callout response message.
8. Based on the correlation token, the response is correlated back to the initiating request from IMS.

Asynchronous callout message flow

An asynchronous callout message is placed in the hold queue and awaits to be pulled by an external application or service. If response data is expected, the external application issues a normal IMS transaction to the appropriate IMS application with the output data.

The following figure shows the typical asynchronous callout message flow:




Figure 117. Asynchronous callout message flow

1. A Java application in WebSphere Application Server starts and obtains a sharable persistent connection to IMS Connect through the IMS TM resource adapter. The application issues a SYNC_RECEIVE_CALLOUT interaction, specifies the tpipe name as the value for the alternate client

ID, and sets a timeout value. The IMS TM resource adapter, in turn, issues a RESUME TPIPE request to the tpipe and waits for the callout request from IMS Connect.

2. An initiating client, such as a terminal or an IMS Connect or OTMA client, starts an IMS application.
3. The IMS application issues an ISRT ALTPCB call to an OTMA destination descriptor, which contains the destination tpipe name. The callout request message is queued in this tpipe.
4. If a callout request is not available at the time of the SYNC_RECEIVE_CALLOUT request, the IMS TM resource adapter is blocked, and the bean waits for the next available callout message or until timeout occurs. When the callout request is available in the tpipe, IMS Connect delivers the callout message to the IMS TM resource adapter.
5. The IMS TM resource adapter receives the callout request message and returns the callout request to the bean. The bean processes the callout request.
6. If the bean receives response data to be returned to IMS, the bean issues a normal IMS transaction request to the appropriate IMS application with the output data.

Managed callout programming model

You can use the *managed callout programming model* to send outbound messages from IMS applications to request services or data from an external message-driven bean (MDB) and receive response data synchronously or asynchronously to the same IMS application in the same transaction.

MDBs are stateless, server-side Java EE components that can be configured as listeners on a Java Connector Architecture (JCA) 1.5 resource adapter for handling inbound messages coming from an enterprise information system (EIS). An MDB is an EJB component that can be activated by message delivery to consume and process messages.

In the JCA 1.5 specification, MDBs are commonly referred to as *message endpoints* or simply *endpoints*. The key benefits of MDBs include a clearer separation of message handling and business processing, and a wider reuse of the business processing by other incoming messages.

Because this programming model takes advantage of the JCA 1.5 standards, the IMS TM resource adapter can manage the callout protocol and the correlation of the response to the appropriate request for you. As a result, the managed callout programming model is preferred over the non-managed callout programming model for its ease of use, cleaner code, greater reuse, and better scalability.

The following diagram shows how a synchronous callout message is processed by an MDB.




Figure 118. Processing a synchronous callout message from IMS by using a message-driven bean

The following diagram shows how an asynchronous callout message is processed by an MDB.




Figure 119. Processing an asynchronous callout message from IMS by using a message-driven bean

In the MDB callout programming model, the MDB must implement:

- The `com.ibm.j2ca.base.ExtendedInboundListener` interface
- The `onMessage()` method for receiving synchronous callout requests
- The `onNotification()` method for receiving asynchronous callout request messages

These methods are invoked by the IMS TM resource adapter when a message comes in for the particular MDB, depending on whether the callout message is asynchronous or synchronous. The IMS TM resource adapter automatically handles the callout protocol. For synchronous callout messages, the IMS TM resource adapter also handles the passing of the correlation token.

You can enable the send-only-with-acknowledgment protocol. Using this protocol, IMS TM Resource Adapter receives an acknowledgment from IMS when IMS receives a synchronous callout response message. In Version 13 and earlier, this acknowledgment is unavailable.

You can enable connection pooling for synchronous callout messages that are sent back to IMS. In Version 13 and earlier, the IMS TM Resource Adapter must create and then discard a new connection whenever a callout response is sent. By enabling connection pooling, you can avoid the overhead of opening and closing multiple connections.

Retrieving IMS callout messages by using a message-driven bean

To retrieve IMS callout messages by using an MDB:

1. Generate the data binding and an MDB by using IBM Rational Application Developer for WebSphere Software or other integrated development environments (IDEs) that include the J2C tools.
 - a) Import the IMS application source file into the IDE.
 - b) Use the J2C wizard to generate data bindings for handling the callout requests and messages.
2. Implement the `com.ibm.j2ca.ExtendedInboundListener` interface to listen to callout messages.
 - Implement the `onMessage()` method for synchronous callout requests.
 - Implement the `onNotification()` method for asynchronous callout request messages
3. Configure a J2C activation specification (an instance of the `IMSActivationSpec` class) in WebSphere Application Server, and specify, at a minimum, a value for the following properties to listen for callout requests:
 - Host name
 - Port number
 - Data store name
 - Queue names (a comma-delimited list of tpipes)
4. Deploy the MDB on WebSphere Application Server by using the WebSphere Application Server administrative console.
 - a) At the **Bind listeners for message-driven beans** panel in the administrative console, enter the JNDI name of the J2C activation specification that you configured previously.
 - b) Start the MDB application.

Configuring a J2C activation specification

You must configure a J2C activation specification to describe the inbound messaging from IMS for message-driven beans.

J2C activation specifications are part of the configuration of inbound messaging support of a JCA 1.5 resource adapter. A JCA 1.5 resource adapter supports inbound messaging by implementing its own message listener interface. The message listener is the interface that the resource adapter uses to communicate inbound messages to the message endpoint. A message-driven bean is a message endpoint and implements one of the message listener interfaces provided by the resource adapter.

When an application that contains a message-driven bean is deployed, the deployer must select a resource adapter that supports the same type of message listener that the message-driven bean implements. As part of the message-driven bean deployment, the deployer needs to specify the properties to set on the J2C activation specification. Later, during application startup, a J2C activation specification instance is created, and these properties are set and used to activate the endpoint.

The `IMSActivationSpec` class is the implementation of J2C activation specifications in IMS TM resource adapter. When you deploy a message-driven bean in WebSphere Application Server, you must select the IMS TM resource adapter, and configure the activation specification by specifying the values for `IMSActivationSpec` properties that describe the inbound communication from IMS.

If multiple data stores and multiple tpipes are involved, the IMS TM resource adapter thread pool might quickly exhaust and the callout processing might stall. In this case, you might want to configure the WebSphere Application Server to add your own thread pool. After the thread pool is added, reference the thread pool in the **Thread pool alias** field for the IMS TM resource adapter.

Related tasks

[“Configuring a thread pool in WebSphere Application Server” on page 772](#)

Configure a thread pool in WebSphere Application Server if you have multiple data stores and multiple tpipes for synchronous callout request processing.

Configuring a J2C activation specification in WebSphere Application Server

You must configure a new J2C activation specification from the IMS TM resource adapter and then specify the values for communication with IMS.

To create and configure a J2C activation specification in WebSphere Application Server:

1. Create a J2C specification for the IMS TM resource adapter.
 - a) Create a J2C specification for the IMS TM resource adapter. In the administrative console (also known as the Integrated Solutions Console) of WebSphere Application Server, select **Resources** > **Resource Adapters** > **J2C activation specifications** in the navigation pane.




Figure 120. J2C activation specification configuration in WebSphere Application Server

- b) Click **New** in the content pane.

A Configuration tab is displayed.
 - c) Under Provider, select the IMS TM resource adapter from the list.

The Message listener type field at the bottom of this page is automatically populated based on the selection.
 - d) Specify a name for this J2C activation specification.

For example, IMSActivationSpec.
 - e) Type the JNDI name.

For example, eis/IMSActivationSpec.
 - f) Click **Apply**.

You are back to the J2C activation specifications page.
 - g) Click **Save** in the message box at the top to save the changes to the master configuration.

The J2C activation specification is created.

New Delete					
Select	Name	JNDI name	Scope	Provider	Description
☐	IMSActivationSpec	eis/IMSActivationSpec	Node=svl-ylaiNode01	IMS TM Resource Adapter	
Total 1					

Figure 121. Newly created J2C activation specification from the IMS TM resource adapter

2. To configure the J2C activation specification for use with the IMS TM resource adapter:

- Click the name of the J2C activation specification for the IMS TM resource adapter. The Configuration tab is displayed.
- Under the Additional Properties section on the side, click the **J2C activation specification custom properties** link.

Custom properties in this instance of J2C activation specification are displayed.

Note: The properties in the following figure might not match the properties that you see, depending on your version of the IMS TM resource adapter.

[J2C activation specifications](#) > [IMSActivationSpec](#) > [Custom properties](#)

Use this page to specify custom properties that your enterprise information system (EIS) requires from providers and resource factories that you configure. For example, most database vendors require a properties for data sources that access the database.

Preferences

Name	Value	Description	Required
queueNames	SYNCTP01, ASYNCTP1		true
portNumber	9999		true
hostName	csdmec22.svl.ibm.com		true
dataStoreName	IMS1		true
SSLKeyStoreName			false
retryInterval	60000		false
SSLKeyStorePassword			false
filterFutureEvents	false		false
pollPeriod	2000		false
pollQuantity	10		false
BONamespace			false
deliveryType	ORDERED		false
SSLEnabled	false		false
assuredOnceDelivery	true		false
stopPollingOnError	false		false
eventTypeFilter			false
userName			false
retryLimit	0		false
SSLTrustStorePassword			false
SSLTrustStoreName			false

Figure 122. Custom properties for the J2C activation specification for the IMS TM resource adapter

- Specify the values for the following required properties by clicking their names in the table and entering the appropriate values.

- **queueName:** A comma-delimited list of the tpipes hold queue names where the callout requests from IMS are being held.
- **portNumber:** the port number for IMS Connect
- **hostName:** the IMS Connect host name
- **dataStoreName:** the IMS data store name.

In WebSphere Application Server Version 8 or later, the data store name can be a comma-delimited list of names, enabling an IMSActivationSpec instance to pull callout messages from multiple data stores.

- If you want the IMS TM resource adapter to use connection pooling for responses that are sent back to IMS, specify values for the following properties:
 - **connectionPoolEnabled:** whether IMS TM resource adapter should use connection pooling (specify `true` to enable this feature)
 - **maxConnections:** the maximum number of connections that the connection pool can hold
 - **poolCleanupFrequency:** the cleanup interval (in minutes) for idle connections in the connection pool
 - **durationOfIdleConnections:** the number of minutes that a connection in the connection pool can remain idle before it is cleaned up at the next cleanup interval
- If you want the IMS TM resource adapter to receive an acknowledgment from IMS when IMS receives a synchronous callout response message, specify `true` for the `sendOnlyWithAck` property.
- Specify the values for the other optional properties based on your environment.
- Click **Apply**.
- Click **Save** in the message box at the top to save the changes to the master configuration.

The J2C activation specification for the IMS TM resource adapter is created and configured.

IMSActivationSpec property configuration for message-driven beans

The values of the properties of the IMSActivationSpec object describe the inbound communication from IMS to be used by message-driven beans.

You must configure an instance of the IMSActivationSpec object when you deploy a message-driven bean in order for the IMS TM resource adapter to communicate with IMS Connect for retrieving and responding to synchronous callout messages.

The following table describes the properties for the IMSActivationSpec object, the implementation of the J2C activation specification in IMS TM resource adapter. Specify and configure the values for these properties by using the administrative console in WebSphere Application Server.

Properties in the J2C activation specification that are not listed in the following table are not supported by the IMS TM resource adapter.

Table 115. Properties for the IMSActivationSpec object

Property	Description
<code>dataStoreName</code>	Specifies the IMS data store name. The name must match the ID parameter of the Datastore statement that is specified in the IMS Connect configuration member when IMS Connect is installed. It also serves as the XCF member name for IMS during cross-system coupling facility (XCF) communications between IMS Connect and OTMA. The data store name can be a comma-separated list of data store names, enabling one instance of the IMSActivationSpec class to pull callout messages from more than one IMS data store.
<code>groupName</code>	Specifies the security authorization facility (SAF) group name.
<code>hostName</code>	Specifies the IMS Connect host name.

Table 115. Properties for the IMSActivationSpec object (continued)

Propertiesk	Description
password	Specifies the SAF password.
portNumber	Specifies the IMS Connect port number.
queueNames	Specifies a comma-delimited list of IMS OTMA tpipe names for the callout (synchronous or asynchronous) messages. Queue names must be 1 to 8 alphanumeric characters (A-Z, 0-9, @, #, \$).
retryInterval	Specifies the time delay in milliseconds before the IMS TM resource adapter tries to check on the availability of the data store.
retryLimit	Specifies the maximum number of times the IMS TM resource adapter will attempt to reconnect to IMS Connect if a connection is lost due to IMS Connect or IMS data store availability issues.
SSLEnabled	Instructs the IMS TM resource adapter to create a Secure Sockets Layer (SSL) socket connection to IMS Connect by using the specified host name and port number in the IMSActivationSpec object. This property is valid for TCP/IP connections only.
SSLEncryptionType	Specifies the SSL encryption type. Valid values are strong and weak.
SSLKeyStoreName	Specifies the name, including the full file path, of the keystore for TCP/IP SSL communications. Private keys and their associated public key certificates are stored in password-protected databases called keystores. Trusted certificates can also be stored in the keystore, and the truststore property can either be empty or could point to the keystore file. An example of a keystore name is c:\keystore\MyKeystore.ks. The file can have other file extensions.
SSLKeyStorePassword	Specifies the password for the keystore for TCP/IP SSL communications.
SSLTrustStoreName	Specifies the name, including the full path, of the keystore file that contains security credentials (certificates) for TCP/IP SSL communications. A value for the SSLTrustStoreName property is not mandatory if a keystore is used. A truststore file is a key database file that contains public keys or certificates. Private keys can also be stored in the truststore, and the keystore property can either be empty or could point to the truststore file. An example of a truststore name is c:\keystore\MyTruststore.ks. The file can have other file extensions.
SSLTrustStorePassword	The password for the truststore for TCP/IP SSL communications.
connectionPoolEnabled	Specifies whether to enable connection pooling. true Enables IMS TM Resource Adapter to create a connection pool for synchronous callout responses that are sent back to IMS. false Disables connection pooling. IMS TM Resource Adapter creates and then discards a new connection whenever a synchronous callout response is sent. The default value is false (for compatibility with earlier versions of IMS TM Resource Adapter).

Table 115. Properties for the IMSActivationSpec object (continued)

Property	Description
durationOfIdleConnections	Specifies the number of minutes that a connection in the connection pool can remain idle before it is cleaned up at the next cleanup interval. The default value is 0, which causes an idle connection to be cleaned up immediately. This property is ignored if the connectionPoolEnabled property is false.
maxConnections	Specifies the maximum number of connections that the connection pool can hold. If the pool is full, all new connections are discarded. The default value is 1. This property is ignored if the connectionPoolEnabled property is false.
poolCleanupFrequency	Specifies the cleanup interval (in minutes) for idle connections in the connection pool. The default value is 20. This property is ignored if the connectionPoolEnabled property is false.
sendOnlyWithAck	Specifies whether to enable the send-only-with-acknowledgment protocol. true Enables the protocol. IMS TM Resource Adapter receives an acknowledgment from IMS when IMS receives a synchronous callout response message. false Disables the protocol. The default value is false (for compatibility with earlier versions of IMS TM Resource Adapter).
userName	SAF user name.

Related information

[Configuring a J2C activation specification](#) For more information, see IBM Documentation for WebSphere Application Server.

Configuring a thread pool in WebSphere Application Server

Configure a thread pool in WebSphere Application Server if you have multiple data stores and multiple tpipes for synchronous callout request processing.

A thread pool enables components of the server to reuse threads, which eliminates the need to create new threads at run time. The IMS TM resource adapter thread pool could quickly exhaust and the callout processing would stall when multiple data stores and tpipes are involved.

If you have multiple data stores and multiple tpipes, configure the WebSphere Application Server to add your own thread pool. The IMS TM resource adapter thread pool could quickly exhaust and the callout processing would stall when multiple data stores and tpipes are involved. A possible thread pool size to start with is $\text{numberOfDatastore} * \text{numberOfQueues} + \text{numberOfActiveWorkload}$ (for example, 20 for 20 message-driven bean requests per second). However, the actual number requires testing based on your environment and usage.

1. Go to the WebSphere Application Server Admin Console.
2. Select **Application servers** > **server_name** > **Additional Properties** > **Thread Pools**.
3. Set the maximum size for the thread pool.

A possible thread pool size to start with is $\text{number\_of\_data\_store} * \text{number\_of\_queues} * \text{number\_of\_active\_workload}$, where the `number_of_active_workload` is the number of request per second that you would like to be processed.

The actual number requires testing based on your environment and usage.

4. Reference the thread pool in the IMS TM resource adapter configuration.

For example, if you add a thread pool called IMSTMRA as shown in the following example:

The screenshot shows the 'Application servers' configuration page in WebSphere. The breadcrumb trail is 'Application servers > server1 > Thread pools > IMSTMRA'. Below the breadcrumb, there is a description: 'Use this page to specify a thread pool for the server to use. A thread pool enables server components to reuse threads instead of creating new threads at run time. Creating new threads is typically a time and resource intensive operation.' The 'Configuration' tab is selected. The page is divided into two sections: 'General Properties' and 'Additional Properties'. Under 'General Properties', there are fields for 'Name' (IMSTMRA), 'Description' (Pool for IMS Resource Adapter), 'Minimum Size' (50 threads), 'Maximum Size' (200 threads), and 'Thread inactivity timeout' (5000 milliseconds). There is also a checkbox 'Allow thread allocation beyond maximum thread size' which is checked. At the bottom, there are buttons for 'Apply', 'OK', 'Reset', and 'Cancel'. Under 'Additional Properties', there is a link for 'Custom properties'.

Application servers > server1 > Thread pools > IMSTMRA

Use this page to specify a thread pool for the server to use. A thread pool enables server components to reuse threads instead of creating new threads at run time. Creating new threads is typically a time and resource intensive operation.

Configuration

General Properties

* Name
IMSTMRA

Description
Pool for IMS Resource Adapter

* Minimum Size
50 threads

* Maximum Size
200 threads

* Thread inactivity timeout
5000 milliseconds

☒ Allow thread allocation beyond maximum thread size

Apply OK Reset Cancel

Additional Properties

■ [Custom properties](#)

Figure 123. Thread pool configuration in WebSphere Application Server

You can add the thread pool to the Thread pool alias field for IMS TM resource adapter.

Resource adapters

Resource adapters > IMS TM Resource Adapter

Use this page to manage resource adapters, which provide the fundamental interface to an Enterprise Information System (EIS). The WebSphere(R) Relational Resource Adapter product to provide access to relational databases. To access another type of stand-alone resource adapter archive (RAR) file. You can configure multiple resource adapters.

Configuration

General Properties

Scope

cells:CL33213KNode01Cell:nodes:CL33213KNode01

Name

IMS TM Resource Adapter

Description

Archive path

\${CONNECTOR_INSTALL_ROOT}/ims1410_denis.rar

Class path

\${CONNECTOR_INSTALL_ROOT}/ims1410_denis.rar

Native library path

☐ Isolate this resource provider

Thread pool alias

IMSTMRA

Apply

OK

Reset

Cancel

Figure 124. Resource adapter configuration in WebSphere Application Server

Sample MDB for receiving callout requests from IMS

Because an MDB implements the `InboundListener` interface, you must implement the `onMessage()` and `onNotification()` methods that are used for synchronous and asynchronous processing, respectively.

The following sample shows a bean implementation class for an MDB (`IMSCalloutIVPMDBBean`) that is configured to listen for both synchronous and asynchronous callout messages from IMS.

```
package ejbs;

import java.io.IOException;
import javax.resource.ResourceException;
import javax.resource.cci.Record;
import com.ibm.connector2.ims.ico.IMSInputStreamRecord;
import commonj.connector.runtime.InboundInteractionSpec;
```

```

/**
 * Bean implementation class for Enterprise Bean: IMSInboundMessage
 */
public class IMSCalloutIVPMDBBean
    implements
        javax.ejb.MessageDrivenBean,
        com.ibm.j2ca.base.ExtendedInboundListener {

    private javax.ejb.MessageDrivenContext fMessageDrivenCtx;

    /**
     * getMessageDrivenContext
     */
    public javax.ejb.MessageDrivenContext getMessageDrivenContext() {
        return fMessageDrivenCtx;
    }

    /**
     * setMessageDrivenContext
     */
    public void setMessageDrivenContext(javax.ejb.MessageDrivenContext ctx) {
        fMessageDrivenCtx = ctx;
    }

    /**
     * ejbCreate
     */
    public void ejbCreate() {
    }

    /**
     * onMessage
     */
    public void onMessage(javax.jms.Message msg) {
    }

    /**
     * ejbRemove
     */
    public void ejbRemove() {
    }

    /**
     * onMessage
     */
    public javax.resource.cci.Record onMessage(
        javax.resource.cci.Record arg0,
        javax.resource.cci.InteractionSpec arg1)
        throws javax.resource.ResourceException {
        return processSyncMessageReceived(arg1);
    }

    /**
     * onNotification
     */
    public void onNotification(
        javax.resource.cci.Record arg0,
        javax.resource.cci.InteractionSpec arg1)
        throws javax.resource.ResourceException {
        processAsyncMessageReceived(arg0);
    }

    public void onNotification(Record arg0) throws ResourceException {
        processAsyncMessageReceived(arg0);
    }

    public Record onMessage(Record arg0) throws ResourceException {
        return processSyncMessageReceived(arg0);
    }

    public Record onMessage(InboundInteractionSpec arg0, Record arg1)
    throws ResourceException {
        return processSyncMessageReceived(arg1);
    }

    public void onNotification(InboundInteractionSpec arg0, Record arg1)
    throws ResourceException {
        processAsyncMessageReceived(arg1);
    }
}

```

```

/*
 * Process synchronous callout requests from IMS and return a response
 */
public Record processSyncMessageReceived(Object event) {
    SYNCCALLOUTREQUEST request = new SYNCCALLOUTREQUEST();
    SYNCCALLOUTRESPONSE response = new SYNCCALLOUTRESPONSE();

    try {
        // Request
        request.setBytes(((IMSInputStreamRecord)event).getBytes());

        System.out.println("Synchronous callout request from IMS: " +
            request.getSync__callout__request__str());

        // Response
        response.setSync__callout__response__str("HELLO FROM WEBSPPHERE MDB");

        System.out.println("Synchronous callout response from WAS MDB: " +
            response.getSync__callout__response__str());

    } catch (IOException e) {
        System.err.println(e);
    }

    return response;
}

/*
 * Receive asynchronous callout requests from IMS
 */
public void processAsyncMessageReceived(Object event) {
    ASYNCCALLOUTREQUEST request = new ASYNCCALLOUTREQUEST();

    try {
        request.setBytes(((IMSInputStreamRecord)event).getBytes());

        System.out.println("Asynchronous callout request from IMS: " +
            request.getAsync__callout__request__str());

    } catch (IOException e) {
        System.err.println(e);
    }
}
}

```

Non-managed (client-managed) callout programming model

You can use the *non-managed callout programming model* to retrieve IMS application callout requests from any external Java applications that do not use a message-driven bean.

For any response data to a callout request, the Java application can send the response data either synchronously to the same IMS application in the same transaction or asynchronously in a different transaction.

For a synchronous callout request, unlike in the managed callout programming model where the callout protocol and response correlation are handled for you through the message-driven bean, you must handle the following issues in your Java application:

- Handle the callout protocol in your application to poll the hold queue for any callout messages.
- Ensure that the correlation token that is attached to the callout request is passed back with the response in order for the response to be correlated back to the initiating request.

Because you must handle the synchronous callout protocol and the correlation of response messages to the request, this programming model is also called *client-managed callout programming model*.

In general, the following tasks are required when you use the client-managed callout programming model:

1. Your WebSphere Application Server administrator must configure a shareable connection factory to be used by the Java application to retrieve the callout requests.
2. Retrieve the callout messages from the IMS OTMA hold queue (tpipe).

- For synchronous callout requests:
 - You must correlate the response from the EJB application to the callout request. For more information, see [“Correlating the response from the Java application to the synchronous callout request” on page 777](#)
 - Follow the steps described in [“Retrieving synchronous callout request messages from non-MDB applications” on page 778](#)
- For asynchronous callout requests, see [“Retrieving asynchronous callout request messages from non-MDB applications” on page 780](#).

You can enable the send-only-with-acknowledgment protocol. Using this protocol, IMS TM Resource Adapter receives an acknowledgment from IMS when IMS receives a synchronous callout response message. In Version 13 and earlier, this acknowledgment is unavailable.

You can enable connection pooling for synchronous callout messages that are sent back to IMS. In Version 13 and earlier, the IMS TM Resource Adapter must create and then discard a new connection whenever a callout response is sent. By enabling connection pooling, you can avoid the overhead of opening and closing multiple connections.

Correlating the response from the Java application to the synchronous callout request

If your Java application generates an output response and you need to correlate the response with the initial synchronous callout request, your IMS application is responsible for such correlation.

For asynchronous callout, correlation can be achieved by defining some data, such as a message identifier or a unique request ID, in the callout request that can correlate with the initial input message.

For synchronous callout, use the correlator token that is passed with the callout message.

The following diagram demonstrates the high-level steps that are involved in a non-MDB application to retrieve an IMS callout request and to send the response back.




Figure 125. Process flow for retrieving and responding to a synchronous callout request from IMS for non-MDB Java applications

1. The Java application issues a SYNC_RECEIVE_CALLOUT request to the OTMA tpipe TP1.
2. An IMS application issues a DL/I ICAL call with the SENDRECV subfunction and specifies the OTMA destination descriptor name and the timeout value.
3. IMS OTMA delivers the request message to the IMSTM Resource Adapter through IMS Connect. A correlation token is sent with the callout request.
4. The Java application receives the callout request message and the correlation token. After processing the request, the Java application sends the response message and the correlation token back to IMS by using the SYNC_SEND request. With the correlation token, IMS correlates the response message back to the corresponding IMS transaction instance.

Retrieving synchronous callout request messages from non-MDB applications

You must modify the Java application to retrieve the callout requests in the hold queue and to handle the correlation of the response to the appropriate request.

Prerequisite: Your WebSphere Application Server administrator must configure a shareable connection factory to be used by the Java application to retrieve the callout requests.

The following steps are key steps in a Java application for retrieving and responding to a synchronous callout request from IMS.

1. To retrieve the synchronous callout request:
 - a) Set the interaction commit mode to 0 (COMMIT_THEN_SEND).

```
interactionSpec.setCommitMode(IMSInteractionSpec.COMMIT_THEN_SEND);
```

- b) Set the interaction sync level to CONFIRM.

```
interactionSpec.setSyncLevel(IMSInteractionSpec.SYNC_LEVEL_CONFIRM);
```

- c) Specify the callout queue name to the OTMA tpipe (hold queue) name.

```
// An 8-character name of the OTMA asynchronous hold queue that the
// messages are to be retrieved from
String calloutQueueName = new String ("CALLOUTQ");
// Set the queue name for the callout messages
interactionSpec.setAltClientID(calloutQueueName);
```

- d) Set the interaction verb to SYNC_RECEIVE_CALLOUT.

```
interactionSpec.setInteractionVerb(IMSInteractionSpec.SYNC_RECEIVE_CALLOUT);
```

- e) Specify the type of callout request to retrieve: synchronous only (CALLOUT_REQUEST_SYNC), asynchronous only (CALLOUT_REQUEST_ASYNC), or both (CALLOUT_REQUEST_BOTH).

The following example retrieves only synchronous callout request messages.

```
// Specify to retrieve only synchronous callout request messages
interactionSpec.setCalloutRequestType(IMSInteractionSpec.CALLOUT_REQUEST_SYNC);
```

- f) Specify the execution timeout value.

The following example sets the timeout value to 5 seconds.

```
interactionSpec.setExecutionTimeout(5000);
```

- g) Execute the interaction.

The following example executes the interaction, and the request is returned as calloutRequestMsg.

```
interaction.execute(interactionSpec, null, calloutRequestMsg);
```

2. Obtain the correlation token from the IMSInteractionSpec instance.

The following example sets the timeout value to 5 seconds.

```
byte[] corrToken = interactionSpec.getSyncCalloutCorrelationToken();
```

3. Process the request.

In the previous example, the request is returned as calloutRequestMsg.

4. Send the response.

- a) Set the interaction verb to SYNC_SEND.

```
interactionSpec.setInteractionVerb(IMSInteractionSpec.SYNC_SEND);
```

- b) Set the correlation token to be sent back with the request.

```
interactionSpec.setSyncCalloutCorrelationToken(corrToken);
```

- c) Execute the interaction.

The following example executes the interaction and sends the callout response message (calloutRespondMsg).

```
interaction.execute(interactionSpec, calloutRespondMsg, null);
```

Sample non-MDB Java application for retrieving and processing synchronous callout requests

To prepare a Java application for inbound requests from an IMS application program, specify the appropriate interaction verb, the tpipe name, and the timeout value in the Java application.

```
// JNDI lookup returns connFactory
InitialContext initCtx = new InitialContext();
ConnectionFactory connFactory =
(ConnectionFactory) initCtx.lookup("java:comp/env/ibm/ims/IMSTarget");
IMSConnectionSpec connSpec = new IMSConnectionSpec();
```

```

Connection connection = connFactory.getConnection(connSpec);
Interaction interaction = connection.createInteraction();
IMSInteractionSpec interactionSpec = new IMSInteractionSpec();

//set the commit mode and sync level
interactionSpec.setCommitMode(IMSInteractionSpec.COMMIT_THEN_SEND);
interactionSpec.setSyncLevel(IMSInteractionSpec.SYNC_LEVEL_CONFIRM);

// An 8-character name of the OTMA asynchronous hold queue that the
// messages are to be retrieved from
String calloutQueueName = new String ("CALLOUTQ");
// Set the queue name for the callout message
interactionSpec.setAltClientID(calloutQueueName);

// Set InteractionVerb for retrieving callout request with a timeout value
interactionSpec.setInteractionVerb(IMSInteractionSpec.SYNC_RECEIVE_CALLOUT);
// Specify to retrieve only synchronous callout request messages
interactionSpec.setCalloutRequestType(IMSInteractionSpec.CALLOUT_REQUEST_SYNC);
interactionSpec.setExecutionTimeout(5000);

// Execute the interaction
interaction.execute(interactionSpec, null, calloutRequestMsg);

// Get correlation token
byte[] corrToken = interactionSpec.getSyncCalloutCorrelatorToken();

// Further processing on the request (calloutRequestMsg)
:

// Send back the response (calloutRespondMsg) by using the SYNC_SEND interaction
interactionSpec.setInteractionVerb(com.ibm.connector2.ims.ico.IMSInteractionSpec.SYNC_SEND);
// SYNC_SEND does not support alternate client ID
interactionSpec.setAltClientID(null);
interactionSpec.setSyncCalloutCorrelatorToken(corrToken);

// Execute the interaction
interaction.execute(interactionSpec, calloutRespondMsg, null);

interaction.close();
connection.close();

```

Retrieving asynchronous callout request messages from non-MDB applications

You must modify your Java application to retrieve the callout requests in the hold queue, and if a response message is expected, correlate the response to the appropriate request.

Prerequisite: Your WebSphere Application Server administrator must configure a shareable connection factory to be used by the Java application to retrieve the callout requests.

Modify the Java application to set the following property values:

1. Specify an alternate client ID whose value is the name of the OTMA asynchronous hold queue (tpipe) on which the callout request is to be queued.
The tpipe that the callout Java application listens to must be reserved for receiving callout requests. The name of the OTMA asynchronous hold queue is:
 - The name of the ALTPCB to which the callout request is inserted, or
 - The name of the tpipe specified in the OTMA destination descriptor.
2. Set the interactionVerb property to SYNC_RECEIVE_CALLOUT (which corresponds to an IMS Connect Resume Tpipe Single Wait request).
3. Specify an execution timeout value that you would like to wait for a callout message in the hold queue.
4. For added security, you can optionally specify a user ID in your Java application if you have configured your IMS security so that only authorized users can retrieve the callout request message from the asynchronous hold queue.
This user ID must be specified in the connectionSpec object in your application or the connection factory used by your application.

Sample Java application for retrieving asynchronous callout requests

To prepare a non-MDB Java application for inbound requests from an IMS application program, specify the appropriate interaction verb, the asynchronous hold queue name, and a timeout value in the Java application.

- Specify the asynchronous hold queue name.

```
interactionSpec.setAltClientID(calloutQueueName);
```

- Set the interactionVerb property to SYNC_RECEIVE_CALLOUT.

```
interactionSpec.setInteractionVerb(com.ibm.connector2.ims.ico.  
IMSInteractionSpec.SYNC_RECEIVE_CALLOUT);
```

- Specify an execution timeout value (in milliseconds).

```
interactionSpec.setExecutionTimeout(3600000);
```

Specify a timeout value to wait for a callout request message in the hold queue. In this example, a large execution timeout value is specified. A large execution timeout value helps to minimize the looping required in a callout EJB component when there might be long periods of time when no callout requests occur. The length of the timeout depends on the maximum length of time that you expect to elapse between your callout requests and what you want your EJB component to do if that length of time is exceeded.

The following sample code shows an example of a Java application that is configured to receive the callout requests.

```
// JNDI lookup returns connFactory  
InitialContext initCtx = new InitialContext();  
ConnectionFactory connFactory =  
    (ConnectionFactory) initCtx.lookup("java:comp/env/ibm/ims/IMSTarget");  
  
IMSConnectionSpec connSpec = new IMSConnectionSpec();  
Connection connection = connFactory.getConnection(connSpec);  
Interaction interaction = connection.createInteraction();  
IMSInteractionSpec interactionSpec = new IMSInteractionSpec();  
  
//set the commit mode and sync level  
interactionSpec.setCommitMode(0);  
interactionSpec.setSyncLevel(1);  
  
// An eight-character queue name that the asynchronous messages to be retrieved from  
String calloutQueueName = new String ("CALLOUTQ");  
// Set the asynchronous queue name for the callout message  
interactionSpec.setAltClientID(calloutQueueName);  
  
// Set interactionVerb to retrieve async output with a timeout  
interactionSpec.setInteractionVerb(com.ibm.connector2.ims.ico.  
IMSInteractionSpec.SYNC_RECEIVE_CALLOUT);  
interactionSpec.setExecutionTimeout(3600000);  
  
for (;;) {  
    try {  
        // Execute the interaction  
        interaction.execute(interactionSpec, null, calloutMsg);  
  
        // Further processing on the calloutMsg  
        :  
    } catch (Exception e) {  
        // if the exception is an execution timeout error,  
        // you can either do nothing and continue to loop  
        // or process the error and then break the loop  
        break;  
    }  
}
```

```
interaction.close();
connection.close();
```

Commit mode and sync level processing

The commit mode that you choose determines the type of commit mode processing that IMS performs.

The Java client specifies the commit mode protocol that is used when it submits a transaction request to IMS. IMS Connect and IMS support two types of commit mode processing:

- Commit mode 0 (commit-then-send), where IMS commits the IMS database changes and then sends the output to the client
- Commit mode 1 (send-then-commit), where IMS sends the output to the client and then commits the database changes

Associated with the commit mode protocols, IMS Connect and IMS also support three synchronization levels (sync levels): NONE, CONFIRM, and SYNCPT (the sync-point protocol). All three sync levels can be used with commit mode 1. Only CONFIRM can be used with commit mode 0. The Java client cannot explicitly set the sync level to SYNCPT. The IMS TM resource adapter supports only the following combinations:

Commit mode	Sync level = NONE supported	Sync level = CONFIRM supported
commit mode 0	No	Yes
commit mode 1	Yes	Yes

Commit mode 1 with sync level SYNCPT is used by the IMS TM resource adapter when it participates in two-phase commit processing with IMS.

Interactions supported with commit mode and sync level combinations

The interactions that are supported by the IMS TM resource adapter depend on the combination of commit mode and sync level that you choose.

The following combinations of commit mode and sync level are supported by the IMS TM resource adapter:

- Commit mode 1 with sync level NONE

This combination is used for non-transactional interactions. For non-conversational applications, use the SYNC_SEND_RECEIVE interaction. For conversational applications, use the SYNC_SEND_RECEIVE or, optionally, the SYNC_END_CONVERSATION interaction.

- Commit mode 1 with sync level CONFIRM

This combination is used for non-transactional interactions. For non-conversational applications, use the SYNC_SEND_RECEIVE interaction. For conversational applications, use the SYNC_SEND_RECEIVE or, optionally, the SYNC_END_CONVERSATION interaction.

- Commit mode 0 with sync level CONFIRM

This combination is used for non-transactional SYNC_SEND_RECEIVE, SYNC_SEND, SYNC_RECEIVE_ASYNCPUT_SINGLE_NOWAIT, SYNC_RECEIVE_ASYNCPUT_SINGLE_WAIT and SYNC_RECEIVE_CALLOUT interactions.

Tip: Commit mode 0 is supported only for non-conversational applications that are running on TCP/IP connections.

For commit mode 1 interactions, the IMS TM resource adapter automatically provides the synchronization level when it communicates with IMS Connect. The synchronization level can also be set explicitly by the Java client by using the `setSyncLevel(int)` method to NONE or CONFIRM. The default sync level is NONE. You can set the interaction verb to SYNC_SEND_RECEIVE without setting the sync level.

For commit mode 0 interactions, only the CONFIRM sync level is supported.

The `setSyncLevel(int)` method throws an exception for any invalid combination of commit mode and sync level.

Related concepts

Java exceptions that involve output messages

If any errors occur when the IMS TM resource adapter passes the message through IMS Connect to IMS, and as IMS processes this transaction and attempts to return the output, the Java client receives an exception.

Specification of transaction pipe, client ID, and interaction verb

You must specify an interaction verb regardless of the commit mode that you use. Specification of the transaction pipe (tpipe) and client ID depends on the commit mode or the type of socket connection that you sue.

In IMS and OTMA terminology, a transaction pipe (tpipe) is a logical connection between a client (IMS Connect) and the server (IMS OTMA). The `IMSConnectionSpec` class has a `clientId` property that identifies the connection associated with that `IMSConnectionSpec`. The name of the IMS OTMA asynchronous output queue or tpipe where recoverable output messages are placed during an interaction is determined in one of two ways, depending on the commit mode that is used for that interaction:

- For commit mode 0 interactions, the tpipe is identified by the client ID that is used for the interaction. Each client ID that is used for a commit mode 0 transaction has its own tpipe.
 - The client ID is automatically generated for interactions on shareable socket connections.
 - The client ID must be set by the user for interactions on dedicated socket connections.

In general, use shareable socket connections rather than dedicated socket connections. Shareable socket connections allow for more efficient use of connections. Shareable socket connections also are less prone to problems in Sysplex Distributor and WebSphere Application Server for z/OS clone environments. In these environments, multiple instances of the WebSphere Application Server might communicate with the same IMS Connect and try to use the same client ID for their connections.

- For commit mode 1 interactions, the tpipe is identified by the IMS Connect port number used for the interaction. Therefore, each port has a tpipe that is used for all clients that run commit mode 1 interactions on that port.

Whether your Java client is running an IMS transaction with commit mode 1 or commit mode 0, the Java client specifies a value for the `interactionVerb` property for the interaction. If a commit mode 0 interaction is specified and a dedicated socket connection is used by the Java client, the Java client also must provide a value for the client ID for the connection.

Socket connections

All socket connections that are created by the IMS TM resource adapter are persistent, and they can be shareable or dedicated.

Shareable persistent sockets can be shared by applications that execute either commit mode 0 or commit mode 1 interactions. Dedicated persistent sockets can be used only for applications that execute commit mode 0 interactions.

The same socket connection between the IMS TM resource adapter and IMS Connect can be serially reused for multiple interactions with IMS Connect. The socket connection will not be closed and reopened between interactions.

Dedicated persistent sockets

A dedicated persistent socket is a socket connection that is assigned to a specific client ID and remains dedicated to that particular client ID until it is disconnected.

Typically, a dedicated persistent socket connection is used (or reused) serially by a single application. However, a dedicated persistent socket connection can also be used serially by multiple applications, which can obtain the socket by using the same client ID.

If different applications attempt to use a dedicated persistent socket connection by using the same client ID at the same time, the attempt by the first application will succeed. Subsequent attempts by other applications when the first application is still using its connection will receive duplicate client errors (DUPCLNT) from IMS Connect.

A duplicate client error might also occur when an application that obtains a dedicated persistent socket connection by using a given client ID runs in an environment in which dedicated persistent socket connections originate from multiple WebSphere Application Server instances. For example, an environment that has multiple WebSphere Application Server instances that are managed by a workload manager might encounter duplicate client errors.

Recommendation: Do not deploy your application to this type of environment if it uses a dedicated persistent socket connection.

Dedicated persistent socket connections are intended to be used in an environment in which a single instance of WebSphere Application Server is connected to a single IMS Connect. In some situations, a single instance of WebSphere Application Server can be connected to multiple instances of IMS Connect. In this configuration, if Sysplex Distributor is used between WebSphere Application Server and IMS Connect, the following interactions are not supported:

- Two-phase commit recovery
- IMS conversational transactions

Supported commit mode and interactions

A dedicated persistent socket can be used only for Java applications that execute commit mode 0 interactions.

Table 116. Supported commit modes, interaction verbs, and sync levels by socket type

Commit mode	Socket type	Interaction verb	Client ID	Sync level
CM1	Shareable persistent	• SYNC_END_CONVERSATION (for conversational transactions) • SYNC_SEND_RECEIVE	Automatically generated and managed	None or Confirm

Table 116. Supported commit modes, interaction verbs, and sync levels by socket type (continued)

Commit mode	Socket type	Interaction verb	Client ID	Sync level
CM0	Shareable persistent	• SYNC_SEND • SYNC_SEND_RECEIVE • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT • SYNC_RECEIVE_CALLOUT (for non-managed callout programming model)	Automatically generated and managed	Confirm
	Dedicated persistent	• SYNC_SEND • SYNC_SEND_RECEIVE • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT	Specified by client applications (the clientID property)	Confirm

Table 117. Supported commit modes, interaction verbs, and sync levels by socket type

Commit mode	Socket type	Interaction verb	Client ID	Sync level
CM1	Shareable persistent	• SYNC_END_CONVERSATION (for conversational transactions) • SYNC_SEND_RECEIVE	Automatically generated and managed	None or Confirm
CM0	Shareable persistent	• SYNC_SEND • SYNC_SEND_RECEIVE • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT • SYNC_RECEIVE_CALLOUT (for non-managed callout programming model)	Automatically generated and managed	Confirm
	Dedicated persistent	• SYNC_SEND • SYNC_SEND_RECEIVE • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT	Specified by client applications (the clientID property)	Confirm

The SYNC_RECEIVE_ASYNCOUTPUT interaction verb is deprecated and replaced by the SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT interaction verb.

Message retrieval

SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions (along with the deprecated SYNC_RECEIVE_ASYNCOUTPUT interaction) on dedicated persistent sockets enable client applications to retrieve messages from the following sources:

- Messages that were placed on an IMS OTMA asynchronous output queue as a result of a failed commit mode 0 interaction
- Messages that were from an IMS application that issued an insert to an Alternate Program Communication Block (ALTPCB)
- Messages from the reroute of the output from a transaction that was executed on a shareable socket connection

To retrieve these messages, the client application must provide the client ID, which represents the tpipe that asynchronous output messages are queued. Undelivered output messages from interactions on dedicated persistent sockets cannot be rerouted or purged.

You might not be able to retrieve all available output messages with the SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT interactions when you use multiple IMS systems without the OTMA Super Member.

Dedicated persistent socket connections

Dedicated persistent socket connections are created by an IMS connection factory with values for at least the following custom properties:

- Host name: TCP/IP host name of the system that is running IMS Connect
- Port number: associated port number
- Data store name: name of the target IMS
- CM0 dedicated: `true`

A value of `true` for the `CM0Dedicated` property ensures that the connection factory creates dedicated persistent socket connections.

If more than one connection factory is configured to create dedicated persistent sockets to the same IMS Connect instance, only one connection factory can dedicate a socket to a particular client ID at one time. For example, assume that a connection factory successfully creates a socket connection that is dedicated to a client ID, `CLIENT01`. A second connection factory also tries to create a socket connection that is dedicated to `CLIENT01`. If the socket connection that is created by the first connection factory is still connected to IMS Connect, the second connection will receive the following exception:

```
javax.resource.spi.EISSystemException: IC00001E:
com.ibm.connector2.ims.ico.IMSTCIPManagedConnection@23766050.processOutputOTMAMsg
(byte [], InteractionSpec,Record) error. IMS Connect returned error: RETCODE=[8],
REASONCODE=[DUPECLNT].
Duplicate client ID was used; the client ID is currently in use.
```

Related tasks

[Establishing dedicated persistent socket connections](#)

Use WebSphere Application Server to configure a J2C connection factory that creates dedicated persistent socket connections.

Related reference

[Client ID \(clientID\)](#)

A value for this property must be provided when an application component uses a dedicated persistent socket connection to IMS Connect to identify this particular client.

Shareable persistent sockets

A shareable persistent socket is a socket that can be shared (serially reused) by multiple applications that execute either commit mode 1 (CM1) or commit mode 0 (CM0) interactions.

For an application that executes a commit mode 0 interaction on a shareable persistent socket, the IMS TM resource adapter automatically generates a client ID with the prefix HWS. This client ID represents and identifies the socket connection and the associated OTMA tpipe. For this type of socket, only client IDs that are generated by the IMS TM resource adapter are allowed. If you specify a client ID for a connection that has not been defined as CM0 on a shareable persistent socket, an exception is thrown and returned to the client.

Restriction: IMS application programs that insert messages to an alternate PCB must not use names that begin with HWS for the alternate PCBs. IMS TM resource adapter cannot retrieve HWS messages.

Tip: Client IDs are different from alternate client IDs. The alternate client ID is a property of the `IMSInteractionSpec` class. Alternate client IDs are used to retrieve asynchronous output messages from OTMA hold queues (tpipes).

Supported commit mode and interaction

A shareable persistent socket can be used for Java applications that execute either commit mode 1 or commit mode 0 interactions.

Table 118. Supported commit modes, interaction verbs, and sync levels by socket type

Commit mode	Socket type	Interaction verb	Client ID	Sync level
CM1	Shareable persistent	• SYNC_END_CONVERSATION (for conversational transactions) • SYNC_SEND_RECEIVE	Automatically generated and managed	None or Confirm
CM0	Shareable persistent	• SYNC_SEND • SYNC_SEND_RECEIVE • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT • SYNC_RECEIVE_CALLOUT (for non-managed callout programming model)	Automatically generated and managed	Confirm
	Dedicated persistent	• SYNC_SEND • SYNC_SEND_RECEIVE • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT	Specified by client applications (the <code>clientId</code> property)	Confirm

Table 119. Supported commit modes, interaction verbs, and sync levels by socket type

Commit mode	Socket type	Interaction verb	Client ID	Sync level
CM1	Shareable persistent	• SYNC_END_CONVERSATION (for conversational transactions) • SYNC_SEND_RECEIVE	Automatically generated and managed	None or Confirm
CM0	Shareable persistent	• SYNC_SEND • SYNC_SEND_RECEIVE • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT • SYNC_RECEIVE_CALLOUT (for non-managed callout programming model)	Automatically generated and managed	Confirm
	Dedicated persistent	• SYNC_SEND • SYNC_SEND_RECEIVE • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT • SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT	Specified by client applications (the clientID property)	Confirm

The SYNC_RECEIVE_ASYNCOUTPUT interaction verb is deprecated and replaced by the SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT interaction verb.

Message retrieval

Any output message that cannot be delivered to a Java client that executes a CM0 interaction on a shareable persistent socket can be queued for later retrieval. Secondary output is always treated as CM0 output and can be requested for subsequent retrieval.

Consider the situation in which a CM1 or CM0 interaction on a shareable persistent socket executes a program-to-program switch. If both programs return output, the second output message returned is treated as a CM0 output. It can be queued again and then retrieved later.

If a CM1 or CM0 interaction on a shareable persistent socket creates a program-to-program switch that invokes another CM0 interaction that results in secondary output (which is always CM0), the output can be queued again for later retrieval.

If undelivered output messages that are queued in the IMS OTMA asynchronous hold queue or tpipe are not retrieved by using the alternate client ID, the original interaction and the retrieval must be invoked within the same client application by using the same socket connection. The same generated client ID that identifies the shareable socket connection and the associated OTMA tpipe must be used for both the original interaction and the retrieval of any asynchronous output from that interaction.

On shareable persistent sockets, undelivered output messages can be handled in several ways:

- Retrieve undeliverable output by specifying the OTMA hold queue name that has the asynchronous output messages as the alternate client ID property value of the IMSInteractionSpec object.
- Purge the undelivered output. To purge undelivered output messages, the value of the IMSInteractionSpec purgeAsyncOutput property must be set to `true`. This input property determines if IMS Connect purges the undelivered I/O PCB output. The purgeAsyncOutput property is valid with the SYNC_SEND_RECEIVE interaction verb only. If the property is not specified on SYNC_SEND_RECEIVE, the default is `true`.

- Reroute the messages to another destination. You can reroute the undelivered output message to a different destination by setting the `IMSInteractionSpec reRoute` property to `true`. This property is valid only for the `SYNC_SEND_RECEIVE` and `SYNC_SEND` interaction verbs. If the `reRoute` property is set to `true`, the undelivered output message is queued to a named destination that is usually provided by the client application. This destination is specified in the `IMSInteractionSpec reRouteName` property. If the `reRoute` property is set to `true` and no `reRouteName` is provided, the value of the `reRouteName` property is the value specified in the IMS Connect configuration file. If no value is specified in the IMS Connect configuration file, the default value `HWS$DEF` is used.

Shareable persistent socket connections

Shareable persistent socket connections are created by an IMS connection factory with required values for the following custom properties:

- Host name: TCP/IP host name of the system that is running IMS Connect
- Port number: associated port number
- Data store name: name of the target IMS
- CM0 dedicated: `false`

The default value for the `CM0Dedicated` property is `false`, which ensures that the connection factory creates a shareable persistent socket connection unless the `CM0Dedicated` property has been explicitly set to `true`.

Related tasks

Establishing shareable persistent socket connections

Use WebSphere Application Server to configure a J2C connection factory that creates shareable persistent socket connections.

Releasing and reconnecting persistent sockets

A TCP/IP connection between the IMS TM resource adapter and IMS Connect is persistent and remains open for both a shareable persistent socket connection and a dedicated persistent socket connection unless an error occurs.

The socket might close when either the IMS TM resource adapter or IMS Connect encounter an error, or when WebSphere Application Server disconnects the socket for connection pool management.

In a dedicated persistent socket connection, the socket connection can be used only by interactions that have the same client ID that was used to establish the connection. The number of socket connections increases as new client IDs are used for interactions on dedicated persistent socket connections.

If you set both the maximum connections property and the connection timeout property in the WebSphere Application Server administrative console to a non-zero value, when the maximum connections value is reached and all the connections are in use, the application gets a `ConnectionWaitTimeoutException` exception after the connection timeout has elapsed. This behavior is a standard behavior for WebSphere Application Server. The `ConnectionWaitTimeoutException` applies to both dedicated persistent sockets and shareable persistent sockets.

However, if the maximum connections value is reached and one of the persistent socket connections is not in use, WebSphere Application Server would disconnect that socket to respond to the request to create a new persistent socket connection. This standard behavior of WebSphere Application Server applies to both dedicated persistent sockets and shareable persistent sockets.

Tip: WebSphere Application Server has several connection pool settings that can help manage the pool and redistribute persistence sessions. Use the aged timeout setting to specify the interval in seconds before a physical connection is discarded.

Important: If you are not using an application server to help manage the connection (also referred to as a *non-managed* environment), for commit mode 0 (CM0, or commit-then-send) transactions, you must manage the connection pooling yourself. Because CM0 transactions are recoverable, IMS Connect creates

a separate TPIPE for each client that uses CM0. Without an application server to manage the connection pool, many TPIPEs would be created and ultimately overload the system.

Socket reconnect

The IMS TM resource adapter has a socket reconnect feature that tries to reestablish a stale connection in a connection pool when one of the connections encounters a communication problem in the process of sending a request to or receiving a response from IMS Connect.

When a request comes in, if the IMS TM resource adapter detects a stale connection, the adapter throws an exception. When the next request comes in and tries to use the stale connection, the IMS TM resource adapter checks to see if IMS Connect is up. If IMS Connect is up, the IMS TM resource adapter would reconnect before submitting the interaction request. When the connection is restored, an ICO0140I informational message is logged.

This socket reconnect feature increases the fault-tolerance of the IMS TM resource adapter when a temporary network connection issue occurs. More importantly, with this feature you can recycle IMS Connect as part of your system maintenance without having to resubmit any IMS TM resource adapter interactions from the client application.

When a sysplex distributor is used between the IMS TM resource adapter and WebSphere Application Server, a WebSphere Application Server connection pool might have connections that are connected to more than one IMS Connect. If one of the IMS Connect instances is being recycled, new requests are directed to other IMS Connect instances. New requests would go to the recycled IMS Connect only when this IMS Connect instance comes back up and if other IMS Connect instances can no longer handle the workload.

To help redistribute the workload among the IMS Connect instances, you can use the aged timeout property in WebSphere Application Server connection pool settings. Because connections with an expired age are discarded, a new connection helps redistribute the connection to recycled IMS Connect instances.

For a callout request, the IMS TM resource adapter would attempt to reconnect based on the values of the IMSActivationSpec retryLimit and retryInterval properties. The retryLimit property specifies the maximum number of times the IMS TM resource adapter would attempt to reconnect to IMS Connect if a connection is lost, or to IMS if the connection to the IMS data store is lost. The retryInterval property specifies the time delay before the next attempt to check the IMS Connect status.

Related information

Aged timeout information in WebSphere Application Server V8 documentation For more information about aged timeout, see the WebSphere Application Server Version 8 documentation.

IMSInteractionSpec property configuration

You must provide a configured IMSInteractionSpec object in your Java application to interact with IMS.

The values of the IMSInteractionSpec properties describe the interaction with IMS, such as the maximum amount of time the IMS TM resource adapter waits for a response from IMS Connect, and the commit mode and sync level for the interaction.

Values for IMSInteractionSpec properties can be provided directly through the set methods by an application that uses the Common Client Interface. They can also be provided to a wizard of an IDE that generates code for the application.

Some of the properties are input-only properties, some are input and output properties, and some are output-only properties. Input-only properties provide input to the IMS TM resource adapter, and their values do not change. Output-only properties can be interrogated by the Java application (also called application component) to determine additional information about the interaction.

Related reference

[IMS interaction specification properties](#)

The properties of the `IMSInteractionSpec` object describe the interaction with IMS.

Submitting commands to IMS

Although the IMS TM resource adapter is intended primarily for you to run transactions on a host IMS system through a Java EE application, you can also issue IMS commands that are supported by IMS OTMA from your Java applications.

The IMS TM resource adapter uses the host product, IMS Connect, to access IMS. IMS Connect uses the cross-system coupling facility (XCF) to access IMS through OTMA.

Only certain IMS commands can be submitted through the IMS OTMA interface. Because the IMS TM resource adapter accesses IMS through OTMA, IMS commands supported by OTMA are the only commands that can be submitted to IMS by an application that uses the IMS TM resource adapter.

The output of an IMS command is a message that consists of one or more segments of data. The output of some IMS commands is a DFS message. For example, the output of most **/START** commands is usually the message `DFS058I START COMMAND COMPLETED`. Other IMS commands do not return DFS messages. For example, **/DISPLAY** commands return multiple segments of data representing lines of display information. To treat both types of output the same, you must set the `imsRequestType` property of the `IMSInteractionSpec` class to 2 (`IMS_REQUEST_TYPE_IMS_COMMAND`). This value indicates to the IMS TM resource adapter that the interaction is an IMS command, and to treat DFS messages as normal output and not as Java exceptions.

Commands that can be submitted to IMS by an application that uses the IMS TM resource adapter can be found in the IMS commands reference.

Configuring IMS connection factories

You use the IMS connection factory in your application to interact with the IMS TM resource adapter. Through the IMS connection factories, you create pre-configured connections to the IMS Transaction Manager (IMS TM).

Connections between the IMS TM resource adapter and IMS Connect can be managed or non-managed.

Managed connections are connections that are managed by a Connection Manager in a Java EE application server such as WebSphere Application Server. In a managed environment, applications do not need to create EIS connections. Instead, they request connections from the Java EE Connection Manager. *Non-managed connections*, however, are obtained directly by applications from the IMS TM resource adapter.

Use TCP/IP socket to connect the IMS TM resource adapter with IMS Connect. Your application uses a Java EE Connection Architecture (JCA) connection factory to obtain a connection to IMS Connect.

When your application uses managed connections, all connections between the IMS TM resource adapter and IMS Connect are persistent. A connection is not opened and closed for each interaction. Instead, it can remain open and can be serially reused by multiple interactions, which provides better utilization of CPU and memory resources. Your Java EE application server administrator configures the connection factory at deployment time.

A non-managed connection must be opened and closed by the application that uses it. Your application must create and configure the connection factories that it needs at run time.

Recommendation: Use managed connections, and get a reference to the appropriate connection factory by using JNDI lookup.

The IMS TM resource adapter supports only persistent connections to IMS Connect. In a managed environment, persistent connections are serially reused by application components, without the need to disconnect and reconnect between each use. When using TCP/IP socket connections, application components can use either dedicated or shareable persistent socket connections between the IMS TM resource adapter and IMS Connect. Shareable persistent socket connections use client IDs that are generated by the IMS TM resource adapter and are not visible to the application that uses the shareable socket.

Recommendation: Always use shareable TCP/IP sockets unless there is some overriding requirement to use dedicated socket connections. Dedicated persistent socket connections use client IDs that are specified by the application.

TCP/IP connections to IMS Connect

TCP/IP can be used for the connections between the IMS TM resource adapter and IMS Connect in both distributed and z/OS environments.

When your application is on a distributed platform, the connection between the IMS TM resource adapter and IMS Connect must be a TCP/IP connection. For example, if the IMS TM resource adapter is installed on WebSphere Application Server on Windows, AIX, Solaris, Linux, Linux for System z, or HP-UX, the IMS TM resource adapter must connect with IMS Connect by using TCP/IP connections.

TCP/IP connections to IMS Connect are associated with an identifier called the *client ID*. IMS Connect ensures the uniqueness of these client IDs for all of the socket connections to it.

TCP/IP connections to IMS Connect can be one of the following two types:

- Dedicated persistent socket connections
- Shareable persistent socket connections

The client ID is determined differently for the two types of socket connections.

Related concepts

Shareable persistent sockets

A shareable persistent socket is a socket that can be shared (serially reused) by multiple applications that execute either commit mode 1 (CM1) or commit mode 0 (CM0) interactions.

Dedicated persistent sockets

A dedicated persistent socket is a socket connection that is assigned to a specific client ID and remains dedicated to that particular client ID until it is disconnected.

Establishing dedicated persistent socket connections

Use WebSphere Application Server to configure a J2C connection factory that creates dedicated persistent socket connections.

To establish dedicated persistent socket connections:

1. In WebSphere Application Server administrative console, configure a J2C connection factory.
2. Specify a value of TRUE for the CM0Dedicated property of the connection factory. For managed connections, your application uses JNDI to locate this connection factory.
3. In your application, set the commitMode property of the IMSInteractionSpec class to 0.
4. Provide a value for the clientID property of the IMSConnectionSpec class that is used to obtain the connection. If your application uses generated code, you can set this value by exposing the clientID property as data. If your application uses the Common Client Interface (CCI) of the IMS TM resource adapter, set this value by using the setClientID method of the IMSConnectionSpec object that is passed in the getConnection method.

Related concepts

Dedicated persistent sockets

A dedicated persistent socket is a socket connection that is assigned to a specific client ID and remains dedicated to that particular client ID until it is disconnected.

Establishing shareable persistent socket connections

Use WebSphere Application Server to configure a J2C connection factory that creates shareable persistent socket connections.

To establish shareable persistent socket connections:

1. In WebSphere Application Server administrative console, configure a J2C connection factory.

2. Specify a value of FALSE for the `CM0Dedicated` property of the connection factory. For managed connections, your application uses JNDI to locate this connection factory.
3. For shareable persistent socket connections, the client ID is set for you automatically by the IMS TM resource adapter. If your application uses generated code, the shareable persistent socket connection is obtained by the generated code. If your application uses the Common Client Interface (CCI) of the IMS TM resource adapter, ensure that you do not provide a value for the `clientId` property of the `IMSConnectionSpec` object that is passed in the `getConnection` method.

Related concepts

[Shareable persistent sockets](#)

A shareable persistent socket is a socket that can be shared (serially reused) by multiple applications that execute either commit mode 1 (CM1) or commit mode 0 (CM0) interactions.

IMS connection factory

The properties of an IMS TM resource adapter connection factory must be configured properly to match the characteristics of the target Enterprise Information System (EIS).

When you create an IMS service definition or define an IMS connection factory to WebSphere Application Server (managed connections), you must provide values for certain properties of the connection between the IMS TM resource adapter and IMS Connect. A connection factory provides the required information, such as host name and data store name, to connect to the EIS. Other connection information might be needed depending on the security settings and the type of connections.

If you must use non-managed connections, the `IMSConnectionSpec` class provides the API to set the connection properties.

Related tasks

[Creating a connection factory in WebSphere Application Server](#)

After you install and deploy the IMS TM resource adapter on WebSphere Application Server, create a connection factory for the IMS TM resource adapter, if it does not already exist.

Related reference

[IMS connection factory properties](#)

The properties of an IMS TM resource adapter connection factory describes the characteristics of the target Enterprise Information System.

Input and output message formats

You must consider the platforms where your target host application and the Java client run, and any special characteristics your input and output message might have when you develop an application that uses the IMS TM resource adapter.

The Java application that you create is running on a different platform (for example, Windows, AIX, and Solaris) than the enterprise information system (EIS) platform. IMS Transaction Manager (IMS TM) is an IMS application program on z/OS). Because of the platform differences, you must convert the text data in your messages from text data in UNICODE at the client end to text data in EBCDIC that is used by your IMS application.

Depending on the Endian value for the platform where the target host application runs, and the code page that is used by your IMS host system, format conversion might be required. The default is US English (037).

Recommendation: Let the integrated development environment (IDE) handle the format conversion for you. For example, the J2C wizard in various WebSphere or Rational development environment creates Java data bindings from imported C, COBOL, and PL/I data structures. The data bindings do all the format conversion for your messages at run time. These data bindings are created based on the options that you select in the Importer wizard.

IMS requires that all transaction messages be prefixed by a 2-byte LL field, a 2-byte ZZ field, followed by the transaction code. The LL field specifies the length of the message. The ZZ is a field of binary zero.

When you build input messages to IMS, you must set the LL value to the size of the input message, and set ZZ to 0.

Multi-segment and variable length output

If the IMS application returns multi-segment output or output with variable lengths, you can use the COBOL definition OUTPUT-MSG to define the output of the transaction, or create an output message for the output of the transaction.

In the following COBOL application, the output message returned by IMS consists of three fixed-length segments. The total output message returned by this IMS application is a fixed size of 99 bytes and is represented by the COBOL 01 structure OUTPUT-MSG. OUTPUT-SEG1 is 16 bytes in length, OUTPUT-SEG2 is 31 bytes, and OUTPUT-SEG3 is 52 bytes.

```
LINKAGE SECTION.

01  INPUT-MSG.
    02  IN-LL          PICTURE S9(3) COMP.
    02  IN-ZZ          PICTURE S9(3) COMP.
    02  IN-TRCD        PICTURE X(5).
    02  IN-DATA1        PICTURE X(6).
    02  IN-DATA2        PICTURE X(6).

01  OUTPUT-MSG.
    02  OUT-ALLSEGS    PICTURE X(99) VALUE SPACES.

01  OUTPUT-SEG1.
    02  OUT-LL          PICTURE S9(3) COMP VALUE +0.
    02  OUT-ZZ          PICTURE S9(3) COMP VALUE +0.
    02  OUT-DATA1       PICTURE X(12) VALUE SPACES.

01  OUTPUT-SEG2.
    02  OUT-LL          PICTURE S9(3) COMP VALUE +0.
    02  OUT-ZZ          PICTURE S9(3) COMP VALUE +0.
    02  OUT-DATA1       PICTURE X(13) VALUE SPACES.
    02  OUT-DATA2       PICTURE X(14) VALUE SPACES.

01  OUTPUT-SEG3.
    02  OUT-LL          PICTURE S9(3) COMP VALUE +0.
    02  OUT-ZZ          PICTURE S9(3) COMP VALUE +0.
    02  OUT-DATA1       PICTURE X(15) VALUE SPACES.
    02  OUT-DATA2       PICTURE X(16) VALUE SPACES.
    02  OUT-DATA3       PICTURE X(17) VALUE SPACES.
```

You can use the J2C wizard to create a J2C bean that runs the IMS transaction by :

- Creating a message buffer class to store the output message returned from the IMS application
- Importing the COBOL file to do the COBOL-to-Java mapping for both the input and output messages
 - Creating the input binding operation by selecting the input message structure (represented by the COBOL 01 structure INPUT-MSG)
 - Creating the output binding operations for the segments of the output message by using the data mapping wizard, available by selecting **File > New > Other > CICS/IMS Java Data Binding**. For each segment, use the data binding wizard to select the correct segment as the data structure and specify the parameters such as the code page, endian name, quote name, and trunc name.
- Invoking the J2C bean method that runs the IMS transaction, and populating the output segments from the buffer of data returned by the IMS transaction

```
package sample.ims;

import com.ibm.etools.marshall.util.MarshallIntegerUtils;
import sample.ims.data.*;

public class TestMultiSeg
{
    public static void main (String[] args)
    {
        byte[] segBytes = null;
        int srcPos = 0;
        int dstPos = 0;
```

```

int totalLen = 0;
int remainLen = 0;
byte[] buff;
short LL = 0;
short ZZ = 0;

try
{
    // -----
    // Populate the IMS transaction input message with
    // data. Use the input message format handler method
    // getSize() to set the LL field of the input message.
    // -----
    InputMsg input = new InputMsg();
    input.setIn_ll((short) input.getSize());
    input.setIn_zz((short) 0);
    // -----
    // find out the transaction code from your IMS
    // administrator
    // -----
    input.setIn_trcd("SKS6 ");
    input.setIn_data1("M2 SI1");
    input.setIn_data2("M3 SI1");

    // -----
    // Run the IMS transaction. The multi-segment output
    // message is returned.
    // -----
    MSOImpl proxy = new MSOImpl();

    sample.ims.CCIBuffer output = proxy.runMultiSegOutput(input);

    // -----
    // Retrieve the multi-segment output message as a
    // byte array using the output message format
    // handler method getBytes().
    // -----
    System.out.println(
        "\nSize of output message is: " + output.getSize());
    segBytes = output.getBytes();

    srcPos = 0;
    dstPos = 0;
    totalLen = segBytes.length;
    remainLen = totalLen;

    // -----
    // Populate first segment object from buffer.
    // -----
    buff = null;
    // Get length of segment.
    LL =
        MarshallIntegerUtils.unmarshallTwoByteIntegerFromBuffer(
            segBytes,
            srcPos,
            true,
            MarshallIntegerUtils.SIGN_CODING_TWOS_COMPLEMENT);

    // Put segment in byte array.
    buff = new byte[LL];
    System.arraycopy(segBytes, srcPos, buff, dstPos, LL);
    remainLen -= LL;

    // Create and populate segment object from byte array.
    OutputSeg1 S1 = new OutputSeg1();
    S1.setBytes(buff);
    System.out.println(
        "\nOutputSeg1 LL is:      "
        + S1.getOut_ll()
        + "\nOutputSeg1 ZZ is:      "
        + S1.getOut_zz()
        + "\nOutputSeg1_DATA1 is: "
        + S1.getOut_data1());

    // -----
    // Populate second segment object from buffer.
    // -----
    srcPos += LL;
    buff = null;
    // Get length of segment.
    LL =
        MarshallIntegerUtils.unmarshallTwoByteIntegerFromBuffer(

```

```

        segBytes,
        srcPos,
        true,
        MarshallIntegerUtils.SIGN_CODING_TWOS_COMPLEMENT);

// Put segment in byte array.
buff = new byte[LL];
System.arraycopy(segBytes, srcPos, buff, dstPos, LL);
remainLen -= LL;

// Create and populate segment object from byte array.

OutputSeg2 S2 = new OutputSeg2();
S2.setBytes(buff);
System.out.println(
    "\nOutSeg2 LL is:      "
    + S2.getOut__ll()
    + "\nOutSeg2 ZZ is:    "
    + S2.getOut__zz()
    + "\nOutSeg2_DATA1 is: "
    + S2.getOut__data1()
    + "\nOutSeg2_DATA2 is: "
    + S2.getOut__data2());
// -----
// Populate third segment object from buffer.
// -----
srcPos += LL;
buff = null;
// Get length of segment.
LL =
    MarshallIntegerUtils.unmarshallTwoByteIntegerFromBuffer(
        segBytes,
        srcPos,
        true,
        MarshallIntegerUtils.SIGN_CODING_TWOS_COMPLEMENT);

// Put segment in byte array.
buff = new byte[LL];
System.arraycopy(segBytes, srcPos, buff, dstPos, LL);
remainLen -= LL;

// Create and populate segment object from byte array.
OutputSeg3 S3 = new OutputSeg3();
S3.setBytes(buff);
System.out.println(
    "\nOutSeg3 LL is:      "
    + S3.getOut__ll()
    + "\nOutSeg3 ZZ is:    "
    + S3.getOut__zz()
    + "\nOutSeg3_DATA1 is: "
    + S3.getOut__data1()
    + "\nOutSeg3_DATA2 is: "
    + S3.getOut__data2()
    + "\nOutSeg3_DATA3 is: "
    + S3.getOut__data3());
}
catch (Exception e)
{
    System.out.println("\nCaught exception is: " + e);
}
}
}

```

For variable length messages, multiple segment messages, and messages that contain arrays, see the IMS tutorials and samples in the online help or IBM Documentation in the WebSphere or Rational development environment.

- IMS tutorials are available in **Tutorials > Do and Learn**.
- IMS samples are available in **Samples > Technology samples > J2C samples**.

These tutorials and samples provide guidance information and sample code for handling special message formats.

Related information

[Tutorials in IBM Documentation for Rational Application Developer V9](#) See the IMS tutorials under the Java category in tutorials section in IBM Documentation.

Measuring IMS TM resource adapter performance

To measure IMS TM resource adapter performance, use the interaction execution as the measurement point.

Measuring the time between the opening and closing of the socket does not provide useful data because various nuisances that have nothing to do with the actual transaction response time could be involved. A meaningful transaction response time should be calculated from the start of the send request to the end of the request by setting the start and stop timers around the `IMSInteraction.execute()` method.

The measured data would be similar to what would be measured in IMS Connect HWS trace and other mainframe performance measurement tools.

Securing interactions with the IMS Transaction Manager

You can secure the interactions with IMS Connect and IMS by using the different security features in IMS TM resource adapter.

IMS TM resource adapter security

The Java EE Connector Architecture (JCA) specifies that the application server and the Enterprise Information System (EIS) must collaborate to ensure that only authenticated users are able to access the EIS.

The JCA security architecture extends the end-to-end security model for Java EE-based applications to include integration with EISs. The IMS TM resource adapter follows the Java EE Connector Architecture security architecture, and works with the WebSphere Application Server Java 2 Security Manager.

EIS signon

The JCA security architecture supports the user ID and password authentication mechanism that is specific to an EIS. The user ID and password that are used to sign on to the target EIS are supplied either by the application component (*component-managed signon*) or by the application server (*container-managed signon*).

For the IMS TM resource adapter, IMS is the target EIS. The security information provided by the application component or the application server is passed to the IMS TM resource adapter. IMS TM resource adapter then passes it to IMS Connect. IMS Connect uses this information to perform user authentication, and passes that information to IMS OTMA. IMS OTMA can then use this information to verify authorization to access certain IMS resources.

In a typical environment, the IMS TM resource adapter passes on the security information (user ID, password, and optional group name) that it receives to IMS Connect in an IMS OTMA message. Depending on its security configuration, IMS Connect might then call the Security Authorization Facility (SAF) on the host.

- For WebSphere Application Server on distributed platforms or z/OS with TCP/IP, using either component-managed signon or container-managed signon:
 - If `RACF=Y` is set in the IMS Connect configuration member, or if the IMS Connect command **SETRACF ON** has been issued, IMS Connect calls the SAF to perform authentication using the user ID and password that are passed by the IMS TM resource adapter in the OTMA message. If authentication succeeds, the user ID, optional group name, and `UTOKEN` returned from the IMS Connect call to the SAF are passed to IMS OTMA for verifying authorization to access IMS resources.
 - If `RACF=N` is set in the IMS Connect configuration member, or if the IMS Connect command **SETRACF OFF** has been issued, IMS Connect does not call the SAF. However, the user ID and group name, if specified, are passed to IMS OTMA for authorization to access IMS resources.

- You can provide the user identity to the application server in two ways:
 - The user ID and password can be provided in a Java Authentication and Authorization Service (JAAS) alias. The JAAS alias is associated with either the connection factory that is used by the application that accesses IMS or, depending on the version of WebSphere Application Server, with the EJB resource reference that is used by the application. The application server creates and passes the user token that represents the user identity in the alias to the IMS TM resource adapter.
 - WebSphere Application Server for z/OS can be configured to obtain the user identity that is associated with the thread of execution of the application. You have the option to synchronize a Java thread identity or an operating system thread identity (SyncToOSThread), in which user authentication is performed by the application server. The application server creates and passes the user token that represents this user identity to the IMS TM resource adapter. The IMS TM resource adapter then passes the user token to IMS Connect for authentication. To bypass authentication, modify the HWSJAVA0 exit with the trusted users. IMS Connect then passes the user token to IMS OTMA to verify authorization to access IMS resources.

The level of authorization checking that IMS completes is controlled by the IMS command, **/SECURE OTMA**.

Distributed network security credentials

You can configure WebSphere Application Server or WebSphere Liberty to enable IMS TM resource adapter to pass between Java EE applications and IMS the network security credentials that are entered from an Java EE applications. The network security credentials can include a network session ID and a network user ID. The distributed network security credentials are propagated by the IMS TM resource adapter to IMS Connect, which then passes the credentials to IMS, when a Java EE application accesses IMS transactions.

To enable IMS TM resource adapter to pass network security credentials from a Java EE application to IMS, you must configure and link to your application the Java Authentication and Authorization Service (JAAS) login module that is provided with IMS TM resource adapter. After you link your application to the JAAS login module, users must enter their security credentials when they invoke an IMS transaction for authentication by an external user account registry. The external user account registry can be any user account registry that is supported by WebSphere Application Server or WebSphere Liberty such as an LDAP server. After the credentials are successfully authenticated, IMS TM resource adapter sends the distributed credentials to IMS Connect by using the security-data section of the OTMA message prefix.

You can also enable IMS TM resource adapter to support network security credentials when IMS applications that run in IMS dependent regions make synchronous or asynchronous callout requests to external Java EE applications.

The network session ID that is extracted by the IMS TM resource adapter is, by default, the IP address and port of the authentication server that is used. The maximum length for a network session ID is 254 bytes.

The network user ID that is extracted by the IMS TM resource adapter is one of the distinguished names that are defined in the authentication server. The maximum length for a network user ID is 246 bytes.

Secure Sockets Layer (SSL) Communications

You can configure IMS TM resource adapter and IMS Connect, if properly configured, are able to use the TCP/IP SSL protocol to secure the communications between them.

SSL connections are more secure than non-SSL TCP/IP connections, and provide authentication for the IMS Connect server and, optionally, for the IMS TM resource adapter client. Messages that flow on SSL connections might also be encrypted.

SSL with null encryption provides an intermediate level of security in which the authentication occurs but the messages are not encrypted. Non-encrypted SSL communications offer higher throughput because of the elimination of the overhead that is required to encrypt each message that flows between the IMS TM resource adapter and IMS Connect.

Related concepts

Container-managed EIS signon

With container-managed EIS signon, the security manager in the application server manages the security information for your application.

Secure Sockets Layer (SSL) support

SSL provides security for your interactions by securing the TCP/IP connection between the IMS TM resource adapter and IMS Connect.

Related tasks

Configuring container-managed EIS signon

In WebSphere Application Server Version 6 and later, the container-managed authentication is superseded by the specification of a login configuration on the resource-reference mapping at deployment time.

Related information

[Java thread identity and an operating system thread identity \(WebSphere Application Server\)](#)

Container-managed EIS signon

With container-managed EIS signon, the security manager in the application server manages the security information for your application.

When the directive `<res-auth>Container</res-auth>` is specified in the deployment descriptor of the application, container-managed EIS signon is used. In this case, the application server (the container) provides the security information (user ID and password).

For TCP/IP, the application server passes the security information in the alias to the IMS TM resource adapter. The IMS TM resource adapter passes the security information to IMS Connect for authentication. IMS Connect authenticates the user and passes the security information for signon to IMS. If IMS Connect cannot authenticate the user, a security failure is returned to the IMS TM resource adapter which, in turn, passes an exception back to the application.

With container-managed signon, if your application does pass security information to the IMS TM resource adapter by using the `userName`, `password` or `groupName` properties of the `IMSCConnectionSpec` class, this information is ignored. However, if you pass other information in the `IMSCConnectionSpec` object, such as the client ID with commit mode 0 interactions, this information is used by the IMS TM resource adapter.

Configuring container-managed EIS signon

In WebSphere Application Server Version 6 and later, the container-managed authentication is superseded by the specification of a login configuration on the resource-reference mapping at deployment time.

The `<res-auth>Container</res-auth>` directive in the application deployment descriptor file that you set in your Rational or WebSphere development environment is replaced by the setting of the resource reference that is used by your application.

Configuring container-managed EIS signon consists of:

1. Creating a Java Authentication and Authorization Service (JAAS) authentication data entry alias to be used to provide the user ID and password used for EIS signon.
2. Configuring your application or connection factory to use the appropriate JAAS authentication alias.

To configure for container-managed EIS signon:

1. Create a JAAS - J2C authentication data entry alias and specify the user ID and password.
 - a) In the WebSphere Application Server administrative console, click **Security**.
 - b) Click **Global security**.
 - c) In the global security page, under the Authentication section, click to expand **Java Authentication and Authorization Service** and then click **J2C authentication data**.

- d) In the JAAS - J2C authentication data page, click **New**.
- e) Specify an alias name, and specify the user ID and the password to be used by your application for signon when connecting to IMS.

Global security

[Global security](#) > [JAAS - J2C authentication data](#) > **New**

Specifies a list of user identities and passwords for Java(TM) 2 connector security to use.

General Properties

* Alias
ims

* User ID
usrt001

* Password
.....

Description

Apply **OK** **Reset** **Cancel**

Figure 126. Specification of a security alias to use for J2C authentication

- f) Click **OK**, and click **Save** at the top of the page.
2. Configure your application or connection factory to use the appropriate JAAS authentication alias.
 - To configure an application, in the administrative console:
 - a. Click **Applications** > **Application Types** (for example, **WebSphere enterprise applications**).
 - b. Click the name of your application from the list of applications.
 - c. Under the References section, click **Resource references**.
 - d. Click **Modify Resource Authentication Method**.
 - e. Select the authentication method, and select the module in the table that the authentication method applies to.

Specify authentication method:

☐ None

☒ Use default method (many-to-one mapping)

Authentication data entry
svl-hfungNode01/ims

☐ Use trusted connections (one-to-one mapping)

Authentication data entry
Select...

☐ Use custom login configuration

Application login configuration
Select...

Apply

Figure 127. Specification of the authentication method for your application

Tip: When no authentication is specified, the deprecated container-managed authentication alias on the data source or connection factory is used. Use the online help that is provided to assist you with your configuration.

- f. Click **Apply**, and click **Save** to save the changes.
- To configure a connection factory, in the administrative console:
 - a. Click **Resources > Resource adapters > J2C connection factories**.
 - b. Click the name of the connection factory.
 - c. Specify the component-managed authentication alias, or the container-managed authentication alias.

General Properties

* **Scope**
cells:svl-hfungNode01Cell:nodes:svl-hfungNode01

* **Provider**
IMS TM Resource Adapter

* **Name**
MyIMS

JNDI name
eis/MyIMS

Description

* **Connection factory interface**
javax.resource.cci.ConnectionFactory

Category

Security settings
Select the authentication values for this resource.

Authentication alias for XA recovery
(none)

Component-managed authentication alias
svl-hfungNode01/ims

Mapping-configuration alias
DefaultPrincipalMapping

Container-managed authentication alias
svl-hfungNode01/ims

The additional properties for this item:
Additional Properties

- Connection pool
- Advanced connect
- Custom properties

Related Items

- JAAS - J2C authentication

Figure 128. Specification of the security setting and authentication alias for a J2C connection factory

Tip: The container-managed authentication alias field is used only in the absence of a login configuration on the component resource reference. Use the online help that is provided to assist you with your configuration.

- d. Click **OK**, and click **Save** at the top of the page to save the changes.

Tip: The process for configuring container-managed signon in a stand-alone WebSphere Application Server is the same as the process for a WebSphere Application Server in a unit test environment.

Related information

Configuring Java EE Connector connection factories in the administrative console (WebSphere Application Server Version 8) [For more information, see IBM Documentation for WebSphere Application Server.](#)

J2C connection factories settings (WebSphere Application Server Version 8) [For more information, see IBM Documentation for WebSphere Application Server.](#)

Resource references (WebSphere Application Server Version 8) [For more information, see IBM Documentation for WebSphere Application Server.](#)

Component-managed EIS signon

With component-managed EIS signon, your application manages the security information provided to IMS Connect and IMS. In general, your application supplies the security information to be used for EIS signon.

To specify component-managed EIS signon for your application, type the value `application` for the `<res-auth>` element in the resource reference of the deployment descriptor of your application.

Your application (the component) provides the security information (user ID, password, and optional group name) that is used for EIS signon:

- If your application uses the Java EE Connector Architecture Common Client Interface (CCI), it performs component-managed EIS signon by using the following methods:
 - `IMSConnectionSpec.setUsername()`
 - `IMSConnectionSpec.setPassword()`
 - `IMSConnectionSpec.setGroupName()`

These methods populate an `IMSConnectionSpec` object with security information. After the application establishes a connection to IMS, it passes the `IMSConnectionSpec` object as a parameter of the `IMSConnectionFactory.getConnection` method. The IMS TM resource adapter passes this security information to IMS Connect for use in signing on (authentication and authorization) to IMS.

- If your application is generated by a Rational or WebSphere development environment, the security information is passed as application input data. To pass the security information as input data you must expose the `username`, `password`, and `groupName` properties of the `IMSConnectionSpec` class.

If your application does not use one of the methods to provide security information, WebSphere Application Server obtains the security information from the J2C connection factory custom properties.

Tip: If you specified a component-managed JAAS alias when you set up your connection factory, the user ID and password in the alias override the `username` and `password` values in the connection factory custom properties during the startup of WebSphere Application Server.

Configuring component-managed EIS signon

You specify component-managed EIS signon during application development and configure it during application deployment.

When you create a Java EE application by using a Rational or WebSphere development environment, you are able to select either component-managed or container-managed EIS signon. Setting the authentication directive in the resource reference of your application deployment descriptor to `Application` configures your application for component-managed EIS signon. Similar steps would be used for other resources and other IDEs.

The following example explains how to verify or change this setting for an EJB Project in a Rational or WebSphere development environment:

1. Set the value of the element to **Application**.
 - a. In the Java EE perspective, in the Project Explorer view, expand your **EJB project** in **EJB Projects**.
 - b. Right-click **Deployment Descriptor: your_EJB_project** and select **Open With > Deployment Descriptor Editor**.

- c. In the EJB Deployment Descriptor view, click the **References** tab, expand the name of the EJB component used by your application, and select the resource reference for the EJB. When you select the resource reference for the EJB, fields on the right side of the EJB Deployment Descriptor view are displayed with values.
- d. Select **Application**, if it is not already selected, in the Authentication field. This field maps to the `<res-auth>` element.
- e. Close the EJB Deployment Descriptor Editor and click **Yes** to save your changes. The following code is added to the deployment descriptor of your EJB application:

```
<res-auth>Application</res-auth>
```

Typically, component-managed signon does not require further configuration because the security information is provided by the application in the `IMSSConnectionSpec` object. However, if your application does not provide an `IMSSConnectionSpec` object, or the user ID is not specified in the `IMSSConnectionSpec` object that is provided, the IMS TM resource adapter will obtain default security values from the connection factory that is used by your application.

2. If the application component does not provide a user ID, or the user ID provided is null or blanks, the default security values is used. The default value is specified in the connection factory and can be provided in two ways:
 - Specify a component-managed authentication alias.
 - a. To use a component-managed authentication alias, you must define a JAAS authentication alias.
 - i) In the Servers view, right-click the server and select **Run administrative console**.
 - ii) Expand Resources and select **Resource Adapters**.
 - iii) Click the resource adapter you want to modify.
 - iv) Under Additional Properties, click **J2C Connection factories**.
 - v) Under Related Items, click **J2C Authentication Data Entries**.
 - vi) Above the list of aliases, click **New**.
 - vii) Enter an alias name, your user ID, password, and optional description. Select **OK**.
 - b. Select the JAAS authentication alias for the component-managed authentication alias property of the J2C connection factory that is used by your application. You can select the JAAS authentication alias when you first create the connection factory, or later by editing the connection factory. To edit the connection factory:
 - i) In the administrative console for the server, navigate to the connection factory that you want to modify by selecting **Resource Adapters > server_name > J2C connection factories > connection_factory_name**.
 - ii) In the Component-managed Authentication Alias drop-down list, select the JAAS authentication alias to be used for component-managed authentication by applications that use that connection factory.
 - iii) Select **OK**.

The user ID and password that are associated with the component-managed authentication alias are used to set (or override if applicable) the default values in the custom properties of the associated connection factory during application server startup.
 - Define default values in the connection factory custom properties.
 - If you do not assign a valid JAAS authentication alias to the Component-managed Authentication Alias field of your J2C connection factory, you can assign values for the `userName`, `password`, and `groupName` fields on the Custom Properties page of your J2C connection factory.
 - To create a connection factory, use the `IMSSConnectionSpec` API to specify connection properties. Using a component-managed authentication alias is preferred over specifying values in the custom properties of your J2C connection factory. The component-managed authentication alias

provides greater security for the user ID and password because the user name and password values of a JAAS authentication alias are visible only to server administrators.

The process for configuring component-managed signon in the test environment of an IDE and the process for configuring component-managed signon in a stand-alone WebSphere Application Server are similar.

Related tasks

[Creating a connection factory in WebSphere Application Server](#)

After you install and deploy the IMS TM resource adapter on WebSphere Application Server, create a connection factory for the IMS TM resource adapter, if it does not already exist.

Related reference

[IMS connection factory properties](#)

The properties of an IMS TM resource adapter connection factory describes the characteristics of the target Enterprise Information System.

Secure Sockets Layer (SSL) support

SSL provides security for your interactions by securing the TCP/IP connection between the IMS TM resource adapter and IMS Connect.

The SSL protocol ensures that the transfer of sensitive information over the Internet is secure. SSL protects information from:

- Internet eavesdropping
- Data theft
- Traffic analysis
- Data modification
- Trojan horse browser/server

The IMS TM resource adapter can communicate with IMS Connect through TCP/IP sockets. If the IMS TM resource adapter uses TCP/IP, you can use SSL to secure the TCP/IP communication between the two entities. The SSL support that is provided by the IMS TM resource adapter and IMS Connect uses a combination of public and private keys, along with symmetric key encryption schemes, to achieve client and server authentication, data confidentiality, and integrity. SSL runs in a layer above the TCP/IP communication protocol and allows an SSL-enabled server to authenticate itself to an SSL-enabled client and vice versa.

For an SSL connection between the IMS TM resource adapter and IMS Connect, the IMS TM resource adapter is considered the client and IMS Connect is considered the server. After authentication is complete, the server and client can establish an encrypted connection that also preserves the integrity of the data.

For SSL support in a WebSphere Application Server environment, the IMS TM resource adapter uses the IBM implementation of Java Secure Socket Extension (IBM JSSE). The SSL library is included in WebSphere or Rational development environments and in WebSphere Application Server.

Important:

- The IMS TM resource adapter supports only X.509 certificates and the JKS keystore type on distributed platforms (which include Linux for System z) and the JKS keystore type, or RACF keyrings on z/OS.
- Transport Layer Security, Version 1 (TLS V1) is the successor to SSL 3.0 protocol. The IMS TM resource adapter supports TLS V1, and newer versions. TLS V1 is compatible with the earlier SSL 3.0 protocol.

SSL concepts

The SSL protocol involves key concepts such as certificates, certificate authority, certificate management, key stores, and trust stores.

Certificate

A digital certificate is a digital document that validates the identity of the certificate's owner. A digital certificate contains information about the individual, such as their name, company, and public key. The certificate is signed with a digital signature by the Certificate Authority (CA), which is a trustworthy authority.

Certificate authority

A certificate authority (CA) is a trusted party that creates and issues digital certificates to users and systems. The CA, as a valid credential, establishes the foundation of trust in the certificates.

Certificate management

Certificates and private keys are stored in files called keystores. A keystore is a database of key material. Keystore information can be grouped into two categories: key entries and trusted certificate entries. The two entries can be stored in the same keystore or separately in a keystore and truststore for security purposes. Keystores and truststores are used by both the SSL client, the IMS TM resource adapter, and the SSL server, IMS Connect.

Keystore

A keystore holds key entries, such as the private key of the IMS TM resource adapter, and the SSL client.

Truststore

A truststore is a keystore that holds only certificates that the user trusts. Add an entry to a truststore only if the user makes a decision to trust that entity. An example of an IMS TM resource adapter (client) truststore entry is the certificate of the target SSL server, IMS Connect.

You can store key entries and trusted certificate entries in either the keystore or the truststore. You can also store them separately. The IMS TM resource adapter supports only X.509 certificates and the JKS keystore type on distributed platforms (which include Linux for System z) and the JKS keystore type, or RACF keyrings on z/OS.

SSL protocol

The SSL protocol consists of server authentication and client authentication, followed by an encrypted conversation (SSL handshake).

Server authentication

SSL server authentication allows a client to confirm the identity of a server. SSL-enabled client software uses standard techniques of public-key cryptography to ensure that a server's certificate and public ID is valid, and that the certificate and ID were issued from one of the client's list of trusted certificate authorities (CA).

Client authentication

SSL client authentication allows a server to confirm a client's identity. Using the same techniques used for server authentication, SSL-enabled server software verifies that a client's certificate and public ID is valid and that the certificate and ID was issued by one of the server's list of trusted certificate authorities (CA).

Null Encryption

Null encryption allows for authentication to occur during the SSL handshake. After the SSL handshake completes, all messages flow without being encrypted over that socket.

SSL handshake

Both the client, the IMS TM resource adapter, and the server, IMS Connect, store their certificates and private keys in keystores. The SSL session between the IMS TM resource adapter and IMS Connect is established by following a handshake sequence between the client and the server. The sequence varies,

depending on whether the server is configured to provide just a server certificate, or to provide a server certificate and request a client certificate, and which cipher suites are available for use. A cipher is an encryption algorithm. The SSL protocol determines how the client and the server negotiate the cipher suite to be used, authenticate one another, transmit certificates, establish session keys, and transmit messages. Some of the algorithms used in cipher suites include:

- DES - Data Encryption Standard
- DSA - Digital Signature Algorithm
- KEA - Key Exchange Algorithm
- MD5 - Message Digest algorithm
- RC2 and RC4 - Rivest encryption ciphers
- RSA - A public key algorithm for both encryption and authentication
- RSA key exchange - A key-exchange for SSL based on the RSA algorithm
- SHA-1 - Secure Hash Algorithm
- SKIPJACK - A classified symmetric-key algorithm implemented in FORTEZZA-compliant hardware
- Triple-DES - DES applied three times.

SSL handshake and authentication process

To use SSL, both the client (the IMS TM resource adapter) and the server (IMS Connect) must be configured for SSL handshake.

At run time, when the Java client application executes an interaction with IMS, the interaction flows on a secure (SSL) connection between the SSL client, the IMS TM resource adapter, and the SSL server, IMS Connect. To open an SSL connection between the client and the server, an SSL handshake process occurs. This SSL handshake, which is transparent to the Java client application, occurs as follows:

1. The SSL client, the IMS TM resource adapter, initiates a connection by sending a client hello message. The server, IMS Connect, replies with a server hello message and its certificate that contains its public key.
2. If this certificate is authenticated successfully by the server, a session key is established at both ends and a cipher specification is negotiated that determines the type of encryption to be used on the connection. The cipher can be STRONG, WEAK, or ENULL. The SSL handshake is then completed if the server does not require client authentication.
3. If the server requires client authentication, the client authenticates the server's certificate by using the server's public key from its certificate. If this authentication is successful, a client certificate is sent from the client's keystore. If this certificate is authenticated successfully by the server, a session key is established at both ends, and a cipher specification is negotiated that determines the type of encryption to be used on the connection. The SSL handshake is then completed.
4. The client and server are ready to send and receive encrypted data.

Important: When the client applications run in a managed environment, which is particularly preferred with SSL connections, the IMS TM resource adapter must use persistent socket connections to communicate with IMS Connect. However, in a non-managed environment, these persistent connections are disconnected by the application after each use rather than being made available for reuse by another application.

When the WebSphere Application Server Connection Manager is used, connections can be serially reused by other client applications. The connection manager creates connections if necessary, and provides them to the applications as needed. When an application is finished using a connection, the connection manager returns that connection to the free pool, making it available for reuse by any other application that requires that type of connection. However, client and server authentication only occurs one time for each socket during the handshake that occurs when that socket is first created and initialized as an SSL socket. When a socket is reused, the SSL client, the IMS TM resource adapter, and the server, IMS Connect, do not change. Consequently, it is not necessary to re-authenticate the client and the server

(go through the handshake process again) when a socket is reused. The client ID that identifies a socket remains the same each time a socket is reused.

If the SSLEncryption value is set to ENULL, performance is faster than SSL connections that use Strong or Weak encryption. The level of performance improvement depends on several factors, including whether hardware or software encryption is used. In general, hardware encryption is faster than software encryption.

Configuring the client and the server for SSL support

To use SSL, both the client (the IMS TM resource adapter) and the server (IMS Connect) must be configured.

To configure the IMS TM resource adapter and IMS Connect for SSL:

1. Decide if client authentication is required on the SSL server, IMS Connect. If client authentication is not required, skip to [step 3](#).

Recommendation: Use client authentication to protect against unauthorized access to IMS Connect.

2. When client authentication is required, the client must have a signed certificate in the server's truststore or keyring.
 - a) Obtain signed certificates and a private key for the client.
 - b) On the client, create a keystore and insert the client's private key and certificate.
 - c) On the server (IMS Connect), insert the client's public key certificate into the keyring. See *IMS Version 14 Communications and Connections* for more information.
3. On the client, create a truststore (another optional keystore) and insert the server's public key certificate. Alternatively, insert the public key certificate into the client keystore if trusted and non-trusted certificates are stored in the same keystore.
4. Decide which IMS Connect SSL port to use. Set up the IMS Connect and SSL configuration members with the appropriate values.

For more information about setting up these configuration members, see *IMS Version 14 Communications and Connections*.

5. Set up the connection factory with the appropriate SSL parameters, including the port number from [step 4](#).
6. Bind the application to the SSL connection factory.

Tip: If the SSLEncryption value is set to ENULL, performance is faster than SSL connections that use Strong or Weak encryption. The level of improvement depends on several factors, including whether hardware or software encryption is used. In general, hardware encryption is faster than software encryption.

Related tasks

[Logging and tracing IMS TM resource adapter information](#)

Configure the logging and tracing setting in WebSphere Application Server to ensure messages are logged, and diagnostic trace information is recorded for the IMS TM resource adapter.

Creating the keystore or truststore for the client

On the SSL client, create a keystore to store the client certificate, and create a truststore for the client to store the server certificate.

Several tools are available for managing keystores. The following steps create a client certificate to store in the client keystore and a client truststore to store server certificates.

1. Create a certificate for the client, the IMS TM resource adapter, and have the certificate signed by a certificate authority (for example, VeriSign). You can also create your own certificate authority (CA) by using software such as OpenSSL to sign your own (self-signed) certificate.
2. Create a keystore by using a key management tool such as Ikeyman or Keytool.
3. Import the client certificate (if one is available) into the keystore.

4. Create a truststore for the client by creating another keystore.
5. Import the certificate for the server (IMS Connect) into the client truststore.

Configuring SSL connections

A secure SSL connection between a Java client application and IMS Connect is created by ensuring that the connection factory used by the Java client application has the appropriate values for its SSL properties.

In WebSphere and Rational development environments, you can set up the SSL properties in multiple ways:

- If you are running your Java client application in the test environment inside of your integrated development environment (IDE), use the tooling in the IDE or the administrative console for your test environment. Map or bind the connection factory resource reference in the Java client application which is installed on an integrated WebSphere Application Server, to an SSL-configured connection factory also deployed in that integrated WebSphere Application Server by setting the JNDI name of the connection factory in the resource reference.
- If you are running your Java application in WebSphere Application Server, you can configure a connection factory that creates SSL connections.

To set a connection factory that creates SSL connections in WebSphere Application Server:

1. In the WebSphere Application Server administrative console, navigate to **Resources > Resource Adapters > yourIMSTMResourceAdapter > J2C Connection Factories > yourJ2CConnectionFactory > Custom Properties**.
2. Map the connection factory resource reference in the Java client application to the SSL connection factory by specifying the JNDIName of the target connection factory in **Applications > Enterprise Applications > yourApplication > Map resource references to resources**.

Related reference

[IMS connection factory properties](#)

The properties of an IMS TM resource adapter connection factory describes the characteristics of the target Enterprise Information System.

Changing RACF passwords

When IMS Connect is configured to perform RACF authentication, you can change RACF password by issuing a password change request to IMS Connect from your Java application.

To change the RACF password of the user ID that is specified in your message from your Java application:

Submit a message that includes the password change request keyword HWSPWCH in the following format:

```
HWSPWCH old-password/new-password/new-password
```

The HWSPWCH keyword must be at the beginning of the application data section of the message and be followed by one or more blank spaces, the current password, a forward slash character, the new password, another forward slash and then the new password again.

- The keyword HWSPWCH must be all uppercase characters.
- You must include one or more blank spaces between HWSPWCH and the old password.
- The old password must be valid password for the user ID.
- All passwords must follow RACF standards.

If the password is successfully changed, a "HWSC0031I PASSWORD CHANGE SUCCESSFUL" message is returned. If the operation fails, the IMS TM resource adapter throws an exception ICO0001E with a reason code indicating the error.

Related concepts

[Container-managed EIS signon](#)

With container-managed EIS signon, the security manager in the application server manages the security information for your application.

Component-managed EIS signon

With component-managed EIS signon, your application manages the security information provided to IMS Connect and IMS. In general, your application supplies the security information to be used for EIS signon.

Related reference

Diagnosing problems with callout requests

Callout request errors are often related to the execution timeout setting, incorrect or corrupted callout messages, or network failures.

Related information

Submitting commands to IMS

Although the IMS TM resource adapter is intended primarily for you to run transactions on a host IMS system through a Java EE application, you can also issue IMS commands that are supported by IMS OTMA from your Java applications.

Securing message retrieval from IMS hold queues

You can specify the user ID and password information to ensure that only authorized users are allowed to retrieve asynchronous output or callout request messages from the IMS hold queue.

You can use both the asynchronous output programming model and the callout programming model of the IMS TM resource adapter to retrieve asynchronous output or callout request messages from the hold queue. To ensure that only an authorized user can retrieve an asynchronous output message or a callout request from the hold queue, you can optionally specify the user ID, together with the tpipe name that is contained in the SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT command message. Authorization is performed by IMS OTMA when the message is retrieved from the hold queue.

If security checking is enabled with IMS Connect or IMS OTMA, you must configure the IMS TM resource adapter with the appropriate user ID for the SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT request to retrieve the asynchronous output message or callout request from an IMS asynchronous hold queue. Specify the user ID and password information in the `IMSConnectionSpec` class or in the authentication alias that is used by the `IMSConnectionFactory` or the application's deployment descriptor.

For more information about OTMA security, see the "Securing messages on the asynchronous hold queue" topics in *IMS Version 14 Communications and Connections*.

Related concepts

Asynchronous output programming model

Use this programming model to retrieve output that has been queued by IMS.

Callout programming models

Use this programming model to send outbound messages from an IMS application to request services or data from message-driven beans (MDBs), Enterprise JavaBeans (EJB) components, or web services and, optionally, receive response data.

Java exceptions that involve output messages

If any errors occur when the IMS TM resource adapter passes the message through IMS Connect to IMS, and as IMS processes this transaction and attempts to return the output, the Java client receives an exception.

Enabling support for distributed network security credentials

IMS TM resource adapter can pass the original, distributed network security credentials, including the network session ID and the network user ID, between Java EE applications and IMS. The network security credentials that are passed by IMS TM resource adapter are written to IMS log records.

The network session ID that can be passed by the IMS TM resource adapter is, by default, the IP address and port of the authentication server that is used. The maximum length for a network session ID is 254 bytes.

The network user ID that can be passed by the IMS TM resource adapter is a distinguished name that is defined in the authentication server. The maximum length for a network user ID is 246 bytes.

Provided with IMS TM resource adapter to support network security credentials is an extendable Java Authentication and Authorization Service (JAAS) login module. You must install the JAAS login module in the application server, which can be either WebSphere Application Server or WebSphere Liberty, and link the login module to a Java EE application. The following process describes how network security credentials are passed by the JAAS login module and IMS TM resource adapter after the login module is installed and configured:

1. The JAAS login module prompts the user to enter network security credentials and captures the credentials.
2. The network security credentials are passed by WebSphere Application Server or WebSphere Liberty to an external user account registry, such as an LDAP server, for authentication.
3. If the security credentials are successfully authenticated by the user account registry, the credentials are set in a Java principal object and passed to IMS TM resource adapter.
4. IMS TM resource adapter extracts the network user ID and the network session ID and includes the security credentials in the security-data section of the OTMA message prefix. The OTMA message is then sent to IMS Connect with the network security credentials.
5. IMS Connect passes the OTMA message with network security credentials to IMS, and the credentials are written to IMS log records.

You can also enable IMS TM resource adapter to support network security credentials from IMS in synchronous and asynchronous callout messages.

Configuring WebSphere Application Server for distributed network security credentials

You can enable IMS to audit the network security credentials that are passed to and from a Java EE application that runs on WebSphere Application Server.

Before you enable support for distributed network security credentials that are passed from your Java EE application, ensure that the following items are enabled:

- Container-managed security for your application. If your application does not have container-managed security, you might have to modify your application's `web.xml` file.
- An external user account registry, such as an LDAP server, that contains authorized users.

The following procedure makes the following assumptions:

- That you use an LDAP server as the external user registry. However, you can use any other user registry that is supported by WebSphere Application Server.
- That you use `IMS_Login` as the alias name for the login module that you create.
- That a login module is linked to a simple `imsicoivp.ear` application, which sends a `/STA OTMA` command.

1. Set up the user registry.

- a) On the left menu, click **Security** and then **Global Security**. Then, in the **User account repository** area, select **Standalone LDAP Registry** from the **Available realm definitions** list. Click **Configure**.
- b) In the configuration menu for your LDAP server, specify the primary administrative user name. This user name must be defined in the LDAP server.

The administrative user name consists of a unique user ID, an organization unit, and one or more domains. In the following example administrative user name, the unique user ID is admin, the organization unit is users, and the domains are security and com:

```
uid=admin,ou=users,dc=security,dc=com
```

- c) Specify the host and port address of your LDAP server.
- d) Specify the base distinguished name of your LDAP server.

The base distinguished name is the directory that holds authorized users. For example, the following base distinguished name might be used:

```
ou=users,dc=security,dc=com
```

If the preceding base distinguished name is used, the following example users are authorized :

```
uid=admin,ou=users,dc=security,dc=com  
cn=Bob,ou=users,dc=security,dc=com
```

- e) Click **Test connection** to test the connection to the LDAP server. Then, click **Apply**.
- f) On the **Global Security** page, select **Standalone LDAP registry** from the **Available realm definitions** list, and then click **Set as current**. Then, click **Apply**.

Your user account registry is set up. The login module uses this user account registry to authenticate distributed network security credentials that are entered.

2. Enable security for WebSphere Application Server and install the custom JAAS login module.

- a) Ensure that your server has Administrative and Application Security enabled. On the left side of the WebSphere Application Server window, click **Security** > **Global Security**, and then select **Enable administrative security** and **Enable application security**.
- b) In the right panel, click **Authentication** and expand **Java Authentication and Authorization Service**. Click **Application logins** and then click **New**.
- c) Create an alias name for the login module. The following steps in this procedure assume that you create an alias name of IMS_Login. Then, in the **JAAS login modules** table, click **New**.
- d) Specify the module class name `com.ibm.ims.login.IMSLoginModule`. This is the login module that is bundled with the IMS TM resource adapter.
- e) From the **Authentication strategy** list, select **REQUIRED**.
- f) In the **Custom properties** section, add a new custom property. Enter `propIdentity` as the name and `Caller` as the value.
- g) If you have a web application and the URI is unprotected, you must allow WebSphere Application Server to prompt for authentication with unprotected URIs. To enable WebSphere Application Server to prompt for authentication, on the **Global Security** page, click **Web and SIP Security**. Select **General Settings**.
- h) In the **Web authentication behavior** area, select **Authenticate when any URI is accessed**.

3. Map the login module to your application.

- a) On the left menu, click **Applications**. Then, click **Application Types** and click the `imsicoivp.ear` application.
- b) Click **Resource References**, select your application, and click **Modify Resource Authentication Method**.
- c) Select **Use custom login configuration**, and then select **IMS_Login** from the **Application login configuration** list. Click **Apply** > **OK**.

After you map the login module to your application, you are prompted to enter a user name and password when you start and open your application in your web browser. Enter your distinguished name as the user name. The following example of a log file snippet shows that the IMS TM resource adapter successfully received the distributed network security credentials:

```
*****
TMRA has received the following credentials:
Security Realm: '0.0.0.0:10389'
Distinguished Name: 'uid=admin,ou=users,dc=security,dc=com'
Authenticated?: 'true'
*****
```

If the credentials are valid and exist in the LDAP registry, the distinguished name and realm are propagated to IMS.

Related concepts

[“Container-managed EIS signon” on page 799](#)

With container-managed EIS signon, the security manager in the application server manages the security information for your application.

Configuring WebSphere Liberty for distributed network security credentials

You can enable IMS to audit the network security credentials that are passed to and from a Java EE application that runs on WebSphere Liberty.

Before you enable support for distributed network security credentials that are passed from your Java EE application, ensure that the following items are enabled:

- Container-managed security for your application.
- An external user account registry, such as an LDAP server, that contains authorized users.

This procedure assumes that you are using an LDAP server as the external user registry. However, you can use any other user registry that is supported by WebSphere Liberty.

1. If the distributed network security credentials are sent from a web application, ensure that the application URI is protected because WebSphere Liberty does not authenticate unprotected URIs. To protect a web application URI, define security roles and authorizations in the `web.xml` file of your web application.

In the following example `web.xml` file, a security role that is called `Testing` is defined, and in the `<security-constraint>` tag, `Testing` is the only role that is authorized to access that application. Any user that wants to access the application must be assigned the `Testing` role.

```
<security-constraint>
<display-name>SampleWebApp</display-name>
  <web-resource-collection>
    <web-resource-name>SampleWebApp</web-resource-name>
    <url-pattern>*/</url-pattern>
  </web-resource-collection>
  <auth-constraint>
    <role-name>Testing</role-name>
  </auth-constraint>
</security-constraint>

<security-role>
  <role-name>Testing</role-name>
</security-role>
```

2. Set up a user registry for your application by creating the registry element in the `server.xml` file.

In the registry tag, specify the following items:

- The LDAP type, port, and host address
- The baseDN attribute for the base distinguished name
- The `<customFilters>` element for custom filters

In this example, the `<ldapRegistry>` element is used because an LDAP server is used as the user registry.


```
<ldapRegistry ldapType="Custom" port="9999" host="0.0.0.0"
  baseDN="ou=users,dc=test,dc=test">
  <customFilters userFilter="(&!(uid=%v)(objectclass=inetOrgPerson))"
    groupFilter="(&!(cn=%v)(!(objectclass=groupOfNames)(objectclass=groupOfUniqueNames)))"
    groupMemberIdMap="ibm-allGroups:member;ibm-allGroups:uniqueMember;groupOfNames:member;groupOfUniqueNames:uniqueMember"></customFilters>
</ldapRegistry>
```

3. Configure the custom login module in the server.xml file.

- Create a `<library>` element that uses a `<fileset>` element, which indicates where the `IMSLogin.jar` file is. The `IMSLogin.jar` file is packaged in the IMS TM resource adapter .rar file. However, WebSphere Liberty cannot scan .rar files, so you must extract the .rar file to a directory that can be scanned by WebSphere Liberty. In this example, the **id** is `customLoginLib`.
- Create the `<jaasLoginModule>` element and set the following attributes:
 - The **id** attribute. In this example, the **id** is `IMS_Login`.
 - On the **className** attribute, specify the `com.ibm.ims.login.IMSLoginModule` class.
 - On the **controlFlag** attribute, specify `REQUIRED` to require a successful authentication of the distributed network security credentials.
 - On the `<libraryRef>` tag, specify `customLoginLib`, the ID of the `<library>` element configured in the previous substep.
- Create a `<jaasLoginContextEntry>` element and specify an ID and a unique name of the system-defined JAAS configuration. In the following example, `IMS_Login_Entry` is used as the name and ID. On the **loginModuleRef** attribute, specify `IMS_Login`, the ID of the `<jaasLoginModule>` element that you created in the previous substep.

```
<jaasLoginContextEntry id="IMS_Login_Entry"
  loginModuleRef="IMS_Login" name="IMS_Login_Entry" />

<jaasLoginModule className="com.ibm.ims.login.IMSLoginModule"
  controlFlag="REQUIRED" id="IMS_Login" libraryRef="customLoginLib">
  <options propIdentity="RunAs" />
</jaasLoginModule>

<library id="customLoginLib">
  <fileset dir="${server.config.dir}/libraries" includes="IMSLogin.jar" />
</library>
```

The server configuration is now set up to use your login module.

- Link your application to the custom JAAS login module. In the web application bind file, manually link the application to the login module by using the `loginConfigurationName` attribute of the `<resRefBindings>` element. For the value of the `loginConfigurationName` attribute, specify the ID of the `<jaasLoginContextEntry>` element that you defined in the previous step.

```
<?xml version="1.0" encoding="UTF-8"?>
<com.ibm.ejs.models.base.bindings.webappbnd:WebAppBinding xmi:version="2.0" xmlns
  <webapp href="WEB-INF/web.xml#WebApp"/>

  <resRefBindings xmi:id="ResourceRefBinding_1264197557812" jndiName="ims_cf2"
    loginConfigurationName="IMS_Login_Entry">
    <bindingResourceRef href="WEB-INF/web.xml#ResourceRef_1264197557812"/>
  </resRefBindings>

</com.ibm.ejs.models.base.bindings.webappbnd:WebAppBinding>
```

Related concepts

[“Container-managed EIS signon” on page 799](#)

With container-managed EIS signon, the security manager in the application server manages the security information for your application.

Enabling IMS TM resource adapter client applications to support network security credentials in callout messages

If you enable IMS TM resource adapter client applications to support network security credentials in IMS callout messages, external web servers can use the credentials to perform authentication and authorization for callout requests.

Network security credentials in synchronous callout messages

To enable IMS TM resource adapter client applications to support network security credentials in synchronous callout messages, set the `resumeTpPipeNsc` property of the `IMSActivationSpec` object to `true`.

Network security credentials in asynchronous callout messages

To enable IMS TM resource adapter client applications to support network security credentials in asynchronous callout messages, you must call the `setResumeTpPipeNSC(int resumeTpPipeNSC)` method for the `IMSInteractionSpec` object and set the value of the `setResumeTpPipeNSC` property to 1. If 1 is set for the `setResumeTpPipeNSC` property, IMS TM resource adapter sets a flag byte in the OTMA message prefix that is sent to IMS to indicate that network security credentials should be included in the callout message.

Related reference

[IMSActivationSpec property configuration for message-driven beans](#)

The values of the properties of the `IMSActivationSpec` object describe the inbound communication from IMS to be used by message-driven beans.

[Resume tp pipe network security credentials \(resumeTpPipeNSC\)](#)

This property is used by a Java application to indicate whether network security credentials are to be included in asynchronous callout requests from IMS.

IMS TM resource adapter timeouts

Several types of timeouts can be used to prevent applications from hanging due to failures during execution of an interaction.

Sometimes network problems affect the transmission of data between a client and IMS. Without a way to handle such unexpected problems, applications can stop or run in an endless loop. One such example is router failure on the TCP/IP path between the IMS TM resource adapter and IMS Connect.

Several timeouts can be used to recover from such a failure.

- The `executionTimeout` property of the `IMSInteractionSpec` class is used to deal with problems between IMS Connect and IMS. Examples of such problems include XCF communications link failure, or a failure in IMS that leads to the target transaction not running or its output not being returned to IMS Connect.
- The `socketTimeout` property of the `IMSInteractionSpec` class is used to address problems affecting TCP/IP communication between the IMS TM resource adapter and IMS Connect.

You can specify values for the `executionTimeout` and `socketTimeout` properties to handle problems that occur during execution of an interaction. The timer for execution timeout runs inside the socket timeout window.

Recommendation: Always set the `socketTimeout` property to a value that is larger than the `executionTimeout` property value.

Execution timeouts

The value for the execution timeout is the maximum amount of time that is allowed for IMS Connect to send a message to IMS and receive a response to that message back from IMS.

The execution timeout setting is used an IMS TM resource adapter client application to instruct IMS Connect to override its current TIMEOUT value.

If an interaction does not complete before execution timeout occurs, IMS Connect returns an error message to the IMS TM resource adapter. The IMS TM resource adapter then returns an exception to the client application, indicating that the duration of time for IMS to respond to IMS Connect has exceeded the timeout value. The error message also specifies the timeout value that was used by IMS Connect.

Tip: Because connections between the IMS TM resource adapter and IMS Connect are persistent, if a connection is in a known state after an execution timeout occurs, the socket is not closed. Instead, the socket remains open and is made available for reuse.

Transaction expiration

You can take advantage of the OTMA transaction expiration enhancement to explicitly instruct IMS Connect to indicate to OTMA to discard or dequeue a transaction when the execution timeout value is reached.

By default, the `transExpiration` property is set to `false` for backward compatibility, and OTMA continues to process the transaction even after the execution times out. To take advantage of the OTMA transaction expiration function to save unnecessary processing costs and CPU cycles for transactions that are no longer needed, set the `transExpiration` property to `true`.

Conversational transactions

For conversational transactions, the execution timeout value applies to each iteration of a conversation. An *iteration* consists of one input message that is sent to IMS and one output message that is received from IMS. If an iteration of the conversation times out due to an execution timeout, the conversation ends and any database updates that have taken place in that conversation are backed out.

Exceptions

If you specify an invalid execution timeout value, the TIMEOUT value that is specified in the IMS Connect configuration member is used and a `javax.resource.NotSupportedException` is thrown.

Related concepts

Other types of timeouts

In addition to execution timeouts and socket timeouts, other types of timeouts, such as those for the J2C connection factory, the Enterprise JavaBeans (EJB) transaction, browsers, the HTTP session, and the EJB session, can impact your application's execution.

Valid execution timeout values

The execution timeout value is represented in milliseconds and must be a decimal integer from 1 through 3600000.

That is, the execution timeout value must be greater than zero and less than or equal to one hour. The execution timeout value can also be -1 if you want an interaction to run without a time limit. The execution timeout value cannot contain non-numeric characters.

If you do not specify an execution timeout value or if the value that you specify is invalid:

- For `SYNC_SEND_RECEIVE` interactions, the timeout value in the IMS Connect configuration member is used and the interaction continues to run.
- For `SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_NOWAIT`, and `SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_WAIT` interactions, IMS Connect sets the timeout value to two seconds and the interaction continues to run.

If an execution timeout is specified for a send-only interaction it is ignored as execution timeout does not apply to send-only interactions.

If you specify an invalid value and a timeout occurs, the timeout value that is specified in the IMS Connect configuration member is used, and the exception `javax.resource.NotSupportedException` is thrown.

Tip: The host system administrator determines the global timeout value in the IMS Connect configuration member. To display this value, issue the **VIEWHWS** command on the z/OS console.

If a valid execution timeout value is set, this value is converted to a value that IMS Connect can use. The following table describes how the values that you specify are converted to the values that IMS Connect uses:

Range of user-specified values	Conversion rule
1 - 250	If the user-specified value is not divisible by 10, it is converted to the next greater increment of 10.
251 - 1000	If the user-specified value is not divisible by 50, it is converted to the next greater increment of 50.
1001 - 60000	The user-specified value is converted to the nearest increment of 1000. Values that are exactly between increments of 1000 are converted to the next greater increment of 1000.
60001 - 3600000	The user-specified value is converted to the nearest increment of 60000. Values that are exactly between increments of 60000 are converted to the next greater increment of 60000.

For example, if you specify a value of 1, this value is converted to 10 (because 1 is not divisible by 10 and 10 is the next increment that is greater than 1). The following examples illustrate how the conversion works for each range of values:

User-specified value (milliseconds)	Converted value (milliseconds)
1	10
11	20
251	300
401	450
1499	1000
1500	2000
60000	60000
89999	60000
3600000	3600000
3750000	3600000

Related concepts

[Other types of timeouts](#)

In addition to execution timeouts and socket timeouts, other types of timeouts, such as those for the J2C connection factory, the Enterprise JavaBeans (EJB) transaction, browsers, the HTTP session, and the EJB session, can impact your application's execution.

Related reference

Execution timeout (`executionTimeout`)

This property specifies the maximum amount of time that is allowed for IMS Connect to send a message to IMS and receive a response

Setting execution timeout values

You can set the execution timeout value in the J2C wizard in your Rational or WebSphere integrated development environment (IDE) or by using the Common Client Interface (CCI) API.

The execution timeout property is a property of the `IMSInteractionSpec` class. The execution timeout value that you set is converted to a value that IMS Connect uses. This conversion occurs to meet the requirements of IMS Connect.

Other timeouts, such as socket timeouts, or the interaction between your client application and WebSphere Application Server for z/OS, can affect your interactions. If other timeout values are less than the execution timeout value you set for your IMS interaction, these other timeouts can mask the fact that IMS did not return a response.

In an IDE, you can set the execution timeout value when you initially define the IMS binding properties for a new J2C Java bean.

To edit the IMS binding properties that are already defined in the IDE for a new J2C Java bean, complete the following steps:

1. Open the appropriate Java binding for IMS by using the Java Editor.
2. Locate the doclet tag for the `IMSInteractionSpec` class.
3. Modify the doclet tag to add the `executionTimeout` property, if it is not listed, and specify a value for the `executionTimeout` property. If the property is listed, modify the value.
4. Close the editor and click **Yes** to save your changes.

You can also code the individual timeout values for different interactions by setting the `IMSInteractionSpec` `executionTimeout` property value. If you code an execution timeout value in your Java client application code, that value overrides any execution timeout value that you set in the IMS binding properties of your J2C Java bean.

With the second approach, use the `setExecutionTimeout` method to set an execution timeout value in an CCI application.

1. First instantiate a new `IMSInteractionSpec` instance or obtain the `IMSInteractionSpec` instance from your specific interaction.
2. Set the `executionTimeout` value by using the `setExecutionTimeout` method. For example:

```
interactionSpec.setExecutionTimeout(timeoutValue);
```

3. Assign this `interactionSpec` to the specific interaction.

Socket timeouts

The value for the execution timeout is the maximum amount of time the IMS TM resource adapter waits for a response from IMS Connect before disconnecting the socket and returning an exception to the client application.

If network problems or routing failures occur, the socket timeout value that you specify can prevent the client application or the IMS TM resource adapter from waiting indefinitely for a response from IMS Connect. Because the `socketTimeout` property is based on the TCP/IP sockets that IMS Connect and the IMS TM resource adapter use to communicate, the `socketTimeout` property does not apply to Local Option connections.

With the `socketTimeout` property, you can set individual timeout values for a particular interaction by using a socket. The value, in milliseconds, can be set on the `socketTimeout` property in the `IMSInteractionSpec` class. If the `socketTimeout` property is not specified for an interaction or it is set to zero milliseconds, no socket timeout value exists, and the connection waits indefinitely.

When determining the socket timeout value, you must consider other existing timeout values.

If a valid socket timeout value is specified for a particular interaction and socket timeout occurs, a `java.io.IOException` and the Java EE J2C exception, `javax.resource.spi.CommException` are thrown. The Java EE J2C exception message indicates that the client has spent more time than was allocated by the `socketTimeout` value to communicate with IMS Connect.

Related concepts

Other types of timeouts

In addition to execution timeouts and socket timeouts, other types of timeouts, such as those for the J2C connection factory, the Enterprise JavaBeans (EJB) transaction, browsers, the HTTP session, and the EJB session, can impact your application's execution.

Related reference

Socket timeout (`socketTimeout`)

The maximum amount of time the IMS TM resource adapter waits for a response from IMS Connect before disconnecting the socket and returning an exception to the client application.

Setting socket timeout values

Set the socket timeout value to a number greater than the execution timeout value.

The `executionTimeout` property is for setting the maximum amount of time that is allowed for IMS Connect to send a message to IMS and receive a response from IMS. Because the socket timeout value encapsulates the execution timeout value, the socket timeout value must be greater than the execution timeout value. A socket might time out unnecessarily if its value is set to less than the execution timeout value.

The following table lists suggested values for socket timeouts based on execution timeout values.

Table 120. Suggested values for socket timeouts based on execution timeout values

Execution timeout value (milliseconds)	Execution timeout behavior	Suggested socket timeout value
0 (or no value)	The default value from the IMS Connect configuration member is used.	The socket timeout value must be greater than the execution timeout default value specified in the IMS Connect configuration file.
1 - 3600000	The wait response times out after the specified millisecond value.	The socket timeout value must be greater than the execution timeout value.
-1	The wait response is indefinite.	Set the socket timeout value to 0 so that the connection waits indefinitely.

You can set the socket timeout value in two ways. You can use the Common Client Interface (CCI) to access the getter and setter methods provided with the `IMSInteractionSpec` class, or use the tooling that is provided by a WebSphere or Rational development environment, with the optional Java EE Connector (J2C) feature installed.

Setting the socket timeout value by using the Common Client Interface

Use the `setSocketTimeout` method of the `IMSInteractionSpec` object to set the socket timeout value.

To use the `setSocketTimeout` method to set the socket timeout value:

1. Instantiate a new `IMSInteractionSpec` object or obtain the `IMSInteractionSpec` object from your specific interaction.
2. set the socket timeout value for `IMSInteractionSpec` by using the `setSocketTimeout` method.
For example:

```
interactionSpec.setSocketTimeout(timeoutValue1);  
interaction.execute(interactionSpec,input,output);  
  
interactionSpec.setSocketTimeout(timeoutValue2);  
interaction.execute(interactionSpec,input,output);
```

Setting the socket timeout value by using a development environment

You can use a Rational or WebSphere development environment to set the socket timeout value when you initially define the IMS binding properties for a new J2C Java bean.

To edit the operation binding properties that are already defined for a new J2C Java bean:

1. Open the appropriate IMS binding Java file by using the Java Editor.
2. Locate the doclet tag for the `IMSInteractionSpec` class.
3. Modify the doclet tag to add the `socketTimeout` property and specify a value. If this property is already listed, modify the value.
4. Select the operation extensibility element again to indicate that changes have been made.
5. Close the editor and click **Yes** to save your changes.

Other types of timeouts

In addition to execution timeouts and socket timeouts, other types of timeouts, such as those for the J2C connection factory, the Enterprise JavaBeans (EJB) transaction, browsers, the HTTP session, and the EJB session, can impact your application's execution.

When you develop your application, you must consider how other types of timeouts might impact the execution of your application.

Timeouts overlap or even encapsulate one another. Such capability provides multiple layers of protection against problems that can cause your application to hang. A longer timeout might mask a shorter timeout. For example, if the socket timeout is set to a lower value than the execution timeout, the socket timeout error would preclude an execution timeout and the absence of a response from IMS would remain unknown.

Other timeout values that might interact with the IMS TM resource adapter timeouts include the following values:

- Connection timeout property of J2C connection factories
- Enterprise JavaBeans (EJB) transaction timeout
- Browser timeout
- Servlet HTTP session or EJB session timeout

One example of a timeout that might impact the execution of your application is the interaction between the applications that run on WebSphere Application Server for z/OS and the IMS TM resource adapter. The WebSphere Application Server for z/OS consists of two parts: a controller and a set of one or more servants. By default, application work is timed, even when it is dispatched to a servant region. When an application that is dispatched to a servant region reaches its timeout, that servant region usually is abended and restarted. The server stays up and continues taking work. For this reason, use care when you choose execution timeout values that are greater than WebSphere Application Server timeout values.

Also use care when you choose the execution timeout value of -1, which tells IMS Connect to wait indefinitely for a response from IMS.

If you disable WebSphere Application Server timeouts, refer to the server documentation for information regarding additional timeout values that are not described in this information.

A second example of a timeout value that would impact the execution of your application is the browser timeout. If the execution timeout value is configured to be greater than the browser timeout, the execution timeout value is not used because the browser timeout occurs before the execution timeout.

Related information

IBM Documentation for WebSphere Application Server Version 8 For more information about server timeouts, see IBM Documentation for WebSphere Application Server.

Conversational programs

An IMS program can support a transaction that is composed of one interaction or multiple interactions. A transaction with multiple iterations is referred to as a *conversational transaction*.

An IMS conversational program is intended to process more complex transactions that take multiple interactions to complete. Each interaction is known as an iteration of the conversation. An IMS conversational program divides the complex processing into a connected series of client-to-program-to-client interactions that are known as *iterations*. In each iteration of a conversation, the IMS conversational transaction program receives a request message from the client, processes that request, and sends a reply to the client. The program also saves any intermediate data from the transaction in the scratchpad area (SPA) to be used in the next iteration of the conversation. The user can modify the data in the reply, or enter additional data and then send in the request to IMS to process the next iteration of the conversation. Using the data in this new request along with any data in the SPA from the last iteration, the IMS conversational program will process the next iteration of the conversation. The processing of these iterations continues until the conversation is ended either by the client application or by IMS.

With IMS TM resource adapter conversational support, a client that interacts with the IMS conversational transaction usually uses one of the following applications:

- Java applications or Web applications
- Business process applications that are based on service component architecture (SCA) in a service-oriented architecture (SOA)

Java applications or Web applications

With Java or Web applications, a user uses the same browser session to conduct a series of requests that iterate through different interactions of an IMS conversation. One possible implementation involves a Java servlet receiving an input request from the browser and using IMS TM resource adapter classes to send a conversational transaction request to IMS Connect through a TCP/IP communication. IMS Connect then forwards the transaction request to both OTMA and IMS and schedules the IMS conversational transaction to run; this process creates a new IMS conversation. The IMS application processes the request and sends the output back to the Java servlet through the conversational support in IMS Connect and in the IMS TM resource adapter. The Java servlet loads the appropriate JavaServer Pages (JSP) file to display the output to the user in the browser.

Business process applications

A business process is an orchestration of a set of services to fulfill a particular business goal. Business process choreography is a key SOA implementation to ensure that business applications are composed of flexible and adaptable elements. IBM WebSphere Process Server provides the runtime infrastructure to execute business processes, and IBM WebSphere Integration Developer provides the modeling tool to author business processes and various other components.

Client-managed and IMS Connect-managed conversation state programming models

An IMS program can support a transaction that is composed of one interaction or multiple interactions. A transaction with multiple iterations is referred to as a conversational transaction.

You can enable your applications for business process choreography by indicating in your Java application that the conversation state will be managed by your application rather than by IMS Connect. By managing a unique conversation ID across iterations of a conversation, you allow different iterations to come in from any connection. This model is known as the *client-managed conversation state programming model*.

By default, the conversation state is still managed by IMS Connect to provide backward compatibility. This model is known as the *IMS Connect-managed conversation state programming model*.

To enable your applications for business process choreography or for use on shareable persistent socket connections, you must manage the conversation state in your client application. In your application, you indicate to IMS Connect that you need IMS to assign a conversation ID for you to pass among the iterations of the conversation by setting the `useConvID` property of the `IMSInteractionSpec` class to `true`. This client-managed conversation state programming model ensures that the different iterations of a conversation can come in from any connections without being rejected.

Recommendation: Use this client-managed conversation state programming model for all new application development for better conversation state management.

Related tasks

[Using the client-managed conversation state programming model](#)

For the client application to manage all iterations in a conversation, the `useConvID` property of the `IMSInteractionSpec` class must be set to `true`.

[Using the IMS Connect-managed conversation state programming model](#)

To specify to have IMS Connect manage all iterations in a conversation, the `useConvID` property of the `IMSInteractionSpec` class must be set to `false`.

Orphaned IMS conversation

When a conversation is not ended explicitly, it continues to exist in the system as an orphaned conversation, and the associated IMS storage continues to be allocated to that conversation.

An IMS conversation is usually ended explicitly in one of two ways:

- The IMS application inserts blanks to the SPA before returning a response to the client.
- The client application program submits a `SYNC_END_CONVERSATION` request.

If the browser was closed before the conversation ends properly, for example, the IMS conversation is not ended explicitly and continues to exist in the system. When an IMS conversation becomes orphaned, you have no way of programmatically continuing or ending that conversation. One measure that you can take to avoid orphaned conversations would be to use timeouts, such as the EJB session timeout, to force the end of a conversation that does not complete in a reasonable amount of time by submitting a `SYNC_END_CONVERSATION` request in the EJB session timeout cleanup code.

If a client application is terminated and a conversation becomes orphaned, the orphaned IMS conversation can be ended only by an IMS restart. You can check for orphaned IMS conversations in the system by issuing an `IMS /DISPLAY CONV` command through an `IMS_REQUEST_TYPE_IMS_COMMAND` interaction. For a list of IMS commands supported by OTMA, see the "Commands Supported from LU 6.2 Devices and OTMA" section in *IMS Version 14 Commands, Volume 2* or in "IMS Commands using OTMA" in *IMS Version 14 Communications and Connections*.

Business process choreography applications

You can run IMS conversational transactions in composite business applications by using IMS TM Resource Adapter.

To allow for business flexibility in a service-oriented architecture (SOA), a key implementation is *business process choreography*. Business process choreography ensures that business applications are made up of flexible and adaptable elements, so you can rapidly change the applications to meet business demands. Instead of having one huge application that cannot be easily modified for new business processes or needs, you compose your business applications based on your business process model. With IMS TM Resource Adapter, you can reuse your IMS conversational transactions in composite business applications that are served by IBM WebSphere Process Server.

To invoke an IMS conversational application through business process choreography:

- You must use the client-managed conversation state programming model.

You must indicate to IMS Connect that this conversation needs a unique conversation ID to keep track of the conversation by setting the `IMSInteractionSpec` class `useConvID` property to `true` in your Service Component Architecture (SCA) component for the business process choreography application.

This setting triggers IMS to assign a unique conversational token and returns it to IMS Connect in the output message of the first iteration. This conversational token can then be passed back and forth between the business process choreography application and IMS Connect, and onto OTMA.

- You must use the dynamic interaction specification (`InteractionSpec`) support that is provided by WebSphere Process Server when you use the SCA enterprise information system (EIS) bindings. The SCA component uses this support to propagate the conversation token to the next component in all subsequent iterations of a conversation or in the `SYNC_END_CONVERSATION` iteration to end a conversation.
- Specify commit mode 1 (CM1) with sync level of NONE or CONFIRM in the `IMSInteractionSpec` object.
- All iterations of a given IMS conversation must be processed by the same IMS Connect and the same IMS in a non-sysplex environment. This unique conversation ID is not shared between multiple IMS systems.
- All iterations of a given IMS conversation must use the same conversational ID, port number, IMS Connect, and data store.

For more information about the conversational support provided by IMS Connect and OTMA and related restrictions, see the topic "IMS Connect conversational support" in *IMS Version 14 Communications and Connections*.

For more information about the dynamic interaction specification support provided in WebSphere Process Server, see the JCA EIS binding-related topics in the WebSphere Process Server documentation.

Related information

[WebSphere Process Server documentation \(all versions\)](#) For more information about JCA EIS binding, see IBM Documentation for the version of WebSphere Process Server that you use.

Enabling your Java client for IMS conversational transactions

You can enable your Java client for IMS conversational transactions by using either the client-managed or the IMS Connect-managed conversation state programming model.

Related concepts

[Conversational programs](#)

An IMS program can support a transaction that is composed of one interaction or multiple interactions. A transaction with multiple iterations is referred to as a *conversational transaction*.

Using the client-managed conversation state programming model

For the client application to manage all iterations in a conversation, the `useConvID` property of the `IMSInteractionSpec` class must be set to `true`.

To enable your Java client for IMS conversational transactions by using the client-managed conversation state programming model:

1. Obtain a connection handle from the `IMSConnectionFactory` object. In this programming model, you can use the same or different connection handles in subsequent iterations. These connection handles must come from the same `IMSConnectionFactory` instance.
2. Set the `IMSInteractionSpec` `useConvID` property to `true`. This property needs to remain true for all iterations of the conversation. If the `useConvID` flag changes during a conversation, an error would occur.
3. Set the `IMSInteractionSpec` `commitMode` property to 1.
4. Set an appropriate value for the `syncLevel` property based on your application design.
5. On the first iteration, set the `IMSInteractionSpec` `convID` property value to an empty string.
6. On all iterations other than the first iteration, set the value of the `convID` property to the value of the `convID` property that is returned in the previous iteration.
7. If the client application controls the ending of a conversation, on the last iteration, set the `IMSInteractionSpec` `interactionVerb` property to `SYNC_END_CONVERSATION`.

Recommendation: With the client-managed conversation state programming model, use local references per request in your client application or servlet instead of application-wide references. This approach is to avoid any race conditions that might occur in the client application due to stress or load that could violate the integrity of the conversational data. If the client application is browser-based, you need to carefully design the navigation so that the conversational request does not get resubmitted accidentally. This scenario can happen if the user presses the back button on the browser, which causes the conversational request to be resubmitted again and therefore results in an additional conversation.

Related concepts

[Client-managed and IMS Connect-managed conversation state programming models](#)

An IMS program can support a transaction that is composed of one interaction or multiple interactions. A transaction with multiple iterations is referred to as a conversational transaction.

Using the IMS Connect-managed conversation state programming model

To specify to have IMS Connect manage all iterations in a conversation, the `useConvID` property of the `IMSInteractionSpec` class must be set to `false`.

To enable your Java client for IMS conversational transactions by using the IMS Connect-managed conversation state programming model:

1. Obtain a connection handle from the `IMSConnectionFactory`. This connection handle must be used by all subsequent iterations for the conversation. If your Java application is a Web application, you might have to save the connection handle in the HTTP Session object for the connection handle to be retrieved and used on subsequent iterations in the same browser.
2. Set the `IMSInteractionSpec` `useConvID` property to `false` (the default value is `false`) in the first iteration. Do not change the value of this property in subsequent iterations of the conversation. If the `useConvID` property changes during a conversation, an error would occur.
3. Set the `IMSInteractionSpec` `commitMode` property to 1. Set an appropriate value for the `syncLevel` property.
4. If the client application controls the ending of a conversation, on the last iteration, set the `IMSInteractionSpec` `interactionVerb` property to `SYNC_END_CONVERSATION`.

5. Close the connection handle at the end of the conversation.

Related concepts

[Client-managed and IMS Connect-managed conversation state programming models](#)

An IMS program can support a transaction that is composed of one interaction or multiple interactions. A transaction with multiple iterations is referred to as a conversational transaction.

Processing global transactions

The IMS TM resource adapter supports global transaction management and two-phase commit processing so that your application can run in a Java EE-compliant application server to access IMS transactions.

Global transaction support with two-phase commit

To protect and maintain the integrity of your business resources, the IMS TM resource adapter supports global transaction management and two-phase commit processing.

Using this support, you can build a Java EE application to group a set of changes into one transaction, or a single unit-of-work, so that all changes within a transaction are either fully completed or fully rolled back. This support means that your application can run in a Java EE-compliant application server (for example, WebSphere Application Server) to access IMS transactions and data in a coordinated manner. Global transaction management ensures the integrity of the data in IMS.

Example of global transaction support

When you make changes to your protected resources, you want to guarantee that the changes are made correctly. For example, as a bank customer you want to transfer money from your savings account to your checking account. You want to be sure that when the money is deducted from your savings account it is added to your checking account simultaneously. You would not want this transaction to be completed only partially, with the money deducted from your savings account but not added to your checking account.

In another example, you need to buy a ticket from San Francisco to Paris but a direct flight is not available. Unless you can successfully reserve a ticket from San Francisco to Chicago and another ticket from Chicago to Paris, you will not commit to your trip to Paris. That is, you will roll back your decision to go to Paris because having a confirmed seat for only one part of your trip is not useful to you.

In both of these examples, several smaller transactions are required to complete one overall transaction. If a problem occurs with one of these smaller transactions, you would not want to commit the overall transaction (such as transferring money or going to Paris). Instead, you would want to roll back every step of the transaction so that none of the smaller transactions are committed. To transfer your money or to go on your trip to Paris successfully, you want the smaller transactions to be managed and coordinated together to complete the overall transaction.

To ensure a coordinated transaction process, the Java EE platform (which consists of a Java EE application server, Java EE application components, and a Java EE connector architecture resource adapter) provides a distributed transaction processing environment where transactions are managed transparently and resources are updated and recovered across multiple platforms in a coordinated manner.

Related concepts

[WebSphere Application Server platform configurations and communication protocol considerations](#)

Use TCP/IP as your communication protocol between IBM WebSphere Application Server and IMS Connect.

Related reference

[Two-phase commit environment recommendations](#)

To run a two-phase commit application, review the following recommendations to avoid a region or message processing from stopping due to contention for resources.

Global transaction and two-phase commit support process

A Java EE-compliant application server uses a Java transaction manager to communicate and coordinate among the application components and the resource managers.

For example, a Java EE-compliant application server such as WebSphere Application Server communicates with the application components (for example, Java servlets or Enterprise JavaBeans components) and the resource managers (for example, IMS or DB2) through the resource adapters (for example, the IMS TM resource adapter) to coordinate a transaction. If a transaction manager coordinates a transaction, that transaction is considered a global transaction. If a transaction manager coordinates a transaction with more than one resource manager, the external coordinator uses the two-phase commit protocol.

Suppose that you want to transfer money from your savings account to your checking account. If your savings account information resides on a separate resource manager from your checking account information (for example, your saving account resides on IMS and your checking account resides on DB2), the transaction manager in the application server (WebSphere Application Server) helps the application to coordinate the changes between IMS and DB2 transparently by using the two-phase commit processing. Specifically, the transaction manager works with the IMS TM resource adapter to coordinate the changes in IMS.

The IMS TM resource adapter is designed to work with the Java transaction manager in the Java EE platform, the Resource Recovery Services (RRS) of z/OS, and IMS Connect to make consistent changes to IMS and other protected resources.

To participate in two-phase commit processing with IMS, the IMS TM resource adapter uses the IMS OTMA synchronization level sync-point protocol. To participate in global transaction and two-phase commit processing when the changes are requested from a remote application, IMS uses RRS on z/OS.

RRS acts as the external coordinator or sync-point manager to coordinate the update and recovery of resources. The IMS TM resource adapter and IMS Connect interact with the Java transaction manager running in the application server and RRS on z/OS to allow a global transaction running on a Java EE platform to participate in a coordinated update with IMS running on the host.

Related concepts

[WebSphere Application Server platform configurations and communication protocol considerations](#)
Use TCP/IP as your communication protocol between IBM WebSphere Application Server and IMS Connect.

Related reference

[Two-phase commit environment recommendations](#)

To run a two-phase commit application, review the following recommendations to avoid a region or message processing from stopping due to contention for resources.

Global transactions support requirements

When setting up a Java EE application to participate in a global transaction, you must select the appropriate communication protocols and RRS must be enabled on IMS.

- IMS must be enabled for Resource Recovery Services (RRS) processing. To ensure that IMS is enabled for RRS processing, check that the RRS value in the startup parameter within your IMS environment is set to Y. This setting will be displayed in the job logs generated when IMS is brought up.
- If you are using two-phase commit processing with TCP/IP, you need to enable IMS Connect for RRS processing by either issuing the IMS Connect command SETRRS ON or set RRS=Y in the IMS Connect configuration file.

Global transaction with TCP/IP

In a global transaction scope, your Java EE application component can access an IMS transaction by establishing a TCP/IP connection with IMS Connect.

Global transaction with TCP/IP

The IMS TM resource adapter interacts with the Java transaction manager by using the X/Open (XA) protocol to manage the global transaction and two-phase commit processing. The XA protocol defines a set of interfaces and interactions that describe how the Java transaction manager and the resource managers interact in a distributed transaction processing environment. The IMS TM resource adapter, together with IMS Connect, uses the XA protocol and works with IMS and Resource Recovery Services (RRS) on z/OS to make consistent changes.

To use global transaction with TCP/IP, RRS must run on the same z/OS image with IMS Connect and IMS.

If both IMS and IMS Connect are at IMS 14 or later, IMS and IMS Connect do not have to be on the same z/OS image if IMS Connect is configured to support cascading global transactions. However, an instance of RRS must be running on each z/OS image.

Global transaction support in client applications

On the Java EE platform, you can use either a programmatic or a declarative transaction demarcation approach to manage transactions in your application.

The programmatic approach is the component-managed (or bean-managed) transaction and the declarative transaction demarcation approach is the container-managed transaction.

Component-managed (or bean-managed) transactions

The Java EE application uses the Java Transaction API (JTA) `javax.transaction.UserTransaction` interface to demarcate a transaction boundary to a set of changes to the protected resource programmatically. Component-managed transactions can be used in both the servlet and the EJB environment. In the case of an EJB component, you set the transaction attribute in its deployment descriptor as `TX_BEAN_MANAGED`.

A transaction normally begins with a `UserTransaction.begin()` call. When the application component is ready to commit the changes, it invokes a `UserTransaction.commit()` call to coordinate and commit the changes. If the application component must roll back the transaction, it invokes `UserTransaction.rollback()` and all changes are backed out. For example:

```
// Get User Transaction
javax.transaction.UserTransaction transaction =
ejbcontext.getUserTransaction();

// Start transaction
transaction.begin();

// Make changes to the protected resources.
// For example, use the Java EE or JCA CCI Interaction interface
// to submit changes to an EIS system(s)
interaction.execute(interactionSpec, input, output);

if (/* decide to commit */) {
// commit the transaction
transaction.commit();
} else { /* decide to roll back */
// rollback the transaction
transaction.rollback();
}
```

Container-managed transactions

Container-managed transactions can be used only in the EJB environment. The EJB component specifies a container-managed transaction declaratively through the transaction attribute in the deployment descriptor (such as TX_REQUIRED). A container-managed transaction is managed by the EJB container. The container calls the appropriate methods (such as begin, commit, or rollback) on behalf of the EJB component. This declarative approach simplifies the programming calls in the EJB component.

Related information

[Java EE architecture and JTA specifications](#) For more information about the Java EE architecture and JTA specifications, see the Java 2 Platform, Enterprise Edition documentation.

Two-phase commit environment recommendations

To run a two-phase commit application, review the following recommendations to avoid a region or message processing from stopping due to contention for resources.

- Have as many message processing program (MPP) regions as possible running to ensure that two-phase commit applications do not contend for a region, because a transaction that is within a two-phase commit application uses an MPP region for the duration of the entire two-phase commit transaction.
- If a number of IMS transactions are performed within a two-phase commit transaction, at least that many MPP regions must be available to avoid stopping the two-phase commit application.
- To safeguard against a transaction that might be waiting for an extensive amount of time for resources, set an appropriate timeout value for each interaction taking place within the global transaction.
- Avoid having an excessive number of database interactions performed in one two-phase commit transaction. If multiple IMS transactions are used within a two-phase commit transaction, they could contend or lock in an attempt to update or modify the same data. To avoid this problem, write an application that will prevent a user from accessing duplicate entries within the same two-phase commit operation.
- Consider configuring your internal resource lock manager (IRLM) or Program Isolation (PI) locking manager to use a block size that is as small as the smallest entry to that database. Larger block sizes might have two transactions contending for entries that might not even be the same and yet reside close to one another on the disk.
- If multiple interactions are performed using the same IMS transaction on the same IMS database within a global transaction (unit-of-work), each interaction with that IMS transaction must run on a separate MPP region. The IMS transaction must have a SCHDTYP=PARALLEL and a PARLIM=0 value, to indicate that the IMS transaction can run on multiple MPP regions and that it will always meet the scheduling requirements (the number of messages will be greater than zero) to process every interaction on a new MPP region.
- If a region is hung, and no execution timeout value has been set, you can end the attempt to run a transaction that is hanging the MPP region by issuing a **/STOP REGION** IMS command with the abend transaction parameter. For example, **/STOP REGION reg# ABDUMP tranname**. This command rolls back the transaction for that particular interaction and free the MPP region.

Other transaction support

The IMS TM resource adapter also supports local transaction, one-phase commit processing, non-global transaction processing, and conversational transaction processing in a global transaction scope.

Local transaction

The Java EE Connection Architecture defines the `javax.resource.cci.LocalTransaction` interface for a resource manager, rather than a transaction manager, to coordinate a transaction locally. However, the IMS TM resource adapter supports only transaction coordination with a transaction manager. Thus, the IMS TM resource adapter does not support the `javax.resource.cci.LocalTransaction` interface. If you call the `IMSConnection.getLocalTransaction()` method you will get a `NotSupportedException`. To use transaction support with the IMS TM resource adapter, you need

to either use the JTA transaction interface, or set an appropriate transaction attribute in the deployment descriptor in your application. See the topic on global transaction support in client applications for more information.

One-phase commit processing

The IMS TM resource adapter supports one-phase commit optimization with a transaction manager. As a result, if all changes inside a transaction scope belong to the same IMS resource, the transaction manager might perform one-phase-commit optimization such that the transaction manager sends the phase two commit request directly to the resource manager for committing the changes without sending the phase one prepare request.

Non-global transaction processing

If no global transaction processing is used in the application (for example, when the transaction attribute is set to TX_NOTSUPPORTED), all non-global transaction processing uses "Sync-On-Return" (OTMA SyncLevel=None). By the time the IMS transaction is committed, the output has been returned to the client.

Conversational transaction processing in global transaction scope

IMS uses a conversational program to divide processing into a connected series of client-to-program-to-client interactions (also called iterations). Each iteration is a type of IMS conversational transaction. Conversational processing is used when one transaction contains several parts. Each part that comprises one large transaction is separately committed or rolled back.

You can run a conversational transaction in the global transaction scope if:

- Each iteration is run under the same transaction level. For example, if the first iteration is processed with a global transaction scope, then all the subsequent iterations in that IMS conversational transaction must be processed at a global transaction level. If you issue the second iteration with no transaction scope, IMS OTMA reports an error.
- Each iteration must be completed with a commit or rollback call before issuing the next iteration in the IMS conversation. You cannot group multiple iterations in a single global transaction scope.

Related concepts

[Global transaction and two-phase commit support process](#)

A Java EE-compliant application server uses a Java transaction manager to communicate and coordinate among the application components and the resource managers.

Common Client Interface (CCI)

You can create an application to use the IMS TM resource adapter to interact with IMS by using the Common Client Interface (CCI).

You can use the application code that is generated by the J2C wizard in a Rational or WebSphere integrated development environment (IDE) to access IMS transactions through the IMS TM resource adapter. No coding is necessary when you use this approach. Alternatively, you can write the application source code without using an IDE.

To write the code yourself, you must use the CCI programming interface. The CCI API provides access from Java EE clients, such as enterprise beans, JavaServer Pages (JSP) pages, and servlets, to backend enterprise information systems (EIS) such as IMS.

Applications that follow the CCI programming interface model have a common structure, independent of the EIS that is being used. The Java EE Connector Architecture (JCA) specification defines two objects that are necessary for the application:

- Connection objects that represent connections to an EIS
- ConnectionFactory objects that create those Connection objects

These objects are what an application server uses to manage security, transaction context, and connection pools for a resource adapter. An application that uses the IMS TM resource adapter CCI programming interface starts by obtaining an `IMSConnectionFactory` object. The `IMSConnectionFactory` object can be obtained in two ways:

- **Managed:** If you are using an application server, the `IMSConnectionFactory` object is normally created from the resource adapter through an administration interface such as the WebSphere Application Server administrative console. This type of environment is called a *managed environment* because an application server is used to manage the qualities of service of the connections in its environment. For example, you would create an `IMSConnectionFactory` object and configure its custom properties by using the WebSphere Application Server administrative console. Custom properties such as the hostname and port number of the target IMS system are configured on the `IMSConnectionFactory` object. When the `IMSConnectionFactory` object is created, it can be made available for use by any enterprise applications through JNDI.
- **Non-managed:** When you are not using an application server to manage the connection (such as running the application as a stand-alone Java application), this type of environment is called a *non-managed environment*. In this type of configuration:
 - You must manually create an `IMSManagedConnectionFactory` object and set its custom properties. You can then create an `IMSConnectionFactory` object from the `IMSManagedConnectionFactory` object.
 - The IMS TM resource adapter uses the `DefaultConnectionManager` class in the JCA 1.5 Connection Management API for connections. This class does not provide any connection pooling, and all the connection pool properties would be ignored. The IMS TM resource adapter would open and close a socket connection with IMS Connect for every transaction request. Opening and closing socket connections introduces resource overhead and is likely to result in performance degradation as compared to a managed scenario. You can build your own connection pooling by implementing a `PoolManager` class to be used with the `DefaultConnectionManager` class. For more information, see the JCA 1.5 Connection Management API section.

Because CM0 transactions are recoverable, IMS Connect creates a separate TPIPE for each client that uses CM0. Without the connection pool management function provided by an application server, too many TPIPEs would be created and overload the system.

IMS connections

An `IMSConnection` object can then be created from that `IMSConnectionFactory` object. The properties of the `IMSConnection` object can either be specified in an `IMSConnectionSpec` object passed as a parameter to the `getConnection` method, or the default values defined in the `IMSConnectionFactory` is used. After an `IMSConnection` is obtained, an `IMSInteraction` instance can be created from the `IMSConnection` instance. An `IMSInteraction` instance represents the interaction that is going to be executed on that connection. As with the connection, interactions can have custom properties taken from the `IMSInteractionSpec` class.

Input and output

To perform the interaction, the application makes a call to the `execute()` method of the `IMSInteraction` object, passing it input and output objects to hold the data. An input byte array must be created containing values for each field in the input message to IMS. Likewise, an output byte array must also be created to hold the response message returned by IMS. The value of each field in the output message is extracted from the output byte array.

You can write the input and output byte arrays yourself, or you can use the J2C options in a Rational or WebSphere development environment to create Java data bindings for the input and output messages of your CCI application.

Requirements for applications that retrieve data from an IMS system:

- Use the `IMSConnectionFactory` object to create an `IMSConnection` object.
- Use the `IMSConnection` object to create an `IMSInteraction` object.

- Use the IMSInteraction object to execute transactions on the backend IMS system.
- Close the IMSInteraction and IMSConnection objects.

Related information

[IMS TM Resource Adapter Version 15 Java API Specification](#) For more information, see the IMS TM Resource Adapter Version 15 Javadoc information.

Sample CCI application code

This sample CCI application code demonstrates how to obtain a connection and then an IMSInteraction object to execute an IMS transaction.

In a managed environment, use the CCI API to look up an IMSConnectionFactory instance from the JNDI namespace and then use it to get an IMSConnection instance.

```
ConnectionFactory cf = null;
if (isManaged) {
    //Use JNDI lookup to get ConnectionFactory instance
    //Assume the connection factory has a JNDI name of MyIMS
    Context ic = new InitialContext();
    cf = (ConnectionFactory) ic.lookup("MyIMS");
}
```

If JNDI is not configured, a CCI application can manually configure an IMSManagedConnectionFactory object and use it to obtain a connection factory.

```
IMSManagedConnectionFactory mcf = new IMSManagedConnectionFactory();
mcf.setDataStoreName("MyDSName");
mcf.setHostName("myHostNm");
mcf.setPortNumber(new Integer(1234));
...
//Create connection factory from ManagedConnectionFactory
cf = (IMSConnectionFactory) mcf.createConnectionFactory();
```

The following example shows the use of the CCI interfaces to execute a command on an EIS.

- Use the IMSConnectionFactory object to create an IMSConnection object.
- Use the IMSConnection object to create an IMSInteraction object.
- Use the IMSInteraction object to execute transactions on the backend IMS system.
- Close the IMSInteraction and IMSConnection objects.

```
public void execute() {
    try {
        ConnectionFactory cf = null;
        if (isManaged) {
            //Use JNDI lookup to get ConnectionFactory instance - assumes
            //connection factory has JNDI name of MyIMS
            Context ic = new InitialContext();
            cf = (ConnectionFactory) ic.lookup("MyIMS");
        } else {
            //Create and set values for ManagedConnectionFactory
            IMSManagedConnectionFactory mcf = new IMSManagedConnectionFactory();
            mcf.setDataStoreName("MyDSName");
            mcf.setHostName("myHostNm");
            mcf.setPortNumber(new Integer(1234));
            //Create connection factory from ManagedConnectionFactory
            cf = (IMSConnectionFactory) mcf.createConnectionFactory();
        }
        // Create an IMSConnection object
        Connection connection = cf.getConnection();

        //Create an IMSInteraction from the connection to
        //interact with IMS to run the IVTNO transaction (Phonebook)
        IMSInteraction interaction = (IMSInteraction) connection.createInteraction();
        IMSInteractionSpec ixnSpec = new IMSInteractionSpec();
        ixnSpec.setInteractionVerb(IMSInteractionSpec.SYNC_SEND_RECEIVE);

        //Create new input record
        input = new PhoneBookInputRecordField("cp037");
        input.setIn__ll((short)59);
        input.setIn__zz((short) 0);
    }
}
```

```

input.setIn__trcd("IVTN0");
input.setTranCodeLength(10);
input.setIn__command("DISPLAY");
input.setIn__name1("LAST3");
input.setIn__name2("");
input.setAllFieldsGiven(false);
PhoneBookOutputRecordField

//Create new output record
output = new PhoneBookOutputRecordField("cp037");

//Execute interaction by calling the execute() method
interaction.execute(ixnSpec, input, output);

//Display output
System.out.println ("Output is: ");
System.out.println("\nMessage: "
+ output.getOut__mesg()
+ "\nName:"
+ output.getOut__name1()
+ " "
+ output.getOut__name2()
+ "\nExtension:"
+ output.getOut__extn()
+ "\nZipcode: "
+ output.getOut__zip());
} catch (Exception e) {
    e.printStackTrace();
} finally {
    //Close both the interaction and the connection
    interaction.close();
    connection.close();
}
}

```

Samples and tutorials

Samples and tutorials for developing Java applications or Web services to access IMS transactions can be found in several WebSphere and Rational integrated development environments (IDEs).

- Tutorials and samples for creating J2C applications for different types of IMS transactions are available in the online help of various WebSphere and Rational development environments.
- WebSphere Integration Developer users can find the examples for creating imports to access IMS transactions in IBM Documentation for WebSphere Integration Developer.
- WebSphere Transformation Extender users can find tutorials for creating maps for various types of IMS transactions.
- An end-to-end callout IVP sample is provided with IMS and the IMS TM resource adapter.
 - In IMS, jobs and tasks are provided in the IMS installation verification program (IVP) that include an IMS callout application and the required IMS OTMA destination descriptor. The OTMA destination descriptor defines the tpipes that callout requests are queued.
 - In the IMS TM resource adapter, a sample application is provided that listens for callout requests from IMS.

Related information

[Tutorials in IBM Documentation for Rational Application Developer V9](#) See the IMS tutorials under the Java category in tutorials section in IBM Documentation.

[Samples in IBM Documentation for Rational Application Developer V9](#) See the IMS samples in J2C samples section IBM Documentation.

Chapter 50. Running your application on a stand-alone WebSphere Application Server

You can test your application by using a stand-alone WebSphere Application Server. You must first export your Java EE application as an Enterprise Application Archive (EAR) file and then install the EAR file on the server.

Prerequisite: If you do not package the IMS TM resource adapter with your enterprise application, you must ensure that:

- The IMS TM resource adapter that your application uses is deployed on the stand-alone WebSphere Application Server.
- A connection factory has been defined to create connections to the IMS Connect and IMS that you want to run the IMS transaction.

After the prerequisites are met, you can proceed to install your EAR file on WebSphere Application Server.

Related tasks

[Installing the resource adapter on WebSphere Application Server](#)

Use the WebSphere Application Server administrative console to deploy the IMS TM resource adapter RAR file.

[Creating a connection factory in WebSphere Application Server](#)

After you install and deploy the IMS TM resource adapter on WebSphere Application Server, create a connection factory for the IMS TM resource adapter, if it does not already exist.

Installing your EAR file on WebSphere servers

You can export your application as an EAR file and deploy it in your WebSphere servers and then run the application.

Prerequisite: You must have installed the IMS TM resource adapter RAR and created a connection factory for the IMS TM resource adapter.

The steps to install your application EAR file is similar to how you deploy the IMS TM resource adapter IVP EAR file. To install your application EAR file and run your application, follow the steps described in the following topics:

- [“Deploying the IVP EAR file to WebSphere Application Server ” on page 740](#)
- [“Deploying the IVP EAR file to WebSphere Liberty servers” on page 741](#)

Chapter 51. Diagnosing problems

Errors might come from various sources that are involved in the transaction of the message request or response. You can configure the application server to log and trace the component-specific information, and IMS TM resource adapter also generate messages from Java exceptions.

Diagnosing IVP failures

If an error occurs when you run the IMS TM resource adapter IVP, first check for information in WebSphere Application Server logs.

The WebSphere Application Server logs are located under *WebSphere_install_directory\AppServer\logs*. The *trace.log* file lists exceptions and provides stack traces that are useful for diagnosing the failure.

The following problems are possible causes for IVP failures:

- WebSphere Application Server is not started properly.
- If you receive the following `ResourceException` error about not finding the RAR reference in the context in a cluster environment, ensure that you have created a copy of the RAR file with the appropriate cluster or server scope, in addition to installing the RAR at the node level.

```
Error 500: javax.resource.ResourceException:  
Context: myCell/clusters/myCluster1, name: myIMSTMRRARef:  
First component in name myIMSTMRRARef not found.
```

- The enterprise application for the IVP, `IMSIC0IVPServiceEAR`, is not started. To determine the status of the enterprise application:
 1. In the WebSphere Application Server administrative console, expand **Applications** in the navigation tree on the left pane of the console.
 2. Click the **Enterprise Applications** link. The `IMSIC0IVPServiceEAR` EAR file displays as one of the installed applications. A green arrow indicates that the application is started. If the application is not started, follow the instructions for starting the application in the topic on [“Deploying the IVP EAR file in the Java EE application server”](#) on page 740.
- Invalid data was specified when you configured the connection factory. For example:
 - The host name is misspelled or is not sufficiently qualified (for TCP/IP communication).
 - An incorrect port number was specified for the target IMS Connect (for TCP/IP communication).
 - The IMS Connect name was misspelled.
 - The data store name is invalid for the target IMS or was misspelled. The data store name must be in uppercase characters.
- IMS is not running.
- IMS Connect is not running.
- The IMS Connect port is not active. Use the IMS Connect command **VIEWHWS** to determine if the port is active. Use the IMS Connect command **OPENPORT** to activate an IMS Connect port.
- The target IMS data store is not active. Use the IMS Connect command **VIEWHWS** to determine if the data store is active. Use the IMS Connect command **OPENDS** to activate a IMS data store.
- TCP/IP fails. Issue the ping command before you run the IVP to ensure the connection is working.
- The wrong level of IMS Connect is running on the host. See [“Prerequisites for running the IVP”](#) on page 739 for the required level of IMS Connect.
- If you expect the IVP to return the message `DFS058I hh:mm:ss START COMMAND COMPLETED`, but the message `DFS1292E SECURITY VIOLATION` is returned instead, check for the conditions:

- The username and password that you provided in the custom properties of your IVP connection factory are not authorized to run the **/STA** OTMA command that the IVP issues to IMS.
- Security is not enabled in IMS Connect (RACF=Y in the **VIEWHWS** command output).
- Security is not enabled in IMS OTMA (SECURITY=FULL in the **/DIS** OTMA command output).
- If you receive the following exception, you must configure the connection factory that is used by the IVP to set the value of the CM0Dedicated property to false.

```
javax.resource.ResourceException: IC00087E:
com.ibm.connector2.ims.ico.IMSTCIPManagedConnection@28f39301.call(Connection,
InteractionSpec, Record, Record) error. Protocol violation. Commit Mode 1 is not
allowed for interactions on a dedicated persistent socket
```

Diagnosing problems accessing IMS from Java applications

If you are unable to access IMS from your Java application, verify that IMS TM resource adapter is properly set up and configured, that IMS Connect and IMS are active, and that you can access the host system by using the **ping** command.

- Verify that you have the correct setup and configurations for using the IMS TM resource adapter. See the topic on supported versions, configurations, and platforms for supported configurations.
- Verify that IMS Connect is active by ensuring that the outstanding IMS Connect reply HWSC0000I *IMS CONNECT READY* *ims_connect_name* is displayed on the system console of the target system. For message HWSC0000I to display on the system console when IMS Connect is active, ensure either that **WTORCMD=Y** is specified in the HWS statement of the IMS Connect HWSCFGxx configuration member or that the **WTORCMD** parameter is omitted from the HWS statement.
- Verify that the port and data store are ACTIVE by entering the IMS Connect command **VIEWHWS** at the IMS Connect outstanding reply.
- Verify that IMS is active by ensuring that the outstanding IMS reply DFS996I *IMS READY* is displayed on the system console of the target system.
- Verify that the cross-system coupling facility (XCF) service status of both the IMS and IMS Connect members is ACTIVE by entering the IMS command **/DISPLAY OTMA** at the outstanding IMS reply. The display output will be similar to the following output:

DFS000I SECURITY	GROUP/MEMBER IMS1	XCF-STATUS	USER-STATUS	
DFS000I IMS1	XCFGRPNM			
DFS000I IMS1	-IMSNAME	ACTIVE	SERVER	FULL
DFS000I IMS1	-ICONNAME	ACTIVE	ACCEPT TRAFFIC	
DFS000I	*02033/143629*	IMS1		

- If you are using TCP/IP to communicate between the Java application and IMS Connect, verify that you can successfully access the target host system by using the **ping** command. If you cannot **ping** the host system and you are using a host name rather than an IP address, ensure that the host name is sufficiently qualified.

If your IMS service is not providing the expected output from the IMS transaction, ensure that the output message that is returned by the IMS application program matches the output COBOL definition used by the service. For a Java EE application, you can view the IMS OTMA message containing the message returned by the IMS application program by setting the traceLevel property to 3. See the topic on logging and tracing with the IMS TM resource adapter for instructions on how to turn on the IMS TM resource adapter trace.

Related tasks

[Logging and tracing IMS TM resource adapter information](#)

Configure the logging and tracing setting in WebSphere Application Server to ensure messages are logged, and diagnostic trace information is recorded for the IMS TM resource adapter.

Diagnosing problems with callout requests

Callout request errors are often related to the execution timeout setting, incorrect or corrupted callout messages, or network failures.

- Execution timeout occurs (no callout message is being retrieved within the time interval).

Explanation: Possible reasons include:

- No IMS application issues a callout request and the tpipe does not hold any request.
- The tpipe name that is specified in the alternate client ID field in the Java application does not match the tpipe name that is used by the IMS application for the callout request.
- The tpipe name that is specified in the alternate client ID field in the Java application does not match the tpipe name that is specified in the OTMA destination descriptor.
- The message processing program (MPP) has not started.

User Action: Either increase the timeout interval (the `executionTimeout` property), or verify that the tpipe name matches. Ensure that the MPP is started.

- IMS Connect fails to process the callout request and the message is placed in the Dead Letter Queue.

Explanation: This problem is most likely due to a bad callout message.

User Action: Fix the callout message in the IMS application.

- IMS Connect fails to deliver the callout request to the IMS TM resource adapter.

Explanation This problem is most likely due to a network failure.

User Action: If a network failure occurs, you might receive a Java exception. Modify the Java application to catch the exception.

- The IMS TM resource adapter or the Java application fails to process the callout request.

Explanation Possible reasons include:

- The callout request is corrupted.
- The message is not a callout request, but a regular asynchronous output message.

User Action: Verify that the callout request is issued correctly from the IMS application. Verify that the tpipe name that is specified in the Java application is correct, and that it is not the name for another queue for a regular asynchronous output message.

Java exceptions that involve output messages

If any errors occur when the IMS TM resource adapter passes the message through IMS Connect to IMS, and as IMS processes this transaction and attempts to return the output, the Java client receives an exception.

The type of exception that is thrown indicates whether an output message is available for retrieval. For example, if the Java client receives an `IMSDFSMessageException` exception that indicates that the transaction is stopped, the application was not run. Therefore, no output message is available for retrieval. However, if the transaction runs but the `executionTimeout` value expires before the output message is returned to IMS Connect, the Java client receives an `EISSystemException` exception. This exception indicates that an execution timeout has occurred. In this case, the output message is queued to the appropriate IMS OTMA asynchronous output queue or tpipe for later retrieval.

Asynchronous output interaction errors

In general, you can use the asynchronous output interactions, `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT` and `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT`,

to retrieve output messages queued for any client ID, regardless of how those messages were queued to the associated client ID. The messages might be queued either as a result of a failed commit mode 0 transaction, or from an IMS application that issued an INSERT call to an Alternate Program Communication Block (ALTPCB).

For example, in retrieving an output message from a failed commit mode 0 transaction, the client ID that is specified in the `IMSConnectionSpec` class for retrieval request must match the client ID that was specified on the failed commit mode 0 transaction. If nothing is in the OTMA asynchronous output queue for that particular client ID, an execution timeout exception occurs. The timeout exception can mean one of the following two scenarios:

- The queue has no messages.
- The timeout value did not provide enough time for IMS Connect to retrieve the message from the queue.

For both asynchronous output interactions, execution timeout is the length of time IMS Connect waits for a response from IMS. If you do not specify an execution timeout value for a retrieval request, the default execution timeout value is used. The default timeout value is the IMS Connect configuration member `TIMEOUT` value. You might need to experiment with the execution timeout value to ensure that output messages are returned for all types of interactions.

Related concepts

[Interactions supported with commit mode and sync level combinations](#)

The interactions that are supported by the IMS TM resource adapter depend on the combination of commit mode and sync level that you choose.

[Asynchronous output programming model](#)

Use this programming model to retrieve output that has been queued by IMS.

[Securing message retrieval from IMS hold queues](#)

You can specify the user ID and password information to ensure that only authorized users are allowed to retrieve asynchronous output or callout request messages from the IMS hold queue.

Logging and tracing IMS TM resource adapter information

Configure the logging and tracing setting in WebSphere Application Server to ensure messages are logged, and diagnostic trace information is recorded for the IMS TM resource adapter.

Application logging captures various events that happen in the execution of an application. The IMS TM resource adapter provides Java logging support that builds on the `java.util.logging` package included in WebSphere Application Server. Because Java message logging is included in the trace log, logging and tracing is used together to describe the troubleshooting support feature. While the message logs are for end users, application developers, or system administrators, diagnostic trace information is for technical support personnel to obtain finer and verbose details of events that took place.

In WebSphere Application Server, tracing for all components is disabled by default. After you configure the logging and tracing setting, tracing is enabled. A trace file is created when logging and tracing is enabled, and a Java application is run in WebSphere Application Server.

For standalone applications, you can also create your own logger by using the `java.util.logging.Logger` object, create a `FileHandler`, pass the log file to the `FileHandler`, and set the logger to the `ManagedConnectionFactory` object.

Tip: Ensure that only one client is running when the trace is on.

Logging and tracing in WebSphere Application Server

To turn on logging and tracing for IMS TM resource adapter information in WebSphere Application Server, set the trace output to a file or the memory buffer, and specify the log level for the packages that contain IMS TM resource adapter information.

To set controls for logging and tracing IMS TM resource adapter information:

1. On the welcome page of WebSphere Application Server administrative console, expand **Troubleshooting**, and click **Logs and Trace** to open the Logging and Tracing page.
2. On the Logging and Tracing page, click the server that you want to trace (for example, server1), and click **Diagnostic Trace**.

The Diagnostic trace service page opens.

3. On the **Configuration** tab, set Trace Output to **Memory Buffer** or **File**

Tips:

- Changes made in the **Configuration** tab persist after the server restarts. To temporarily change the trace level, make the changes in the **Runtime** tab.
 - Trace information are events at the fine, finer and finest levels, and can be written only to the trace log. For differences among the different log levels, see [Log level settings](#).
4. If you select **File** to store trace output, you can accept the default name and location of the trace output file or you can modify it. To modify the default name and location, enter a different name and location of the file in the File Name field.
 5. Click **OK**.
 6. On the Logging and Tracing page, click **Change Log Detail Levels**.
The Change Log Detail Levels page displays.
 7. To enable logging and tracing for the IMS TM resource adapter:
 - a) Under General Properties, click **Components**.
 - b) Enter the following string in the text box:

```
*=info: com.ibm.j2ca.RAIMSTM=finest
```

You can combine the string with others to enable tracing in other components.

8. Click **OK**. If you are making changes on the **Configuration** tab, click **Save** at the top of the page to save the changes to the master configuration.
9. Run your Java application and then examine the trace file.

Logging and tracing in WebSphere Liberty

To turn on logging and tracing for IMS TM resource adapter information in WebSphere Liberty, specify the log level for the packages that contain IMS TM resource adapter information.

To print out the logging and tracing information for the IMS TM resource adapter, add a `<logging>` entry in the `server.xml` file.

1. Open the `server.xml` file.
2. Add a `<logging>` for the IMS TM resource adapter by specifying the `com.ibm.j2ca.RAIMSTM` package for the `traceSpecification` element, and set the trace level to `FINEST`:

```
<logging traceSpecification="*=info:com.ibm.j2ca.RAIMSTM=finest"></logging>
```

For more information about WebSphere Liberty server logging and tracing configuration, the log level settings and their content, and what the primary log files, see the documentation for the version of WebSphere Liberty server that you are using.

Related information

[Liberty profile logging and tracing \(V8.5.5\)](#)

Creating a stand-alone Logger with output sent to a file

This code sample demonstrates how to create a stand-alone Logger with output sent to a file.

This sample uses the `java.util.logging.Logger` object, sets the logger to the `ManagedConnectionFactory` object and the log level to `FINEST`:

1. In the application, create a `java.util.logging.Logger` object and a `FileHandler`, pass the output file to the `FileHandler`, set the logger to the `ManagedConnectionFactory` object, and set the log level:

```
com.ibm.connector2.ims.ico.IMSManagedConnectionFactory mcf =
    new com.ibm.connector2.ims.ico.IMSManagedConnectionFactory();
mcf.setHostName("yourHostName");
mcf.setDataStoreName("yourDataStore");
mcf.setPortNumber(new Integer(yourPortNumber));

// Create a file logger
Logger logger = Logger.getLogger("myLogger");
FileHandler fh = new FileHandler("output.txt");
SimpleFormatter formatter = new SimpleFormatter();
fh.setFormatter(formatter);
logger.addHandler(fh);

// Set the file logger to the IMSManagedConnectionFactory object
// Set log level to finest
LogUtils logUtils = new LogUtils(logger, "log", "yourProduct", "yourVersion");
mcf.setLogUtil(logUtils);
mcf.getLogUtil().setLoggingLevel(Level.FINEST);
connFactory = (ConnectionFactory) mcf.createConnectionFactory();
```

2. Add your other code for the application.
3. Run the application.

As the application runs, information is logged to the `output.txt` file.

Analyzing the trace data

The trace data in the `trace.log` file in WebSphere Application Server includes the time and sequence of method calls by the base classes and additional classes based on your configuration of the logging and tracing setting.

The trace output format consists of:

- a timestamp
- a thread ID
- an abbreviated short name of the logging component
- an event type indicator
- the class that issued the message or trace event
- additional text message

The following trace entries demonstrate the information that is logged when WebSphere Application Server is properly configured for the IMS TM resource adapter.

- RAIMSTM is the short name for the IMS TM resource adapter.
- Each time a major event takes place, information is logged for its entrance and exit.

```
[7/5/11 14:24:04:968 PDT] 00000012 RAIMSTM      2
    com.ibm.ims.ico.IMSTCPIPAdapter generateClientID() Entering method.
[7/5/11 14:24:04:968 PDT] 00000012 RAIMSTM      3
    com.ibm.ims.ico.IMSTCPIPAdapter generateClientID()
    LocalPort = [3107]
[7/5/11 14:24:04:968 PDT] 00000012 RAIMSTM      3
    com.ibm.ims.ico.IMSTCPIPAdapter generateClientID()
    IP Address = [9030020219]
[7/5/11 14:24:04:968 PDT] 00000013 RAIMSTM      3
    com.ibm.ims.ico.IMSTCPIPAdapter generateClientID()
    Generated ID = [HWSY6N4P]
[7/5/11 14:24:04:968 PDT] 00000013 RAIMSTM      <
    com.ibm.ims.ico.IMSTCPIPAdapter generateClientID() Exiting method.
[7/5/11 14:24:04:968 PDT] 00000013 RAIMSTM      3
    com.ibm.ims.ico.IMSTCPIPAdapter connect()
    <-- [com.ibm.ims.ico.IMSTCPIPAdapter@1eca167.connect()]
```

If the connection to the IMS host system is successful, the trace shows the message is sent to IMS and the client ID is generated.

```
[7/5/11 14:24:04:968 PDT] 00000013 RAIMSTM      2
    com.ibm.connector2.ims.ico.inbound.IMSInboundUtil .sendIMSMessage()
    Entering method.
[7/5/11 14:24:04:968 PDT] 00000013 RAIMSTM      3
    com.ibm.connector2.ims.ico.inbound.IMSInboundUtil .sendIMSMessage() Mode
is:
    MODE_RECEIVE_ASYNCOUTPUT_AUTO
[7/5/11 14:24:04:968 PDT] 00000012 RAIMSTM      3
    com.ibm.ims.ico.IMSTCPIPAdapter generateClientID()
    Generated ID = [HWSYEW06]
[7/5/11 14:24:04:968 PDT] 00000012 RAIMSTM      <
    com.ibm.ims.ico.IMSTCPIPAdapter generateClientID() Exiting method.
[7/5/11 14:24:04:968 PDT] 00000012 RAIMSTM      3
    com.ibm.ims.ico.IMSTCPIPAdapter connect()
    <-- [com.ibm.ims.ico.IMSTCPIPAdapter@1ee8e00.connect()]
```

The following trace entries demonstrate a failed connection to IMS (host dev555.vmec.ibm.com; port 9999). An ICO0003E message is logged. When the connection fails, no message is sent.

```
connect() -> [com.ibm.ims.ico.IMSTCPIPAdapter@2037bc2.connect():
    HostName=dev555.vmec.ibm.com PortNumber=9999] SocketTimeout = [0]
[7/5/11 14:24:16:734 PDT] 00000021 RAIMSTM      E
    com.ibm.connector2.ims.ico.IMSTCPIPManagedConnection connect() ICO0003E:
    com.ibm.connector2.ims.ico.IMSTCPIPManagedConnection@2037bae.connect()
error.
    Failed to connect to host [dev555.vmec.ibm.com], port [9999].
    [java.net.UnknownHostException: dev555.vmec.ibm.com]
    Explanation=The IMS TM resource adapter was unable to connect to the host
and
    port combination. java_exception indicates the reason for the failure to
    connect.
    UserAction=Examine the exception to determine the reason for the failure to
    connect to the host.
[7/5/11 14:24:16:734 PDT] 00000021 RAIMSTM      2
    com.ibm.connector2.ims.ico.IMSManagedConnection errorOccurred(Exception)
    Entering method.
[7/5/11 14:24:16:734 PDT] 00000021 ConnectionEve W
    J2CA0206W: A connection error occurred.
    To help determine the problem, enable the Diagnose Connection Usage option
    on the Connection Factory or Data Source.
[7/5/11 14:24:16:734 PDT] 00000021 ConnectionEve A
    J2CA0056I: The Connection Manager received a fatal connection error from
    the Resource Adapter for resource
```

With a successful connection, when the message is sent, the trace data shows the buffer that is sent. This information is the same information you see in an IMS Connect recorder trace, unless you have customized exit routine.

```
[7/5/11 14:24:04:984 PDT] 00000012 RAIMSTM      2
    com.ibm.ims.ico.IMSTCPIPAdapter send(byte[]) Entering method.
[7/5/11 14:24:04:984 PDT] 0000000b ApplicationMg A   WSVR0221I: Application started:
    query
[7/5/11 14:24:04:984 PDT] 0000000b CompositionUn A   WSVR0191I: Composition unit
    WebSphere:cuname=query in BLA WebSphere:blaname=query started.
[7/5/11 14:24:04:984 PDT] 00000013 RAIMSTM      3
    com.ibm.ims.ico.IMSTCPIPAdapter send(byte[]) Buffer sent:
    [
        0000001f2 001c0100 5cc8e6e2 d1c1e55c | ...2. ..*HWSJAV*| : 16
        000000000 c0ff0000 c8e6e2e8 f6d5f4d7 | ....{...HWSY6N4P| : 32
        01100000 28004040 40404040 4040a0e0 | ..... μ\ | : 48
        000000000 00000000 00000000 00010000 | ..... | : 64
        00480240 01104040 40404040 40400000 | .ç. .. | : 80
        000000000 00000000 00000000 00000000 | ..... | : 96
        000000000 00000000 00000000 00000000 | ..... | : 112
        000000000 00000000 00000000 00004040 | ..... | : 128
        40404040 40400000 006ac614 09024040 | ...|F... | : 144
        40404040 40400903 40404040 40404040 | .. | : 160
```

```

00000000 00000000 00000000 00000000 | .....| : 176
00000000 00000000 00000000 00000000 | .....| : 192
00000000 00000000 00000000 00000000 | .....| : 208
00000000 00000000 00000000 00000000 | .....| : 224
00000000 00000000 00000000 00000000 | .....| : 240
00000100 0000c9d4 e2f14040 4040c8e6 | .....IMS1 HW| : 256
e2e8f6d5 f4d70000 00000000 00000000 |SY6N4P.....| : 272
00000000 00000000 00000000 00000000 | .....| : 288
00000000 00004040 40404040 40401040 | .....| : 304
00000000 00004040 40404040 40400000 | .....| : 320
00000300 00000000 00000000 0000e3d4 | .....TM| : 336
d9c1e3d7 f4400000 00000000 00000000 |RATP4 .....| : 352
00000000 00000000 00000000 00000000 | .....| : 368
00000000 00000000 00000000 00000000 | .....| : 384
00000000 00000000 00000000 00000000 | .....| : 400
00000000 00000000 00000000 00000000 | .....| : 416
00000000 00000000 00000000 00000000 | .....| : 432
00000000 00000000 00000000 00000000 | .....| : 448
00000000 00000000 00000000 00000000 | .....| : 464
00000000 00000000 00000000 00000000 | .....| : 480
00000000 00000000 00000000 00000000 | .....| : 496
0000
]
[7/5/11 14:24:04:984 PDT] 00000013 RAIMSTM
< com.ibm.ims.ico.IMSTCPIPAdapter send(byte[]) Exiting method.

```

For more information about how to work with trace, see IBM Documentation for WebSphere Application Server.

Related information

[Troubleshooting and support](#)

IMS TM resource adapter messages and exceptions

IMS TM resource adapter messages are Java exceptions that are thrown by the IMS TM resource adapter itself or by class libraries that it uses (such as the Java class libraries), and have the prefix of ICO.

This topic provides information about exceptions generated by IMS TM resource adapter J2C applications.

The following terms, in *italics* in the message descriptions that follow, are replaced by specific values at run time.

hostname

The TCP/IP host name of the system that is running IMS Connect.

java_exception

The Java exception that is thrown.

length

The length of the data.

libraryFileName

The Local Option native library file name.

llvalue

The value of LL.

maxlength

The maximum valid length of the data.

methodname

The name of the method that is throwing this exception.

mode

The type of interaction between the IMS TM resource adapter and the IMS Connect component on the host (as defined in the `interactionSpec` object).

nativeMethodName

The Local Option native method name.

portnumber

The port number that is assigned to IMS Connect.

propertyname

The name of the property.

propertyvalue

The value of the property.

reasoncode

The reason code that is returned by IMS Connect.

For OTMA sense codes, a reason code of 0 indicates that the sense code does not have an associated reason code.

rectype

The type of the record.

returncode

The return code, formatted in decimal, that is returned by IMS Connect.

sensecode

The sense code, formatted in decimal, that is returned from IMS OTMA

socketexception

The socket exception.

source_exception

The exception thrown when the error first occurred in an internal method.

source_methodname

The internal method in which the error first occurred.

state

The internal state of the IMS TM resource adapter.

Exceptions generated by IMS TM resource adapter J2C applications

The following exception messages are produced by applications built with the Java EE Connector Architecture class libraries when an error condition is detected.

ICO0001E **javax.resource.spi.EISSystemException: ICO0001E: *methodname* error.**
IMS Connect returned an error: RETCODE=[*returncode*], REASONCODE=[*reasoncode*]. *reasoncode_string*.

Explanation

IMS Connect returned an error. The connection that is in error will not be reused. *reasoncode_string* provides a brief description of the *reasoncode*, if available.

When IMS TM Resource Adapter Version 12 or later is used with IMS Version 12 or later, RACF return codes and failure descriptions are returned. With older version of the resource adapter or IMS, one single return code and reason code for all RACF authentication errors are returned.

User response

Check the z/OS console for associated IMS Connect error messages. IMS Connect error messages begin

with the prefix HWS. For diagnostic information about the return code (*returncode*) and reason code (*reasoncode*) values, as well as IMS Connect error messages, see the integrated IMS Connect return and reason codes section in *IMS Messages and Codes*.

ICO0002E **javax.resource.spi.EISSystemException: ICO0002E: *methodname* error.**
IMS OTMA returned an error: SENSECODE=[*sensecode*], REASONCODE=[*otmareasoncode*]. [*source_methodname:source_exception*]

Explanation

IMS OTMA returned a negative acknowledgment (NAK) error.

User response

The IMS TM resource adapter displays *sensecode* and *otmareasoncode* in decimal. A reason code of 0 indicates that the sense code does not have

an associated reason code. If the application is running with two-phase commit, you might receive the following sense code values with the NAK error.

Table 121. Sense code values if the application is running with two-phase commit

Sense code	Description
17 (decimal, 11 hex)	Your IMS is not enabled with resource recovery services (RRS) processing. Ensure your IMS has Protected Conversation processing with RRS enabled. See the topics on global transactions and two-phase commit support for more information.
46 (decimal, 2E hex)	RRS and two-phase commit processing is not supported by IMS Connect and the IMS TM resource adapter.
51 (decimal, 33 hex)	If an authorization failure occurs when an asynchronous output message or a callout request is being retrieved from the hold queue, the OTMA reason code might be one of the following values: • 1 – A security segment is required. • 2 – A user ID is required. • 3 – A RACF group name is required. • 4 – A UTOKEN is required. • 5 – A TPIPE name is required. • 6 – A system error occurred in RACF. • 7 – The RACF profile for the USERID is not in RACF database. • 8 – An authorized user ID is required. • 9 – A valid TPIPE name is required.

For diagnostic information about the sense code (*sensecode*) and OTMA reason code (*otmareasoncode*) values of the NAK error, see the *IMS Messages and Codes* information.

Related concepts

[Global transaction and two-phase commit support process](#)

A Java EE-compliant application server uses a Java transaction manager to communicate and coordinate among the application components and the resource managers.

IC00003E	javax.resource.spi.CommException: IC00003E:methodname error. Failed to connect to host [hostname],

port [portnumber].
[java_exception]

Explanation

The IMS TM resource adapter was unable to connect to the host and port combination. *java_exception* indicates the reason for the failure to connect.

User response

Examine *java_exception* to determine the reason for the failure to connect to the host. Some values for *java_exception* are:

Table 122. Exceptions for IC00003E

Exception	Description
java.net .UnknownHost Exception: <i>hostname</i>	The host name that you specified when you configured the connection factory that is used by your application is invalid, or your application specified an invalid host name. Check the spelling of the host name. You might need to use the fully qualified path for the host name or the IP address.
java.net .ConnectException : Connection refused	Some possible reasons for the exception are: • The port number is invalid. Ensure that you are using a valid port number for the IMS Connect that is indicated by <i>hostname</i>. • The specified port is stopped. Check the status of the port by using the IMS Connect command VIEWHWS. If the port is stopped, its status is NOT ACTIVE. To start the port, use the IMS Connect command, OPENPORT dddd, where <i>dddd</i> is the specified port number. • IMS Connect on the specified host is not running. Start IMS Connect on the host system. • TCP/IP was restarted without canceling and restarting IMS Connect or issuing STOPPORT followed by OPENPORT on the host.

Table 122. Exceptions for ICO0003E (continued)

Exception	Description
java.net .SocketException: connect (code=10051)	Some possible reasons for the exception are: • The system with the specified host name is unreachable on the Internet Protocol network. Verify that the host system is accessible from the Internet Protocol network by issuing the ping command to the specified host system. Enter the ping command on the system on which the IMS TM resource adapter is running. Start TCP/IP on the host, if it is not started. • TCP/IP was restarted but the status of the port that is used by the application was NOT ACTIVE. To correct this situation, take one of the following actions: – Use the IMS Connect command OPENPORT dddd, where <i>dddd</i> is the port number, to activate the port. – Restart IMS Connect.

ICO0005E **javax.resource.spi.CommException:**
ICO0005E:methodname error.
A communication error occurred during sending or receiving the IMS message.
clientID=[clientid][java_exception]

Explanation

The IMS TM resource adapter was unable to successfully complete a send and receive interaction with the target IMS Connect. The *clientid* in the message is the client ID of the connection on which the communication exception occurred. *java_exception* indicates the reason for the failure to complete the interaction.

User response

Use the client ID to help analyze the trace data from the different components that are involved. Examine *java_exception* to determine the reason for the failure. Some values for *java_exception* are described in the following table.

Table 123. Java exceptions for ICO0005E

Java exceptions	Description
java.io .EOFException	Some possible reasons for the exception are: • The timeout value that is specified in the IMS Connect configuration member is exceeded before IMS Connect receives a response from IMS. Exceeding a timeout value typically occurs when a region is not available in IMS to run the IMS transaction that processes the client request. If so, ensure that an appropriate region is started and available to process the request. Exceeding a timeout value can also occur if the IMS application program that is associated with the transaction is stopped. If so, use the IMS command /START PROGRAM to start the IMS application program. • A Java client tries to use a previously active client (for example, a connection from the pool) for which an IMS Connect STOPCLNT command has been issued.

Table 123. Java exceptions for ICO0005E (continued)

Java exceptions	Description
java.net .SocketException: Connection reset by peer: socket write error	Some possible reasons for the exception are: - A Java client attempts to use a connection for which the underlying socket is no longer connected to IMS Connect. The socket connection might be lost if IMS Connect is recycled, but the application server is not. After IMS Connect is restarted, the connections that were formerly successfully connected to IMS Connect are still in the connection pool. As clients attempt to reuse each of these connections, the exception java.net.SocketException is thrown, and the connection object is removed from the connection pool. You can change this behavior in WebSphere Application Server by setting the purge policy of the connection factory that is used by the Java application to the entire pool. - TCP/IP on the host is coming down.

ICO0006E	javax.resource.ResourceException : ICO0006E:methodname error. The value for <i>DataStoreName</i> is null or an empty string.

Explanation

The method indicated in *methodname* was invoked using an empty *DataStoreName* parameter. This error message is displayed in the trace log when a connection factory with an empty *DataStoreName* parameter is started. This message is followed by a Java EE Connector warning:

```
J2CA0007W: An exception occurred while invoking
method setDataStoreName on
com.ibm.connector2.ims.ico
.IMSManagedConnectionFactory used by resource
Connection_Factory_JNDI_name.
```

Processing will then continue leading to other error messages after IMS Connect sends a response indicating that a data store with a null name cannot be found. The underlying message which triggers the other messages is:

```
javax.resource.spi.EISSystemException: ICO0001E:
com.ibm.connector2.ims.ico
.IMSTCPIPManagedConnection@.processOutputOTMAMsg
(byte[], InteractionSpec, Record) error.
IMS Connect returned error: RETCODE=[4],
REASONCODE=[NFNDDST ].
Datastore not found.
```

When this error occurs, a corresponding HWSS0742W warning message is displayed on the z/OS console of the host system where IMS Connect is running. This HWSS0742W message includes a field showing the data store name that it attempted to find, in this case all blanks:

```
DESTID=          ;
```

User response:

Provide a valid name for the *DataStoreName* parameter. In a managed environment, the *DataStoreName* is specified when you configure a connection factory to be used by WebSphere Application Server. In a non-managed environment, the *DataStoreName* is specified in your Java application.

ICO0007E	javax.resource.NotSupportedException: ICO0007E:methodname error. The [propertyname] property value [propertyvalue] is not supported.

Explanation:

The value *propertyvalue* specified for the property *propertyname* is not supported.

User response:

Provide a supported value for the named property. For example, certain values of the *interactionVerb* property of the *InteractionSpec* class that are defined in the J2C architecture are not supported by the *IMSInteractionSpec* class in the IMS TM resource adapter. Also, a value of true for the *reRoute* property is not supported on dedicated persistent socket connections.

ICO0008E	javax.resource.ResourceException : ICO0008E:methodname error. The value [propertyvalue] of the [propertyname] property exceeds the maximum allowable length of [maxpropertylength].

Explanation:

The length of the value *propertyvalue* that is specified for the property *propertyname* exceeds *maxpropertylength*, the maximum length that is allowed for the value of property *propertyname*.

User response:

For the named property, provide a value that does not exceed *maxpropertylength*.

IC00009E **javax.resource.ResourceException**
:
IC00009E:methodname error.
The [propertyname] property value
[propertyvalue] is not valid.

Explanation:

The value *propertyvalue* specified for the property *propertyname* is not valid.

User response:

Provide a valid value for the named property. For example, valid values for the *interactionVerb* property of the *InteractionSpec* class of IMS TM resource adapter are listed in the Javadoc documentation for the *IMSInteractionSpec* class.

IC00010E **javax.resource.spi.IllegalStateException:**
IC00010E:methodname error.
The method was invoked on an
invalid IMSConnection instance.

Explanation:

The method indicated in *methodname* was invoked on an invalid *IMSConnection* instance. If the *methodname* is *lazyEnlist*, an attempt was made to enlist a connection in the current transaction that could not be enlisted.

User response

The named method was most likely issued on an *IMSConnection* instance that was already closed.

- If the *methodname* is not *lazyEnlist*, ensure that the *IMSConnection* instance is not already closed before you attempt to use it or close it.
- If the *methodname* is *lazyEnlist*, ensure that your application is not using non-managed connections in a managed environment, because only managed connections are eligible for lazy transaction enlistment optimization. For more information, please refer to IBM Documentation for WebSphere Application Server.

Related information

Lazy transaction enlistment optimizationFor more information about lazy transaction enlistment optimization, see IBM Documentation for WebSphere Application Server.

IC00011E **javax.resource.spi.IllegalStateException:**
IC00011E:methodname error.
The method was invoked on an
invalid IMSInteraction instance.

Explanation:

The method indicated in *methodname* was invoked on an invalid *IMSInteraction* instance.

User response:

The named method was most likely issued on an *IMSInteraction* instance that was already closed. Ensure that the *IMSInteraction* instance is not already closed before you attempt to use it or close it.

IC00012E **javax.resource.ResourceException**
:
IC00012E:methodname error.
The value provided for HostName
is null or an empty string.

Explanation:

The method indicated in *methodname* was invoked with a null or empty **HostName** parameter.

User response:

Provide a valid value for the **HostName** parameter. In a managed environment, the property value is specified when you are configuring a Connection Factory to be used by WebSphere Application Server. In a non-managed environment, the property value is specified in your Java application.

IC00013E **javax.resource.ResourceException**
:
IC00013E:methodname error.
The ConnectionManager is null.

Explanation:

The application server invoked the *createConnectionFactory* method of the *IMSManagedConnectionFactory* class with a null *ConnectionManager* object.

User response:

This form of the *createConnectionFactory* method is typically used in a managed environment rather than invoked by a client program. Contact the service personnel for your application server.

IC00014E **javax.resource.ResourceException**
:
IC00014E:methodname error.
The input record contains no data.

Explanation:

The method that is indicated in *methodname* was invoked with an input record that contained no data.

User response:

Verify that the input record that you provide is not empty.

IC00015E **ResourceAdapterInternalException**
IC00015E: methodname error.

An unexpected error occurred while the OTMA message was being processed.
[*java_exception*]

Explanation:

An unexpected internal error was encountered when the OTMA message was being processed. The local Java client was unable to locate code page 1047 or 037.

User response:

Ensure that the local Java client has code pages 1047 and 037 installed and that they are not corrupted.

IC00016E **javax.resource.ResourceException**
:
IC00016E:methodname error.
The message was encoded using an unsupported code page.
[*java_exception*].

Explanation:

The message cannot be encoded using the specified code page, or the code page was not found.

User response:

Ensure that the local Java client has code pages 1047 and 037 installed and that they are not corrupted.

IC00017E **ResourceAdapterInternalException**
IC00017E:methodname error.
Invalid value provided for TraceLevel.

Explanation:

An invalid trace level was specified.

User response:

Specify a valid trace level. Optionally, this exception can be ignored because the default trace level is used for this connection factory. In this case, the connection factory is still usable but the trace level is the default trace level.

IC00018E **javax.resource.ResourceException**
:
IC00018E:methodname error.
The value provided for PortNumber is null.

Explanation:

The method indicated in *methodname* was invoked using a null PortNumber.

User response:

Provide a valid PortNumber parameter. In a managed environment, the property value is specified when you are configuring a Connection Factory to be used by WebSphere Application Server. In a non-managed environment, the property value is specified in your Java application.

IC00020E **javax.resource.ResourceException**
:
IC00020E:methodname error.
Alternate client ID is not supported for interaction [*interactionverb*].

Explanation:

The value of *interactionverb* that is specified with the alternate client ID is not valid.

User response:

Provide a valid *interactionverb* value. The alternate client ID can only be specified for SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT, and SYNC_RECEIVE_CALLOUT.

IC00024E **javax.resource.ResourceException**
:
IC00024E:methodname error.
Invalid segment length (LL) of [*llvalue*] in the input object.
[*java_exception*]

Explanation:

The input message provided by the Java program for the IMS application program contains a value for its segment length (*llvalue*) that is negative, 0, or greater than the number of bytes of data in the message segment.

User response:

Provide the correct value for the segment length of the input message.

IC00025E **javax.resource**
.IllegalArgumentException:
IC00025E:methodname error.
Invalid segment length (LL) of [*llvalue*] in the OTMA message.

Explanation:

The output message provided by the IMS application program contains a value for its segment length (*llvalue*) that is negative, 0, or greater than the number of bytes of data in the message segment. The output message provided by the IMS application program is contained in the OTMA message.

User response:

Ensure that your IMS application program provides valid lengths for the segments of its output message.

IC00026E **javax.resource.ResourceException**
:
IC00026E:methodname error.
An error was encountered while processing the IMS message.
[*source_methodname:source_exception*]

Explanation:

An error occurred when the IMS transaction message is being processed. *source_exception* provides additional information regarding the cause of the error.

User response

Examine *source_exception* for additional information regarding the cause of the error. Suggested actions to take, based on the value of *source_exception* are described in the following table.

Table 124. Source exceptions for ICO0026E

Source exceptions	Description
java.io .IOException	Error preparing input or output record. Ensure that the objects you are providing to the IMS TM resource adapter for use as the IMS transaction input and output are defined properly for the J2C architecture. For example, they must implement the interfaces javax.resource.cci.Record and javax.resource.cci.Streamable.
com.ibm.ims.i co .IMSConnResou rce Exception	The OTMA message containing the IMS transaction output message contained an invalid length field (for example, LLLL was less than or equal to 0). If this error continues to occur after verifying that your IMS application program is returning a valid output message, contact your IBM service representative.
java.lang .IllegalArgum ent Exception	The output message returned from IMS Connect is invalid. Ensure that the release levels of the IMS TM resource adapter and IMS Connect are compatible. For example, this exception occurs if you built a transactional-required EJB application to perform a two-phase commit transaction through TCP/IP by using a newer version of the resource adapter, but you are using an out-of-service version of IMS TM Resource Adapter at run time. To resolve the problem, update to the same version of IMS TM Resource Adapter in the runtime environment.

ICO0027E

javax.resource.ResourceException
:
ICO0027E:methodname error.
The OTMA header of the IMS
output message did not contain a
segment_name segment.

Explanation:

The OTMA header in the IMS transaction output message was not built properly. It did not contain the segment identified by *segment_name*.

User response:

Correct the problem with how the message header is built. An example of a situation in which this error could occur is executing a SYNC_RETRIEVE_ASYNCOUTPUT_WAIT or SYNC_RETRIEVE_ASYNCOUTPUT_NOWAIT interaction that retrieves a message that was queued by a non-IMS TM resource adapter client (for example, a 3270 terminal interaction). Messages that are queued by non-IMS TM resource adapter clients typically do not have the OTMA headers that are required by the IMS TM resource adapter. To resolve the problem, ensure that the IMS TM resource adapter and non-IMS TM resource adapter clients use separate queues for their respective asynchronous output.

ICO0028E

javax.resource.ResourceException
:
ICO0028E:methodname error.
The Prefix flag in the OTMA header
Message-Control Information
segment of the
IMS output message is not valid.

Explanation:

The OTMA header in the IMS transaction output message was not built properly. The Prefix flag in the OTMA header specified which segments are included in the OTMA header of the message, but the total length of the message (which must include both the headers and the message data) is less than the sum of the lengths of the specified headers. The total message length must be equal to or greater than the total length of the headers, depending on whether the message contains any data, but never less than the total length of the headers.

User response:

Correct the problem with how the message header is built. This situation typically is a result of either a user message incorrectly modifying the message header or an internal error in either IMS Connect or the IMS TM resource adapter. If your user exits are being used as shipped, without modification, or you are unable to determine why the headers are not being built correctly, contact IBM Software Support.

ICO0030E

javax.resource.spi

.ApplicationServerInternalException:
ICO0030E:methodname error.
[source_methodname:source_exception]

Explanation:

A run time error or exception was detected in *methodname* during the interaction. *source_methodname:source_exception* indicates where the error or exception that was detected in *methodname* originally occurred and might provide additional information regarding the cause of the error.

User response

Examine *source_exception* for additional information regarding the cause of the error. The actions to be taken depend on the value of *source_methodname:source_exception*. Suggested actions to take, based on the value of *source_methodname:source_exception* are described in the following table.

Table 125. Source exceptions for ICO0030E

Java exceptions	Description
java.lang. .OutOfMemory Error	This error is thrown when the Java Virtual Machine cannot allocate an object because it is out of memory, and no more memory could be made available by the garbage collector. Increase the amount of memory that is available to the virtual machine used by WebSphere Application Server.
java.io. .Interrupted IO Exception	An InterruptedException is thrown to indicate that an input or output transfer has been terminated because the thread performing it was terminated. Investigate reasons why the thread was terminated.
ICO0031E	javax.resource.spi.IllegalStateException: ICO0031E:methodname error. Protocol violation. The interaction verb [interactionverb] is not allowed for the current state [state]. [java_exception]

Explanation:

The interaction attempted by the application resulted in a protocol violation. *[interactionverb]* is

the value of the *interactionVerb* property of the *IMSInteractionSpec* object that was used for the interaction. *[state]* is the current state of the protocol that is used for the interactions between the IMS TM resource adapter and IMS Connect.

User response:

Ensure that you are using an appropriate value for the *interactionVerb* property of the *IMSInteractionSpec* class.

ICO0034E	javax.resource.NotSupportedException: ICO0034E:methodname error. Auto-commit is not supported.

Explanation:

Auto-commit is not supported by the IMS TM resource adapter.

User response:

Ensure that your Java application uses classes and methods that are appropriate for the level of support that is provided by the IMS TM resource adapter.

ICO0035E	javax.resource.NotSupportedException: ICO0035E:methodname error. Local transaction is not supported.

Explanation:

Local transactions are not supported by the IMS TM resource adapter.

User response:

Ensure that your Java application uses classes and methods that are appropriate for the level of support that is provided by the IMS TM resource adapter.

ICO0037E	javax.resource.NotSupportedException: ICO0037E:methodname error. ResultSet is not supported.

Explanation:

ResultSets are not supported by the IMS TM resource adapter.

User response:

Ensure that your Java application uses classes and methods that are appropriate for the level of support that the IMS TM resource adapter provides.

ICO0039E	javax.resource.spi.IllegalStateException: ICO0039E:methodname error. Not in CONNECT state.

Explanation:

The sequence of interactions between the IMS TM resource adapter and IMS Connect is invalid. The current state of the protocol used for the interactions between the IMS TM resource adapter and IMS

Connect is not CONNECT as it needs to be at this point in the interactions.

User response:

This is most likely an error in the IMS TM resource adapter or in IMS Connect. Contact IBM Software Support.

IC00040E	javax.resource.NotSupportedException: IC00040E:methodname error. IMSConnector does not support this version of execute method.

Explanation:

The IMS TM resource adapter does not support the form of the execute method that takes two input parameters and returns an object of type `javax.resource.cci.Record`.

User response

Use the supported form of the execute method in class `IMSInteraction`. The supported form of the execute method has the following signature:

```
boolean execute(InteractionSpec, Record input,
                Record output)
```

IC00041E	javax.resource.ResourceException : IC00041E:methodname error. An invalid interactionSpec [interactionSpec] was specified.

Explanation:

An invalid `InteractionSpec` object was passed to the execute method of class `com.ibm.connector2.ims.ico.IMSInteraction`.

User response:

Ensure that the `InteractionSpec` object that you pass to the execute method of class `com.ibm.connector2.ims.ico.IMSInteraction` is of type `com.ibm.connector2.ims.ico.IMSInteractionSpec`.

IC00042E	javax.resource.ResourceException : IC00042E: methodname error. The input is not of type Streamable.

Explanation

The input that is provided to the execute method of `com.ibm.connector2.ims.ico.IMSInteraction` for the Input parameter either was null or did not implement the interface `javax.resource.cci.Streamable`. This exception occurs most likely when an application is written to use the Java EE Connector Architecture Common Client Interface (CCI). This exception does not occur

if a Rational or WebSphere integrated development environment is used to build the input message.

The execute method allows null input objects for some types of interactions. For example, interactions with `interactionVerb` values of `SYNC_END_CONVERSATION`, `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT`, and `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT` allow null input objects.

User response:

Ensure that you provide a valid `javax.resource.cci.Record` object for the input parameter to the execute method. For example, ensure that this object implements the interfaces `javax.resource.cci.Record` and `javax.resource.cci.Streamable`.

IC00043E	javax.resource.ResourceException : IC00043E: methodname error. The output is not of type Streamable.

Explanation:

The output object provided to the execute method of `com.ibm.connector2.ims.ico.IMSInteraction` either was null or did not implement the interface `javax.resource.cci.Streamable`. This exception occurs most likely when an application is written to use the Java EE Connector Architecture Common Client Interface (CCI). This exception does not occur if a Rational or WebSphere integrated development environment is used to build the output message.

User response:

Ensure that you provide a valid output object to the execute method.

IC00044E	javax.resource.NotSupportedException: IC00044E:methodname error. RecordFactory is not supported by IMS TM Resource Adapter.

Explanation:

`RecordFactory` is not supported by the IMS TM resource adapter.

User response:

Ensure that your Java application uses classes and methods that are appropriate for the level of support that the IMS TM resource adapter provides.

IC00045E	javax.resource.NotSupportedException: IC00045E:methodname error. Invalid type of ConnectionRequestInfo.

Explanation:

An invalid `ConnectionRequestInfo` object was passed to an IMS TM resource adapter method.

User response:

This is most likely an error in the IMS TM resource adapter. Contact IBM Software Support.

IC00049E **javax.resource.NotSupportedException:**
IC00049E:methodname error.
The security credentials passed to getConnection do not match existing security credentials.

Explanation:

The security credentials in the request do not match the security credentials of the IMSManagedConnection instance that is used to process the request.

User response:

Contact IBM Software Support.

IC00050E **IC00050E: methodname error.**
Invalid RACF user id is specified in SSLKeyStoreName or SSLTrustStoreName when specifying a RACF keystore or truststore.

Explanation:

The user ID specified in the SSLKeyStoreName or the SSLTrustStoreName property for the RACF keystore or truststore is not valid.

User response:

Specify a valid user ID and ensure that the RACF user ID is less than 8 characters long.

IC00053E **javax.resource.ResourceException**
:
IC00053E: methodname error.
The client ID is not value. The prefix HWS is reserved by the IMS TM resource adapter.

Explanation:

The value specified for the clientID property is not valid. The prefix HWS is reserved by the IMS TM resource adapter.

User response

Provide a valid value for the clientID property. A valid value follows the following rules:

- Is not a null string
- Does not start with a blank field
- Does not start with the IMS TM resource adapter reserved prefix HWS
- Is 8 characters long
- Uses valid characters: A - Z, 0 - 9, @, #, and \$

IC00054E **javax.resource.ResourceException**
:

IC00054E:methodname error.

Invalid ConnectionSpec.

Explanation:

The IMS TM resource adapter was unable to cast the provided connectionSpec for this connection to type IMSConnectionSpec. Although the Common Client Interface accepts a connectionSpec object for any supported connector, the IMS TM resource adapter works only with an IMSConnectionSpec or a derivative of IMSConnectionSpec as its connectionSpec.

User response:

Ensure that the connectionSpec used by your application is an IMSConnectionSpec or inherits from IMSConnectionSpec.

IC00055E **javax.resource.ResourceException**
:
IC00055E:methodname error.
Failed to cast the connection object to IMSConnection.

Explanation:

The IMS TM resource adapter was unable to cast the connection object allocated by the ConnectionManager for this connection to type IMSConnection. The IMS TM resource adapter works only with an IMSConnection or a derivative of IMSConnection as its connection object. This error might be the result of a problem with the ConnectionManager instance.

User response:

Contact IBM Software Support.

IC00056E **javax.resource.ResourceException**
:
IC00056E:methodname error.
IMSConnectName is only valid for Local Option connections which can only be used in z/OS or OS/390.

Explanation:

Setting the IMSConnectName property of an IMSManagedConnectionFactory instance is required for managed connection factory to be used for Local Option connections. Also, you can use only Local Option to communicate with IMS Connect if the application that is using the IMS TM resource adapter is running on z/OS. This exception indicates that you have specified a value for the IMSConnectName property but your application is not running on z/OS.

User response:

Ensure that your application that is using the IMS TM resource adapter is running on z/OS. The Web application server where your application is running must be running in the same z/OS image as IMS Connect. If, for example, you plan to run your

application on a workstation platform or if the Web server where you plan to run your application is on z/OS but in a different z/OS image from IMS Connect, ensure that the connection factory that is used by your application is set up to use TCP/IP communication.

IC00057E **javax.resource.spi.IllegalStateException:**
IC00057E:methodname error.
Invoked with invalid connection handle.

Explanation:

The application is in an illegal state. The connection handle (IMSConnection instance) that is used for this interaction is not valid. The application might be attempting to use a connection handle that is either for a previously used connection or for the wrong connection if the application has more than one connection open.

User response:

Ensure that the application is using the valid IMSConnection instance for that connection.

IC00058E **javax.resource.ResourceException**
:
IC00058E:methodname error.
Interactions
SYNC_SEND_RECEIVE,
SYNC_SEND,
SYNC_RECEIVE_ASYNCOUTPUT,
SYNC_RECEIVE_ASYNCOUTPUT
_SINGLE_NOWAIT and
SYNC_RECEIVE_ASYNCOUTPUT
_SINGLE_WAIT interactions with
Commit Mode 0
are not supported with Local
Option.

Explanation:

User response:

Ensure that your application is selected with commit mode 1. If you plan to run your application with commit mode 0, correct your application to use TCP/IP communication.

IC00059E **javax.resource.ResourceException**
:
IC00059E: methodname error.
SYNC_END_CONVERSATION
interaction with Commit Mode 0 is
not supported.

Explanation:

Interaction SYNC_END_CONVERSATION with commit mode 0 is not supported.

User response:

The IMS TM resource adapter supports the interaction combination

SYNC_END_CONVERSATION with commit mode 1, SYNC_SEND_RECEIVE with commit mode 0, and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT with commit mode 0.

IC00060E **java.lang.UnsatisfiedLinkError:**
IC00060E:methodname error.
Error loading Local Option native
library: libname=libraryfilename.
[source_exception].

Explanation:

The Local Option native library named *libraryfilename* cannot be found in any of the directories listed in the LIBPATH environment variable.

User response:

Ensure that the Local Option native library exists in one of the directories in the LIBPATH environment variable. If the IMS TM resource adapter is running on WebSphere Application Server for z/OS, ensure that the full name of the directory that contains the Local Option native library file is defined in the LIBPATH environment variable for your Java EE server. For more information about how to prepare the base operation system, see IBM Documentation for WebSphere Application Server.

Related information

Preparing the base z/OS operating systemFor more information, see IBM Documentation for WebSphere Application Server.

IC00061E **javax.resource.ResourceException**
:
IC00061E:methodname error.
Local Option runs only in z/OS and
OS/390.

Explanation:

You can use Local Option to communicate with IMS Connect only if your application using the IMS TM resource adapter is running on z/OS.

User response:

Ensure that your application using the IMS TM resource adapter is running on z/OS. The Web application server where your application is running must be running in the same z/OS image as IMS Connect. If you plan to run your application on a workstation platform or if the Web server where you plan to run your application is on z/OS but in a different z/OS image than IMS Connect, ensure that the connection factory used by your application is set up to use TCP/IP communication.

IC00062E **javax.resource.ResourceException**
:
IC00062E:methodname error.

Error loading Local Option native method:
libfilename=libraryFileName,
methodname=nativeMethodName.
[source_exception].

Explanation:

The Local Option native method named *nativemethodname* cannot be found in the Local Option native library file named *libraryfilename*.

User response:

Verify that you have the correct level of the IMS TM resource adapter and Local Option native library installed on your system. Always use the version of the Local Option native library that shipped with the IMS TM resource adapter that you installed in your WebSphere Application Server for z/OS.

IC00063E **javax.resource.spi**
.ResourceAdapterInternalException:
IC00063E:methodname error.
Exception thrown in native
method. [source_exception].

Explanation:

An internal error occurred in the Local Option native method.

User response:

Contact IBM Software Support.

IC00064E **javax.resource.spi.SecurityException:**
IC00064E:methodname error.
Invalid security credential.

Explanation:

WebSphere Application Server did not provide a security credential that is supported by the IMS TM resource adapter.

User response:

Ensure that you have the correct level of WebSphere Application Server for z/OS installed. Configure WebSphere Application Server for z/OS to provide a security credential that is supported by IMS TM resource adapter (the PasswordCredential for TCP/IP connections and the UToken GenericCredential for Local Option connections).

IC00065E **javax.resource.spi.SecurityException:**
IC00065E:methodname error.
Error obtaining credential data
from the security credential.
[source_exception].

Explanation:

A security-related error occurred in obtaining the credential data from the security credential provided by the application server.

User response:

Ensure that you correctly set up security for your application server so that the user associated with the calling program is authorized to extract the data from a security credential.

IC00066E **javax.resource.ResourceException**
:
IC00066E:methodname error.
Error loading WebSphere
Application Server
Transaction Manager.
[source_exception].

Explanation:

An error occurred when accessing the transaction manager of the WebSphere Application Server for processing the transaction request.

User response:

Ensure that you have the correct level of WebSphere Application Server for z/OS installed.

IC00068E **javax.resource.ResourceException**
:
IC00068E:methodname error.
Error obtaining the transaction
object. [java_exception]

Explanation:

An error occurred in determining if a transaction has been started using the WebSphere Application Server for z/OS transaction manager.

User response:

Ensure that you have the correct level of WebSphere Application Server for z/OS installed.

IC00069E **javax.resource.spi**
.ResourceAllocationException
IC00069E:methodname error.
Error obtaining RRS transaction
context token.
IMSConnResourceException: RRS
retcode=[rrs_routinecode].

Explanation:

An unexpected internal error occurred in obtaining a Recovery Resource Services (RRS) transaction context token for processing the global transaction.

User response:

Check the RRS job log for associated RRS error messages. For diagnostic information about the RRS return code (*rrs_routinecode*) see *z/OS V1R9 MVS Programming: Resource Recovery* or the same manual for your release of z/OS.

Related information

[RRS return code in Resource recovery services \(RRS\) in z/OS MVS Programming: Callable Services for High-Level Languages](#) For more information, see the [Resource recovery services](#)

(RRS) return code information in z/OS MVS Programming: Callable Services for High-Level Languages.

IC00070E **javax.resource.spi.EISSystemException**
IC00070E:methodname error.
IMS Connect reported an RRS error: IMS Connect Return Code=[returncode], RRS Routine name=[rrs_routine], RRS Return code=[rrs_routinecode].

Explanation:

IMS Connect returned an error that resulted from a Recovery Resource Services (RRS) failure.

User response:

Check the z/OS console for associated IMS Connect and RRS error messages. For diagnostic information about the return code (*returncode*) value, as well as IMS Connect error messages, see the *IMS Messages and Codes* information. For diagnostic information about the RRS return code (*rrs_routinecode*) locate the RRS routine name (*rrs_routine*) within *z/OS V1R9 MVS Programming: Resource Recovery* or the same manual for your release of z/OS.

Related information

RRS return code in [Resource recovery services \(RRS\) in z/OS MVS Programming: Callable Services for High-Level Languages](#)For more information, see the [Resource recovery services \(RRS\) return code information in z/OS MVS Programming: Callable Services for High-Level Languages](#).

IC00071E **javax.transaction.xa.xaException**
IC00071E:methodname error.
A communication error occurred when processing the XA commandtype operation.
[java_exception]

Explanation:

A communication failure during the processing of a global transaction can occur when a TCP/IP or socket failure occurs or when IMS Connect is down. The connection in error will not be reused.

User response

- Examine the Java exceptions to determine the reason for the failure to connect to the host.
- Check the z/OS console for associated IMS Connect or TCP/IP error messages.
- Validate that system can be reached through TCP/IP and that IMS Connect has not been brought down.

The command type (*commandtype*) in the error message refers to the stage at which this communication failure occurred during the global transaction: prepare, commit, rollback, recover, or forget.

IC00072E **javax.transaction.xa.xaException:**
IC00072E:methodname error.
The associated UR for the Xid is not found.

Explanation:

During transaction processing a unit of recovery (UR) that was tied to a specific global transaction ID (XID) was eliminated by manual intervention or an error in IMS Connect or Recovery Resource Services (RRS).

User response:

See IBM Documentation for WebSphere Application Server for steps on how to acquire transaction information and Xids within the WebSphere Application Server logs. Refer to the *IMS Version 14 Commands, Volume 3* for IMS commands that list the XIDs and their associated UR. Verify that a UR is listed for that XID. Verify that the global transaction was not left in a state where the transactions need manual completion.

Related information

[IBM Documentation for WebSphere Application Server 8.5.5](#)For more information about resolving transaction information and Xids, see [IBM Documentation for WebSphere Application Server 8.5.5](#).

IC00073E **javax.transaction.xa.xaException:**
IC00073E:methodname error.
RRS is not available.

Explanation:

Recovery Resource Services (RRS) has been brought down, or communication between RRS and IMS Connect has ended.

User response:

Check the z/OS console for associated IMS Connect and RRS error messages. Ensure that RRS has not been brought down on your z/OS system. Refer to *IMS Version 14 Communications and Connections* for IMS Connect commands that can be used to verify that RRS is enabled.

Related information

RRS return code in [Resource recovery services \(RRS\) in z/OS MVS Programming: Callable Services for High-Level Languages](#)For more information, see the [Resource recovery services \(RRS\) return code information in z/OS MVS Programming: Callable Services for High-Level Languages](#).

[IBM Documentation for WebSphere Application Server 8.5.5](#)For more information about

resolving transaction information and Xids, see IBM Documentation for WebSphere Application Server 8.5.5.

IC00074E **javax.transaction.xa.xaException:**
IC00074E: The RRS *rrs_routine*
call returns with a return code
[*rrs_routinecode*].

Explanation:

During the processing of your global transaction, the Recovery Resource Services (RRS) error message was passed by IMS Connect.

User response

Check the z/OS console for associated IMS Connect and RRS error messages. For diagnostic information on the RRS return code (*rrs_routinecode*) locate the RRS routine name (*rrs_routine*) within *z/OS V1R9 MVS Programming: Resource Recovery*, or the same manual for the version of your release of z/OS.

Related information

RRS return code in Resource recovery services (RRS) in *z/OS MVS Programming: Callable Services for High-Level Languages* For more information, see the Resource recovery services (RRS) return code information in *z/OS MVS Programming: Callable Services for High-Level Languages*.

IC00075E **javax.transaction.xa.xaException:**
IC00075E: *methodname* error.
The transaction branch might have
been heuristically completed.
[*rrs_exception*]

Explanation:

An Recovery Resource Services (RRS) error was passed by IMS Connect that indicates that the processing of your transaction might be left in a state where part of the transaction was committed and part of it encountered an error during the commit phase. The *rrs_exception* is an IC00074E error message that indicates the RRS routine and return code that are associated with this issue.

User response

Check the z/OS console for associated IMS Connect and RRS error messages. For diagnostic information on the RRS return code (*rrs_routinecode*) locate the RRS routine name (*rrs_routine*) within *z/OS V1R9 MVS Programming: Resource Recovery*, or the same manual for the version of your release of z/OS.

See IBM Documentation for WebSphere Application Server for steps on how to acquire transaction information and global transaction ID (XID) within

theWebSphere Application Server logs. Refer to *IMS Version 14 Communications and Connections* for IMS Connect commands that list the XIDs and their associated units of recovery (URs). Determine the XIDs and URs that are involved and the result that should have been committed to IMS. Verify values within IMS to determine that a heuristic state, where the completion of the transaction cannot be determined, has occurred. Decide the action to take to rectify the data within IMS so that it matches the result that would have been committed, or to rectify the other databases that are involved to return to a state before the execution of that transaction.

Related information

RRS return code in Resource recovery services (RRS) in *z/OS MVS Programming: Callable Services for High-Level Languages* For more information, see the Resource recovery services (RRS) return code information in *z/OS MVS Programming: Callable Services for High-Level Languages*.

IBM Documentation for WebSphere Application Server 8.5.5 For more information about resolving transaction information and Xids, see IBM Documentation for WebSphere Application Server 8.5.5.

IC00076E **javax.resource.ResourceException**
:
IC00076E: *methodname* error.
An internal error occurred.
[*rrs_exception*]

Explanation:

An internal error occurred in extracting information about an Recovery Resource Services (RRS) error message from IMS Connect. The *rrs_exception* is an IC00074E error message indicating the RRS routine and return code associated with the error.

User response

Check the z/OS console for associated IMS Connect and RRS error messages. For diagnostic information on the RRS return code (*rrs_routinecode*) locate the RRS routine name (*rrs_routine*) within *z/OS V1R9 MVS Programming: Resource Recovery*, or the same manual for the version of your release of z/OS.

Related information

RRS return code in Resource recovery services (RRS) in *z/OS MVS Programming: Callable Services for High-Level Languages* For more information, see the Resource recovery services (RRS) return code information in *z/OS MVS Programming: Callable Services for High-Level Languages*.

ICO0077E	javax.resource.ResourceException : ICO0077E:methodname error. The transaction has already rolled back. [rrs_exception]

Explanation:

An Recovery Resource Services (RRS) error has been passed by IMS Connect, which indicates that the attempt to roll back a transaction has been made a second time on the same unit of recovery (UR). RRS prevents the second rollback from taking place and throws an error indicating that such an action is being attempted. The *rrs_exception* is an ICO0074E error message indicating the RRS routine and return code associated with the error.

User response:

No action is needed as the transaction should be rolled back. Refer to the documentation for ICO0074E for more information about the RRS failure that has taken place. As a precaution, verify that data was not lost or modified before the execution of the transaction.

Related information

[ICO0074E](#)

[javax.transaction.xa.XAException:](#)

ICO0074E: The RRS *rrs_routine* call returns with a return code [*rrs_routinecode*].

ICO0078E	javax.resource.ResourceException : ICO0078E: methodname error. A valid user-specified clientID is required for interactions on a dedicated persistent connection.

Explanation:

A valid, user-specified value is required for the *clientID* property when the commit mode is 0, and the interaction is using a dedicated persistent socket connection. This requirement applies to SYNC_SEND_RECEIVE, SYNC_SEND, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions.

User response

Provide a valid value for the *clientID* property. A valid value must follow the following rules:

- Is not a null string
- Does not begin with a blank field
- Does not begin with the IMS TM resource adapter reserved prefix HWS
- Is eight characters long
- Has valid characters A - Z, 0 - 9, and @, #, \$

ICO0079E	com.ibm.connector2.ims.ico

.IMSDFSMessageException:
ICO0079E:methodname error.
IMS returned a DFS
message:DFS_message

Explanation

IMS returned a DFS message (*DFS_message*) instead of the output of the IMS transaction. This exception is thrown if the interaction uses the value IMS_REQUEST_TYPE_IMS_TRANSACTION for the *imsRequestType* property.

For example, if the Java application attempts to run an IMS transaction that is stopped, this exception is thrown and the value of *DFS_message* might be:

```
DFS064 hh:mm:ss DESTINATION CAN NOT BE FOUND OR
        CREATED, DEST=
DFS065 hh:mm:ss TRAN/LTERM STOPPED
```

User response:

Find the explanation and response that corresponds to *DFS_message* in the *IMS Messages and Codes* documentation, and then address the problem in IMS.

ICO0080E	javax.resource.spi.EISSystemException: ICO0080E:methodname error. Execution timeout has occurred for this interaction. The executionTimeout was [executionTimeout_value] milliseconds. The IMS Connect TIMEOUT was used.

Explanation

The time it took for IMS Connect to send a message to IMS and receive the response (*executionTimeout_value*) was greater than the IMS Connect TIMEOUT value. The IMS Connect TIMEOUT value is:

- Specified in the IMS Connect configuration member for SYNC_SEND_RECEIVE interactions
- 2 seconds for SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT, and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions

The IMS Connect TIMEOUT value was used because the *executionTimeout* property for this interaction was not specified or has been set to zero.

User response:

Ensure that your application has set a valid *executionTimeout* value. To set the *executionTimeout* values, you can use either a WebSphere or Rational integrated development

environment, or use the `setExecutionTimeout` method.

IC00081E `javax.resource.spi.EISSystemException:`
IC00081E:methodname error.
 Execution timeout has occurred for this interaction. The `executionTimeout` value specified was `[executionTimeout_value]` milliseconds. The value used by IMS Connect was `[rounded_executionTimeout_value]` milliseconds.

Explanation:

The amount of time that it took for IMS Connect to send a message to IMS and receive the response (`executionTimeout_value`) was greater than the `executionTimeout` value that was rounded to an appropriate execution timeout interval, `rounded_executionTimeout_value`. After a valid execution timeout value is set, this value is converted into a value that IMS Connect can use.

User response

If the rounded execution timeout value is not what you expected, verify that the value follows the conversion rules:

Table 126. Execution timeout value conversion rules

Range of user-specified values	Conversion rule
1 - 250	If the user-specified value is not divisible by 10, it is converted to the next greater increment of 10.
251 - 1000	If the user-specified value is not divisible by 50, it is converted to the next greater increment of 50.
1001 - 60000	The user-specified value is converted to the nearest increment of 1000. Values that are exactly between increments of 1000 are converted to the next greater increment of 1000.
60001 - 3 600000	The user-specified value is converted to the nearest increment of 60000. Values that are exactly between increments of 60000 are converted to the next greater increment of 60000.

IC00082E `javax.resource.NotSupportedException:`
IC00082E:methodname error.
 Execution timeout has occurred for this interaction. The `executionTimeout` value of `[{executionTimeout_value}]` milliseconds is not supported. The valid range is `[{executionTimeout_waitforever_flag}, 0 to {maximum_executionTimeout_value}]` milliseconds. The IMS Connect TIMEOUT was used.

Explanation:

The execution timeout value that was specified for the `executionTimeout` property did not meet the requirements for the minimum or maximum timeout values.

User response:

Ensure that your application has set a valid value for the `executionTimeout` property. The execution timeout value is represented in milliseconds and must be a decimal integer in the range of 1 to 3600000, inclusively. You can also set the value to -1 if you want an interaction to run without a time limit.

Related reference

[Execution timeout \(executionTimeout\)](#)

This property specifies the maximum amount of time that is allowed for IMS Connect to send a message to IMS and receive a response

IC00083E `javax.resource.ResourceException:`
IC00083E:methodname error.
SYNC_SEND_RECEIVE, SYNC_SEND, SYNC_RECEIVE_ASYNCOUTPUT, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions with **Commit Mode 0** are not valid within the scope of a global transaction.

Explanation:

`SYNC_SEND_RECEIVE, SYNC_SEND, SYNC_RECEIVE_ASYNCOUTPUT` (deprecated), `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT` (replacing `SYNC_RECEIVE_ASYNCOUTPUT`) and `SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT` interactions with commit mode 0 are not valid within the scope of a global transaction. The global transaction requires commit mode 1 interactions with a sync level of `SYNC_LEVEL_SYNCPOINT`.

User response

- If you want to use commit mode 0, ensure that your application is configured as a non-transactional application.
- If you want to run your interactions within the scope of a global transaction, the `commitMode` property value must be 1.

ICO0084E	javax.resource.ResourceException : ICO0084E:methodname error. An unexpected internal IMS TM Resource Adapter error occurred. [source_method] [source_exception]

Explanation:

A `PrivilegedActionException` occurred in executing a `[source_method]` call in `methodname`. This exception occurs if Java EE security is enabled and the user associated with the calling program, `methodname`, or any program in the current call stack is not authorized to execute `[source_method]`.

User response:

Ensure that you have correctly set up security for your application server so that the user associated with the calling program, along with any programs in the current call stack at the time of the exception, are authorized to execute `[source_method]`. Alternatively, turn off Java EE security checking in the application server.

ICO0085E	javax.resource.ResourceException : ICO0085E: methodname error. Protocol violation. A user-specified clientID is not allowed for interactions on a shareable persistent socket.

Explanation:

The value specified for the `clientID` property is not allowed. Because the connection factory is configured for a shareable persistent socket, a user-specified client ID is not allowed within this connection factory.

User response:

For a shareable persistent socket connection factory, the IMS TM resource adapter provides a generated client ID. User-specified client ID is not allowed. To determine if you are using a shareable persistent socket, check for a value of `false` for the `CMODedicated` property of the connection factory that is used by the interaction.

Related reference

[Client ID \(clientID\)](#)

A value for this property must be provided when an application component uses a dedicated

persistent socket connection to IMS Connect to identify this particular client.

ICO0086E	javax.resource.ResourceException : ICO0086E:methodname error. Invalid value was specified for CommitMode property.

Explanation:

The value that you specified in the `commitMode` property field is invalid.

User response

Ensure that your application has set a valid value for the `commitMode` property. Values supported are:

- 1 (`SEND_THEN_COMMIT`) indicates that IMS processes the transaction and sends a response back before committing the data.
- 0 (`COMMIT_THEN_SEND`) indicates that IMS processes the transaction and commits the data before sending a response.

Related reference

[Commit mode \(commitMode\)](#)

This property indicates the type of commit mode processing to be performed for an IMS transaction.

ICO0087E	javax.resource.ResourceException : ICO0087E: methodname error. Protocol violation. commit mode 1 is not allowed for interactions on a dedicated persistent socket.

Explanation:

The value of 1 that is specified for the `commitMode` property is invalid. Because the connection factory is configured for a dedicated persistent socket, commit mode 1 is not allowed.

User response:

For a dedicated persistent socket connection factory, only commit mode 0 interactions are valid. To determine if you are using a dedicated persistent socket, check for a value of `true` for the `CMODedicated` property of the connection factory that is used by the interaction.

Related reference

[CMO dedicated \(CMODedicated\)](#)

This optional property applies to TCP/IP connections only. The default value is `false`, which indicates that the connection factory will generate shareable persistent socket

connections that can be used by commit mode 0 and commit mode 1 interactions.

IC00088E **javax.resource.ResourceException**
:
IC00088E: *methodname* error.
Protocol violation.
SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT interactions are not allowed on a shareable persistent socket.

Explanation:

The value SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT that is specified for the interactionVerb property is not valid. Because the connection factory is configured for a shareable persistent socket, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT is not allowed.

User response:

Determine if you are using a shareable persistent connection by checking for a value of false for the CM0Dedicated property of the connection factory that is used by the interaction. Set the interaction to SYNC_SEND_RECEIVE, SYNC_SEND, and SYNC_END_CONVERSATION for shareable persistent connections.

IC00089I **javax.resource.ResourceException**
:
IC00089I: *methodname*.
Non-persistent socket closed for Commit Mode 0 IMS transaction.

Explanation:

Running commit mode 0 with non-persistent socket (transaction socket), the IMS TM resource adapter forces the removal of managed connection object from the connection pool.

User response:

No action is required.

IC00091E **javax.resource.ResourceException**
:
IC00091E: *methodname* error.SSL client context could not be created. [{1}]

Explanation

An SSL context could not be created due to one of the following reasons:

- The algorithm that checks the integrity of the keystore cannot be found.
- The certificates in the keystore could not be loaded.
- The key cannot be recovered (for example, the specified password is wrong).

User response

Ensure that:

- The algorithm that creates the certificates is supported by the IBM Java Secure Socket Extension provider (IBMJSSE).
- The passwords for the keystore and truststore are correct.

IC00096I **javax.resource.ResourceException**
:
IC00096I: *methodname* Warning. Invalid value provided for SSL parameter.

Explanation:

The method indicated in *methodname* was invoked by using a null or empty SSLKeystoreName, SSLKeystorePassword, SSLTruststoreName, or SSLTruststorePassword parameter. This information message does not terminate the program execution.

User response:

Provide valid values for SSLKeystoreName, SSLKeystorePassword, SSLTruststoreName and SSLTruststorePassword parameters. For convenience, you can store private keys and certificates in either a keystore or a truststore. Only one set of valid values (either SSLKeystoreName and SSLKeystorePassword or SSLTruststoreName and SSLTruststorePassword) are required for proper execution.

IC00097E **javax.resource.ResourceException**
:
IC00097E: *methodname* error.
The given value is invalid for 'SSLEncryptionType'. The value must be 'STRONG' for strong encryption, 'WEAK' for weak (export) encryption or 'ENULL' for null (no) encryption.

Explanation:

A value other than STRONG, WEAK, or ENULL was provided for the SSLEncryptionType parameter.

User response:

Specify STRONG, WEAK, or ENULL for the SSLEncryptionType parameter. The value is not case-sensitive.

Related reference

[SSL encryption type \(SSLEncryptionType\)](#)

This optional property applies to TCP/IP connections only, and specifies the SSL encryption type. The SSLEnabled property must be set to true.

IC00111E **javax.resource.ResourceException**
:

ICO0111E:methodname error.
SSLEnabled must be set to FALSE
when using Local Option.

Explanation:

The property SSLEnabled is set to true, and the property IMSConnectName is set to a non-null value, which indicates that local option connections are used. However, SSL is not supported on local option connections.

User response:

Set the SSLEnabled property to false.

ICO0112E **ICO0112E:methodname error.**
Connection is closed due to
transaction timeout.

Explanation:

The connection is closed due to a WebSphere Application Server transaction timeout.

User response:

Allow more time for your transaction to complete by increasing the WebSphere Application Server transaction timeout value.

ICO0113E **javax.resource.spi.CommException:**
ICO0113E: methodname error.
Socket timeout has occurred
for this interaction. The socket
timeout value
specified was
[socket_timeout_value]
milliseconds.
[source_exception:exception_reason]

Explanation:

The execution timeout value for the IMS TM resource adapter to receive a response from IMS Connect is greater than the value that is specified for socket timeout.

User response:

Ensure that the value of the socketTimeout property is sufficient for the IMS TM resource adapter to receive a response from IMS Connect. If it is not, increase the value. If the value of the socketTimeout property given is sufficient, network problems might be causing delays. Contact your network administrator.

Related reference

Socket timeout (socketTimeout)
The maximum amount of time the IMS TM resource adapter waits for a response from IMS Connect before disconnecting the socket and returning an exception to the client application.

ICO0114E **javax.resource.ResourceException**
:

ICO0114E: methodname error.
The socket timeout value
of [socket_timeout_value]
milliseconds is not valid.
[source_exception:exception_reason]

Explanation:

The value *socket_timeout_value* that was specified for the socketTimeout property is not valid.

User response:

Review the exception reason provided. Ensure that the socketTimeout property is set to a positive numeric value.

Related reference

Socket timeout (socketTimeout)
The maximum amount of time the IMS TM resource adapter waits for a response from IMS Connect before disconnecting the socket and returning an exception to the client application.

ICO0115E **javax.resource.spi.CommException:**
ICO0115E: methodname error.
A TCP Error occurred.

Explanation:

An error occurred in the underlying network protocol.

User response:

Contact your network administrator.

ICO0116E **ICO0116E:methodname error. A**
Common Client Interface error
occurred.

Explanation:

An error occurred in the underlying protocol.

User response:

Contact IBM Software Support.

ICO0117E **javax.resource.ResourceException**
:
ICO0117E: methodname error.
Protocol violation: Commit Mode 1
is not allowed for SYNC_SEND,
SYNC_RECEIVE_ASYNCOUTPUT,
SYNC_RECEIVE_ASYNCOUTPUT
_SINGLE_NOWAIT
and
SYNC_RECEIVE_ASYNCOUTPUT
_SINGLE_WAIT interactions.

Explanation:

The IMS TM resource adapter supports only commit mode 0 for SYNC_SEND interactions.

User response:

Correct either the commit mode or the interaction verb. Commit mode 0 is required for SYNC_SEND,

SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions. Commit mode 1 is valid with SYNC_SEND_RECEIVE and SYNC_END_CONVERSATION interactions.

Related concepts

[Interactions supported with commit mode and sync level combinations](#)

The interactions that are supported by the IMS TM resource adapter depend on the combination of commit mode and sync level that you choose.

IC00118E **javax.resource.ResourceException**
:
IC00118E: *methodname* error. Protocol violation. IMS request type 2 (IMS_REQUEST_TYPE_IMS_COMMAND) is not allowed for SYNC_SEND, SYNC_END_CONVERSATION, SYNC_RECEIVE_ASYNCOUTPUT, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions.

Explanation:

The value 2 (IMS_REQUEST_TYPE_IMS_COMMAND) that was specified for the `imsRequestType` property is invalid.

User response:

The value of 2 (IMS_REQUEST_TYPE_IMS_COMMAND) for the `ImsRequestType` property is valid only with the SYNC_SEND_RECEIVE interaction. The value of 1 (IMS_REQUEST_TYPE_IMS_TRANSACTION) is required for SYNC_SEND, SYNC_END_CONVERSATION, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions.

IC00119E **javax.resource.ResourceException**
:
IC00119E: *methodname* error. A supported SSL provider was not found. [caught_exception]

Explanation:

When attempting to initialize a Secure Sockets Layer (SSL) TCP/IP connection with IMS Connect, the IMS TM resource adapter needs to use one of the two supported providers: `com.ibm.jsse.JSSEProvider` or `sun.security.provider.Sun`. This error indicates that neither of these providers is available.

User response:

The provider `com.ibm.jsse.JSSEProvider` is added by default in an IBM JVM. The provider `sun.security.provider.Sun` is added by default in a Sun JVM. Ensure that you are running the IMS TM resource adapter in a supported IBM JVM if you are running in WebSphere Application Server or a Sun JVM in other application servers.

IC00121E **javax.resource.ResourceException**
:
IC00121E: *methodname* error. Invalid reRoute name value. Prefix HWS is reserved for use by IMS TM Resource Adapter.

Explanation:

The value for the `reRouteName` property is invalid. The prefix HWS is reserved for use only by the IMS TM resource adapter.

User response:

Provide a valid value for the `reRouteName` property. See the `reRouteName` property information for more information.

Related reference

[Reroute name \(reRouteName\)](#)

This property provides the name of the destination to which asynchronous output is queued.

IC00122E **javax.resource.ResourceException**
:
IC00122E: *methodname* error. Invalid reRoute value. When purgeAsyncOutput value is true, reRoute value cannot be true.

Explanation:

The value for the `reRoute` property is invalid because the value for the `purgeAsyncOutput` property is set to `true`, or the default value (`true`) is used for the `purgeAsyncOutput` property.

User response:

Set the `purgeAsyncOutput` property to `false` if you want to set the `reRoute` property to `true`.

Related reference

[Purge async output \(purgeAsyncOutput\)](#)

This input property determines whether or not IMS Connect purges undelivered output.

[Reroute \(reRoute\)](#)

This input property determines whether undelivered output is to be rerouted to a named destination that is specified in the `reRouteName` field.

IC00123E **javax.resource.NotSupportedException:**

**IC00123E : *methodname* error .
A Sync Level value of *synclevel*
is not supported for commit-then-
send
(Commit mode 0) interactions.**

Explanation:

The value specified for the `syncLevel` property is not supported for commit-then-send (commit mode 0) interactions . The value 1 (confirm) is the only supported value of the `syncLevel` property for commit-then-send (commit mode 0) interactions.

User response:

Specify the value of 1 (confirm) for the `syncLevel` property, or accept the default sync level value 1 for commit-then-send (commit mode 0) interactions.

Related reference

Sync level (`syncLevel`)

This input property specifies the synchronization level of the interaction between the IMS TM resource adapter and IMS OTMA.

**IC00124E `javax.resource.ResourceException`
:
**IC00124E : *methodname* error .
SYNC_SEND_RECEIVE interactions
with Sync Level Confirm are not
supported
with Local Option.****

Explanation:

Commit-then-send (commit mode 0) interactions are not supported with Local Option. Send-then-commit (commit mode 1) interactions with a sync level of confirm (1) are not supported with Local Option.

User response:

Do not use Local Option for this functionality.

**IC00125E `javax.resource.EISSystemException`
**on:
IC00125E : *methodname* error.
An internal error occurred. The
status of the IMS transaction
associated
with this SYNC_SEND_RECEIVE
interaction with Commit Mode 1
and Sync Level 1
cannot be determined.****

Explanation:

As part of the send-then-commit (commit mode 1) sync level confirm (1) protocol, the IMS TM resource adapter received an output message from the IMS transaction, sent a positive acknowledgment to IMS Connect, but received an error indication instead of the expected indication from IMS Connect that the transaction was committed. The output message that was received by the IMS TM resource adapter was

not returned to the application component. The IMS transaction might or might not have been committed.

User response:

Determine whether the in-flight transaction was committed or aborted by consulting the type X'01' and X'03' IMS log records on the target IMS system. Perform reconciliation between the state of the IMS system and the external application if necessary.

**IC00126E `javax.resource.ResourceException`
:
**IC00126E : *methodname* error .
IMS Connect reported a Commit
Mode 1 Sync Level Confirm error:
IMS
Connect Return
Code=[*returncode*], Reason
Code=[*reasoncode*].
*reasoncode_string*****

Explanation:

IMS Connect returned an error. The connection in error will not be reused. *reasoncode_string* provides a brief description of the reason code, if available.

User response:

Check the z/OS console for associated IMS Connect error messages. IMS Connect error messages begin with the characters HWS.

**IC00127E `com.ibm.ims.ico`
**.IMSIllegalStateException:
IC00127E : *methodname* error.
Protocol violation. The Mode
[*mode*] is not allowed for the
current state
[*state*]. [*java_exception*]****

Explanation:

The IMS TM resource adapter is in an illegal state.

User response:

This is most likely an error in the IMS TM resource adapter or IMS Connect. Contact your IBM service representative.

**IC00128E `javax.resource.NotSupportedException`
**IC00128E : *methodname* error.
The Sync Level property value of
[*synclevel*] given is invalid. Sync
Level
NONE (0) and Sync Level CONFIRM
(1) are the only values supported
by the
setSyncLevel(int) method.****

Explanation:

A value other than 0 or 1 was specified as input to the `setSyncLevel(int)` method.

User response:

Specify either 0 or 1 as a sync level property value input to the `setSyncLevel(int)` method.

Related reference

[Sync level \(syncLevel\)](#)

This input property specifies the synchronization level of the interaction between the IMS TM resource adapter and IMS OTMA.

IC00129E	javax.resource.ResourceException : IC00129E:methodname error. Specifying the Alternate ClientID is not allowed on a dedicated persistent connection. The Alternate ClientID value is supported on shareable persistent socket connections ONLY.

Explanation:

The retrieval of asynchronous output using an alternate client ID is supported only on shareable persistent socket connections.

User response:

Use a shareable persistent socket connection to retrieve asynchronous output with an alternate client ID.

Related reference

[Alternate client ID \(altClientID\)](#)

This input-only property provides the name of an alternate client ID. This name is used as a tpipe from which asynchronous output is retrieved on a shareable persistent socket connection.

IC00130E	javax.resource.ResourceException : IC00130E:methodname error. The ignorePURGCall property is not allowed on a dedicated persistent connection.

Explanation:

The ignorePURGCall property is not supported on a dedicated persistent connection for SYNC_SEND_RECEIVE interactions. An exception is thrown if you try to set this property on a dedicated persistent connection.

User response:

Either use a shareable persistent socket connection for the interaction (preferred solution) or do not set the ignorePURGCall property if your application must use a dedicated persistent connection.

Related reference

[Ignore PURG call \(ignorePURGCall\)](#)

This property controls whether or not multiple ISRT and PURG calls in an IMS application for a commit mode 0 interaction is to be ignored.

IC00131E	javax.resource.ResourceException : IC00130E:methodname error. The property ignorePURGCall is not supported for interaction verb [SYNC_END_CONVERSATION or SYNC_RECEIVE_ASYNCOUTPUT_*]. The property ignorePURGCall can only be specified for interaction SYNC_SEND and SYNC_SEND_RECEIVE.

Explanation:

The ignorePURGCall property is not supported for SYNC_END_CONVERSATION, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT, and SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions.

User response:

Use the SYNC_SEND or SYNC_SEND_RECEIVE interaction on a shareable persistent socket connection. If your intent is to execute a SYNC_END_CONVERSATION, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT, or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interaction, do not specify a value of true for the ignorePURGCall property.

Related reference

[Ignore PURG call \(ignorePURGCall\)](#)

This property controls whether or not multiple ISRT and PURG calls in an IMS application for a commit mode 0 interaction is to be ignored.

IC00132E	IMS Connect returned an error during a conversational transaction: RETCODE={1}, REASONCODE={2}. {3}

Explanation:

IMS Connect returned an error return code and reason code during a conversational transaction.

User response:

Refer to the *IMS Messages and Codes* documentation and take the appropriate action that is suggested for the specific return code and reason code.

IC00133E	The interactionVerb must be either SYNC_SEND_RECEIVE or SYNC_END_CONVERSATION when the value of the IMSInteractionSpec property useConvID is true (client-managed

conversation state programming model).

Explanation:

The interactionVerb property of the IMSInteractionSpec object is not set to SYNC_SEND_RECEIVE or SYNC_END_CONVERSATION.

User response:

Use either SYNC_SEND_RECEIVE or SYNC_END_CONVERSATION for the interaction verb setting of the IMSInteractionSpec object.

Related reference

Interaction verb (interactionVerb)
Specifies the mode of interaction between the Java application and IMS. The following table describes the values that are supported by the IMS TM resource adapter.

IC00134E	When the value of IMSInteractionSpec property useConvID is true (client-managed conversation state programming model), conversational transactions are supported on shareable persistent socket connections ONLY.

Explanation:

Conversational transactions that use the conversation ID to track the iterations of a conversation (by setting the IMSInteractionSpec useConvID property to true) are supported only on shareable persistent socket connections.

User response:

Use shareable persistent socket connections with the client-managed conversation state programming model.

IC00135E	When the value of IMSInteractionSpec property useConvID is true (client-managed conversation state programming model), conversational transactions are supported with Commit Mode 1 ONLY.

Explanation:

Conversational transactions that use the conversation ID to track the iterations of a conversation (by setting the IMSInteractionSpec useConvID property to true) are supported only for commit mode 1.

User response:

Use commit mode 1 with the client-managed conversation state programming model.

IC00136I	The duplicate [queue] name <i>name</i> was removed. The list of [queue] names was reset to <i>new_list</i>.

Explanation:

Duplicate queue or data store names are not allowed.

User response:

This message is informational, and no user action is required.

IC00137I	The maximum length of the property <i>property_name</i> was exceeded. <i>old_property_value</i> has been truncated to <i>old_property_value</i>.

Explanation:

The maximum length of this property is eight characters.

User response:

This message is informational, and no action is required.

IC00138E	The value for <i>property_name</i> is invalid. The prefix HWS is reserved for use by IMS TM Resource Adapter.

Explanation:

A reserved prefix, HWS, was used in the *property_name* value.

User response:

Provide a value for the *property_name* property that does not use a reserved prefix.

IC00139E	The message endpoint threw an exception. <i>endpoint_exception</i>

Explanation:

A Java exception (*endpoint_exception*) was passed to the IMS TM resource adapter by the message-driven bean (MDB).

User response:

Resolve the exception that is passed to the IMS TM resource adapter by the MDB.

IC00140I	Reconnecting in <i>number_of_seconds</i> seconds...

User response:

This message is informational, and no user action is required.

IC00141E	The endpoint is already active.

Explanation:

An attempt was made to add an endpoint that already exists.

User response:

Do not attempt to add this endpoint until it has been removed.

IC00142E The endpoint is not found.

Explanation:

An attempt was made to access an endpoint that does not exist.

User response:

Do not attempt to remove this endpoint.

IC00143E An unexpected error occurred when the message was dispatched from IMS to the message endpoint. *java_exception*

User response:

Contact IBM Software Support and provide the Java exception that is thrown.

IC00144E The IMS TM resource adapter encountered an error while sending the message to IMS.

Explanation:

An unexpected internal error occurred when the message was sent to IMS.

User response:

Contact IBM Software Support.

IC00145E The IMS TM resource adapter encountered an error while sending the response message for a synchronous callout request to IMS. *java_exception*

Explanation:

The maximum length of this property is eight characters.

User response:

Examine the Java exception to determine the cause for the failure to send the message. Ensure that the queue is not in an invalid state by issuing the **/DISPLAY TMEMBER TPIPE** *queue_name* SYNC command in IMS.

IC00146E An error occurred when the error message for a synchronous

**callout request was back to IMS.
*java_exception***

Explanation:

This message is the result of an encoding or message type error generated within a message that is passed back to IMS during a synchronous callout request.

User response:

Examine the documentation for the Java exception or error message to determine why the original message generated an error.

IC00147E No message listeners were started. There are no queues on which to listen.

User response:

Specify at least one valid queue name in the queueNames property of the IMSActivationSpec class.

IC00148E The ActivationSpec *activationSpec_name* is not supported.

Explanation:

An unsupported ActivationSpec was passed to the method.

User response:

Use the supported ActivationSpec, which is IMSActivationSpec.

IC00150E The IMS TM Resource Adapter is unable to listen for callout messages from IMS on queue *queue_name* because it is unable to connect to host *host_name*, port *port_number*.

Explanation:

The TM Resource Adapter did not receive a response from a host at the specified address.

User response:

Ensure that you have specified a valid host name and port number in the ActivationSpec.

Other exceptions and error messages

You might encounter errors from IMS Connect, or from WebSphere Application Server, as it catches the error messages and exceptions that are thrown by the IMS TM resource adapter or the class libraries used by the IMS TM resource adapter.

J2CA0056I

Exceptions that the IMS TM resource adapter throws can be caught by a component other than your Java application, such as the WebSphere Application Server.

For example, when WebSphere Application Server catches an exception from your Java application, it might then issue its own message, including in it the message from the IMS TM resource adapter

exception. For example, when execution timeout occurs, you see the following message on the WebSphere Application Server console:

```
J2CA0056I: The Connection Manager received a fatal connection error
from the
Resource Adaptor for resource myConnFactory. The exception which
was received is
IC00080E:
com.ibm.connector2.ims.ico.IMSTCPIPManagedConnection@e59583c.
processOutputOTMAMsg(byte[],IMSInteractionSpec, int) error.
Execution timeout has occurred for this interaction.
The executionTimeout was [0] milliseconds. The IMS Connect TIMEOUT
was used.
```

J2CA0056I is an informational message from WebSphere Application Server. The fatal connection error occurs because IMS Connect closes the socket when an execution times out, and the WebSphere Application Server Connection Manager removes the connection object for the socket from the connection pool.

Another example occurs when a transaction (non-persistent) socket is used for a commit mode 0 interaction. In this case, you see the following message on the WebSphere Application Server console:

```
J2CA0056I: The Connection Manager received a fatal connection error
from the
Resource Adaptor for resource myConnFactory. The exception which
was received is IC00089I:
com.ibm.connector2.ims.ico.IMSTCPIPManagedConnection@6db5d83a.call(Connection,
InteractionSpec, Record, Record). Non-persistent socket closed for
Commit Mode 0 IMS transaction.
```

Related information

[WebSphere Application Server Version 8 documentation](#) See the WebSphere Application Server documentation for more troubleshooting information.

WLTC0017E

WLTC0017E is produced by the WebSphere Transaction Monitor to indicate that the resources enlisted with a local transaction containment (LTC) were rolled back instead of committed.

A LTC is used to define the application server behavior in an unspecified transaction context. For example, if a single method within a container-managed enterprise Java bean (EJB) that has a transaction attribute of NotSupported is called outside of any transaction scope, the WebSphere Application Server creates a local transaction to handle resources that are used during the execution of that method.

Message WLTC0017E indicates that the resources that were enlisted with the LTC were rolled back instead of committed because the `setRollbackOnly()` method was called on the LTC. This message does not require any action by the user and is for your information only.

WLTC0017E: Resources rolled back due to `setRollbackOnly()` being called.

The prefix of a WebSphere Application Server message indicates the component that issued the message. You can find documentation of these messages in the WebSphere Application Server Information Center.

Related information

[WebSphere Application Server Version 8 documentation](#) See the WebSphere Application Server documentation for more troubleshooting information.

HWSP1445E

Message HWSP1445E occurs in IMS Connect when you specify to use Secure Sockets Layer (SSL) for the connection, but provide a non-SSL port for your Java application.

When you configure a connection factory for use by your Java application, you specify whether you are using SSL with the `SSLEnabled` property. If you are using SSL (`SSLEnabled=TRUE`), the port number that you provide must be configured as an SSL port in IMS Connect. If you provide a non-SSL port for your Java application, unexpected results will occur when you run your application.

- The IMS TM resource adapter will throw an exception that indicates a communication error:

```
javax.resource.spi.CommException:
IC00003E:
com.ibm.connector2.ims.ico.IMSTCIPManagedConnection@56503fc6.connect()
error.
Failed to connect to host [CSDMEC13], port [9999].
[java.net.SocketException:
Connection reset by peer: socket closed]
```

- The following IMS Connect message is displayed on the z/OS console:

```
HWSP1445E UNKNOWN EXIT NAME SPECIFIED IN MESSAGE PREFIX; MSGID=
/9 * !hR, M=SDRC
```

The first step in establishing an SSL connection involves the SSL handshake protocol, in which the client (the IMS TM resource adapter) sends the server (IMS Connect) an SSL "Hello" message. In this scenario, IMS Connect is waiting for an incoming message on a non-SSL port. When IMS Connect receives the handshake message, it interprets the message as an OTMA message with a valid Exit name in the prefix and issues message HWSP1445E.

HWSSSL00E

Message HWSSSL00E occurs in IMS Connect when you specify not to use Secure Sockets Layer (SSL) for the connection, but the port number you provide for your Java application is an SSL port.

In this error scenario, the following errors will occur:

- The IMS TM resource adapter will throw an exception that indicates a communication error:

```
javax.resource.spi.CommException: IC00005E:
com.ibm.connector2.ims.ico.IMSTCIPManagedConnection@5bcdcd4.receive()
error. A communication error occurred while sending or receiving
the IMS message.
[java.net.SocketException: Connection reset by peer: socket closed]
```

- The following IMS Connect message is displayed on the z/OS console:

```
HWSSSL00E Unable to initialize the SSL socket:Error while reading
or writing data
```

IMS Connect's attempt to initialize the SSL socket fails because IMS Connect does not receive the initial client "Hello" message that is part of the SSL handshake protocol.

Chapter 52. Reference information

The IMS TM resource adapter has the following classes for managing interactions and connections with IMS, as well as inbound communication from IMS to the Java client application.

IMS connection factory properties

The properties of an IMS TM resource adapter connection factory describes the characteristics of the target Enterprise Information System.

The following list describes these connection properties.

Client ID (clientID)

A value for this property must be provided when an application component uses a dedicated persistent socket connection to IMS Connect to identify this particular client.

The client ID value must be specified to identify the client to which the socket connection is dedicated (when the `CM0Dedicated` property is set to `true`). A dedicated persistent socket connection must be used for only commit mode 0 interactions with the following interaction verbs:

- `SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_WAIT`
- `SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_NOWAIT` (replacing the deprecated `SYNC_RECEIVE_ASYNCOUPTUT`)
- `SYNC_SEND_RECEIVE`

The client ID is generated in a shareable persistent socket connection that can be used by commit mode 0 and commit mode 1 interactions (when the `CM0Dedicated` property is set to `false`).

The client ID in a shareable persistent socket can be obtained through the `IMSConnectionMetaData.getClientID` method. You can access the `IMSConnectionMetaData` class from the `IMSConnection` class. If you use this method with IMS Version 12, APAR PM75222 (V12.1.2) is required.

A valid user-specified client ID must meet the following requirements:

- Must be a string of 1 to 8 alphanumeric (A-Z, 0-9) or special (@,#,\$) characters.
- Must not start with the character string HWS.
- Must not be an IMS Connect port number.
- If lowercase letters are provided, the letters are changed to uppercase

Client IDs are dedicated to their sockets, regardless of whether they are shareable or dedicated. Unlike the other properties, the original client ID never changes until the socket is disconnected.

CM0 dedicated (CM0Dedicated)

This optional property applies to TCP/IP connections only. The default value is `false`, which indicates that the connection factory will generate shareable persistent socket connections that can be used by commit mode 0 and commit mode 1 interactions.

When the value is `false`, the IMS TM resource adapter generates a client ID to identify the socket connection.

A value of `true` indicates that the connection factory will generate dedicated persistent socket connections, which require user-specified client IDs to identify the socket connections. Each dedicated persistent socket connection is reserved for a particular client ID. Only commit mode 0 interactions are permitted on dedicated persistent socket connections.

Data store name (dataStoreName)

This required value is the name of the target IMS data store.

The IMS data store name is the ID parameter of the Datastore statement in the target IMS Connect configuration member. This name also serves as the XCF member name for IMS during internal XCF communications between IMS Connect and IMS OTMA. The value for the dataStoreName property is case sensitive.

Group name (groupName)

This optional value is the IMS group name to use for all connections that are created by this connection factory if the default user name is used.

The groupName property can be provided only in a component-managed EIS signon environment.

Host name (hostName)

You must provide the IP address or host name of the system on which the target IMS Connect is running. The value is ignored for Local Option connections.

Password (password)

This optional value is the password to use for connections that are created by this connection factory if the default user name is used.

Password phrase (passwordPhrase)

The password phrase to be used by the Security Access Facility, such as RACF, to authenticate the user that is identified by the userName property.

V12.2,V13.2, and later This property applies to IMS TM Resource AdapterVersion 12.2.0, Version 13.2.0, and later.

The password phrase must be 9 - 100 characters.

Port number (portNumber)

You must provide the port number that is used by the target IMS Connect for TCP/IP connections.

Multiple sockets can be open on a single TCP/IP port, and IMS Connect can be configured to use multiple ports for communications with the IMS TM resource adapter and other clients.

SSL enabled (SSLEnabled)

This optional property applies to TCP/IP connections only. The default value is `false`, which indicates that SSL sockets are used for connecting to the port that is specified in the portNumber property.

A value of `true` indicates that this connection factory is used to create SSL socket connections to IMS Connect by using the host name and port number that are specified in its connection properties. This port number must be configured as an SSL port in the IMS Connect configuration.

SSL encryption type (SSLEncryptionType)

This optional property applies to TCP/IP connections only, and specifies the SSL encryption type. The SSLEnabled property must be set to `true`.

Valid encryption types are STRONG, WEAK and ENULL. STRONG and WEAK reflect the strength of the ciphers, which is related to the key length. All ciphers that can be used for export are in the weak category and the others are in the strong category.

By default, the encryption type is set to WEAK. When ENULL is specified, the IMS TM resource adapter uses a cipher specification whose name contains the string "NULL." Null encryption enables authentication during the SSL handshaking process. After the handshaking process for a socket completes, which includes the required authentication, all messages flows over that socket.

The `SSLEncryptionType` property is case-insensitive.

SSL keystore name (`SSLKeyStoreName`)

This optional property applies to TCP/IP connections only, and only when the `SSLEnabled` property is set to `true`. The value contains the name, including the full file path, of the keystore.

Private keys and their associated public key certificates are stored in password-protected databases called keystores. For convenience, trusted certificates can also be stored in the keystore. The `SSLKeyStoreName` property can either be empty or could point to the keystore file. If the `SSLKeyStoreName` or the `SSLKeyStorePassword` property is empty, an informational message is generated in the server log.

For non-z/OS platforms, specify the fully-qualified path name of your JKS keystore file. An example of a fully-qualified path name of your JKS keystore file is `c:\keystore\MyKeystore.k`s.

For z/OS, the `SSLKeyStoreName` property can be used to specify either a JKS keystore or a RACF keyring. For a JKS keystore, specify the name with the full path of the JKS keystore file. For a RACF keyring, specify the string that provides the information needed to access the RACF keyring. An example of a RACF keyring is `keystore_type;keyring_name;racfid`.

- `keystore_type` must be one of the following values:
 - **JCERACFKS** if software encryption is used for SSL.
 - **JCE4758RACFKS** if hardware encryption is used.
- `keyring_name` is the name of the RACF keyring that you are using as your keystore.
- `racfid` is a RACF ID that is authorized to access the specified keyring.

The following examples show two RACF keyring specifications:

- `JCERACFKS;myKeyring;kruser01`
- `JCE4758RACFKS;myKeyring;kruser01`

On z/OS, if the `SSLKeyStoreName` matches the RACF keyring format, the IMS TM resource adapter uses the specified RACF keyring as its keystore. If the specified keystore type is anything other than **JCERACFKS** or **JCE4758RACFKS**, the IMS TM resource adapter attempts to interpret the `SSLKeyStoreName` that is specified as the name of a JKS keystore file.

The JKS file can have a file extension other than `KS`.

Related reference

[SSL truststore name \(`SSLTrustStoreName`\)](#)

This optional property applies to TCP/IP connections only, and only when the `SSLEnabled` property is set to `true`. The value contains the name, including the full file path, of the truststore.

SSL keystore password (`SSLKeyStorePassword`)

This optional property applies to TCP/IP connections only and specifies the password for the SSL keystore. The `SSLEnabled` property must be set to `true`.

SSL truststore name (SSLTrustStoreName)

This optional property applies to TCP/IP connections only, and only when the `SSLEnabled` property is set to `true`. The value contains the name, including the full file path, of the truststore.

On z/OS, SSL truststore name is the JKS name or the RACF keyring of the truststore. The same format is used for the values of the `SSLKeyStoreName` and `SSLTrustStoreName` properties. See the description of the `SSLKeyStoreName` property for this format.

For other platforms, specify the fully-qualified path name of your JKS truststore file.

A truststore file is a key database file (keystore) that contains public keys or certificates. For convenience, private keys can also be stored in the truststore. The `SSLKeyStoreName` property can be empty or can point to the truststore file. If the `SSLTrustStoreName` or `SSLTrustStorePassword` property is empty, an informational message is generated in the server log.

The JKS file can have a file extension other than `KS`.

Related reference

[SSL keystore name \(SSLKeyStoreName\)](#)

This optional property applies to TCP/IP connections only, and only when the `SSLEnabled` property is set to `true`. The value contains the name, including the full file path, of the keystore.

SSL truststore password (SSLTrustStorePassword)

This optional property applies to TCP/IP connections only and specifies the password for the SSL truststore. The `SSLEnabled` property must be set to `true`.

SSL protocol (SSLProtocol)

This property specifies the protocol to be used for a secure socket connection. `SSLProtocol` can have a value of `TLS`, `TLSv1.0`, `TLSv1.1`, or `TLSv1.2`. The default value is `TLSv1.0`.

User name (userName)

This optional value is the default security authorization facility (SAF) user name to use for connections that are created by this connection factory if no user name is provided by the application component or the container.

IMS interaction specification properties

The properties of the `IMSInteractionSpec` object describe the interaction with IMS.

Values for `IMSInteractionSpec` properties can be provided directly in the application by using the `set` methods. They can also be provided to the J2C wizard in an IDE that generates code for the application.

The following list describes the properties of the `IMSInteractionSpec` object:

Alternate client ID (altClientID)

This input-only property provides the name of an alternate client ID. This name is used as a `tpipe` from which asynchronous output is retrieved on a shareable persistent socket connection.

This property is supported with the following criteria:

- TCP/IP connection with shareable persistent socket.
- Valid interaction verbs for retrieval of asynchronous output messages (resume `tpipe` requests) are:
 - `SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_NOWAIT`
 - `SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_WAIT`
 - `SYNC_RECEIVE_CALLOUT`

- Commit mode 0 interactions cannot be used with Local Option connections.
- This property cannot be used with dedicated socket connections
- The reroute name and the alternate client ID are mutually exclusive and cannot be specified at the same time.

Async output available (asyncOutputAvailable)

This output-only property is used by a Java application to determine whether queued output exists for the tpipe associated with the connection that is used by the Java application.

For dedicated persistent socket connections, the name of the tpipe is the value in the clientID property the of IMSConnectionSpec object.

For shareable persistent socket connections, the name of the tpipe is generated by the IMS TM resource adapter. The value of asyncOutputAvailable is true if messages exist in the queue. The asyncOutputAvailable property is not set on input by the application component.

If your Java application uses this property, it must be exposed as an output property of the IMSInteractionSpec object in your application.

Callout request type (calloutRequestType)

This property specifies which type of callout request messages the Java application is to retrieve.

Valid values are:

- CALLOUT_REQUEST_ASYNC (default): To retrieve only asynchronous callout request messages.
- CALLOUT_REQUEST_SYNC: To retrieve only synchronous callout request messages.
- CALLOUT_REQUEST_BOTH: To retrieve both synchronous and asynchronous callout request messages.

This property is only valid when the interaction verb is set to SYNC_RECEIVE_CALLOUT.

Conversation ended (convEnded)

This is an output-only property that is used by a Java application to determine whether a conversation has been ended by IMS. The convEnded property is not set on input by the application component.

If your Java application uses this property, it must be exposed as an output property of the IMSInteractionSpec object in your application.

Conversation ID (convID)

This property is an 8-byte time stamp (a hexadecimal string) that uniquely identifies an IMS conversation.

When the useConvID property is set to true:

- The IMS TM resource adapter returns this conversation ID to the Java client as part of the output message for each iteration of a conversation.
- The conversation ID is assigned by IMS and set in the response message header at the end of the first iteration of the conversation.
- The Java client is responsible for passing this conversation ID in the input for all subsequent iterations in order for IMS Connect to process all iterations of a conversation that comes in from shareable persistent socket connections.

Restriction: This property is supported for only SYNC_SEND_RECEIVE and SYNC_END_CONVERSATION interaction verbs. It is ignored for other interaction verbs. For restrictions that apply to the use of conversation ID, see the topic on "Conversational transactions for business process choreography."

Related concepts

[Business process choreography applications](#)

You can run IMS conversational transactions in composite business applications by using IMS TM Resource Adapter.

Commit mode (commitMode)

This property indicates the type of commit mode processing to be performed for an IMS transaction.

- The commitMode property can be set to 0 or 1 when the interactionVerb property is set to SYNC_SEND_RECEIVE.
- When the interactionVerb property is set to SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT, SYNC_RECEIVE_CALLOUT, or SYNC_SEND, the IMS TM resource adapter uses commit mode 0.
- Commit mode 1 is required when the interactionVerb property is set to SYNC_END_CONVERSATION.

If the commit mode is 0 and a shareable persistent socket is used for the interaction (when the CMODedicated property of the connection factory is set to `false`), the clientID property must not be specified. If commit mode 0 is specified for an interaction on a shareable persistent socket, the output message from a transaction can be purged or rerouted. The undelivered secondary output from a program-to-program switch can also be purged or rerouted.

If a dedicated persistent socket connection is used for an interaction (when the CMODedicated property of the connection factory is set to `true`), the commit mode must be 0 and the client ID of the IMSConnectionSpec object that is used for the connection must be provided. If a dedicated persistent socket is used for a commit mode 0 interaction, undelivered output messages are always recoverable and cannot be purged or rerouted.

Related reference

[Sync level \(syncLevel\)](#)

This input property specifies the synchronization level of the interaction between the IMS TM resource adapter and IMS OTMA.

CMO response (CMOResponse)

This property indicates whether a DFS2082 message should be issued for a CMO input transaction if the IMS application do not reply to the IOPCB or do a message switch to another transaction.

By default, a CMO transaction does not receive a DFS2082 informational message if the IMS application do not reply to the IOPCB or do a message switch to another transaction. To request for a DFS2082 message, set this property to true.

Execution timeout (executionTimeout)

This property specifies the maximum amount of time that is allowed for IMS Connect to send a message to IMS and receive a response

- The executionTimeout value is represented in milliseconds and must be a decimal integer that is either -2, -1, 0, or between 1 and 3600000 (one hour), inclusively.
- If a value of 0 is set for this property, the actual timeout value is determined by IMS Connect.
- If a -1 value is set for this property, the interaction runs without a time limit.
- If a value of -2 is set for this property, the callout request message is retrieved by using the OTMA RESUME TPIPE call with SINGLE option and a timer of NOWAIT. If no callout request message exists on the OTMA hold queue, the IMS TM resource adapter does not wait for the callout request and immediately returns the control back to the Java application.

Related reference

[Valid execution timeout values](#)

The execution timeout value is represented in milliseconds and must be a decimal integer from 1 through 3600000.

[Socket timeout \(socketTimeout\)](#)

The maximum amount of time the IMS TM resource adapter waits for a response from IMS Connect before disconnecting the socket and returning an exception to the client application.

Ignore PURG call (ignorePURGCall)

This property controls whether or not multiple ISRT and PURG calls in an IMS application for a commit mode 0 interaction is to be ignored.

When an IMS application has multiple ISRT and PURG calls:

- If the client application uses commit mode 1 interaction, IMS sends one response message with multiple output segments.
- If the client application uses commit mode 0:
 - If the ignorePURGCall property is set to `false` (the default), IMS sends multiple output response messages, one for each PURG call. The client application receives the first output message, and the remaining output messages remain in the asynchronous hold queue for that client. The message can then be retrieved by using the SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT, or SYNC_RECEIVE_CALLOUT interaction.
 - If the ignorePURGCall property is set to `true`, the multiple PURG calls in the IMS application are ignored. The output is returned as one single message with multiple segments.

The ignorePURGCall flag is valid for the SYNC_SEND and SYNC_SEND_RECEIVE interactions for commit mode 0 interactions on shareable persistent socket connections. Multiple PURG calls are always ignored, regardless of the value of this flag for commit mode 1 interactions.

The ignorePURGCall flag is not valid for SYNC_END_CONVERSATION, SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT or SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT interactions, or for SYNC_SEND_RECEIVE interactions on dedicated persistent socket connections. An exception is thrown in these interactions.

IMS request type (imsRequestType)

This property indicates the type of IMS request and determines how output from the request is handled by the IMS TM resource adapter.

The following table describes the valid IMS request types.

Table 127. IMS request types

Value	Named constant in IMSInteractionSpecProperties	Description
1	IMS_REQUEST_TYPE_IMS_TRANSACTION	The request is an IMS transaction. Normal transaction output that is returned by IMS is used to populate the application's output message. If IMS returns a DFS message, the IMS TM resource adapter throws an IMSDFSMessageException that contains the DFS message. This value for the imsRequestType property is used for applications that are not generated from MFS message definitions through message input descriptor (MID) and message output descriptor (MOD) names.
2	IMS_REQUEST_TYPE_IMS_COMMAND	The request is an IMS command. Command output that is returned by IMS, including DFS messages, is used to populate the application's output message. The IMSDFSMessageException is not thrown. This value for the imsRequestType property is used for applications that submit IMS commands.
3	IMS_REQUEST_TYPE_MFS_TRANSACTION	The request is from an application that is generated from MFS message definitions through message input descriptor (MID) and message output descriptor (MOD) names. Normal transaction output that is returned by IMS, as well as DFS messages, are used to populate the output message from the application. The IMSDFSMessageException is not thrown.

Interaction verb (interactionVerb)

Specifies the mode of interaction between the Java application and IMS. The following table describes the values that are supported by the IMS TM resource adapter.

Table 128. Mode of interaction between the Java application and IMS

Value	Named constant in <code>IMSInteractionSpecProperties</code>	Description
0	<code>SYNC_SEND</code>	The IMS TM resource adapter sends the client request to IMS through IMS Connect and does not expect a response from IMS. With a <code>SYNC_SEND</code> interaction, the client does not need to synchronously receive a response from IMS. <code>SYNC_SEND</code> is supported on both shareable and dedicated persistent socket connections and is allowed only with commit mode 0 interactions. If the <code>interactionVerb</code> property is set to <code>SYNC_SEND</code>, execution timeout and socket timeout values are ignored. Restriction: Type 2 of IMS request (<code>IMS_REQUEST_TYPE_IMS_COMMAND</code>) is not allowed with the <code>SYNC_SEND</code> interaction and will generate an exception.
1	<code>SYNC_SEND_RECEIVE</code>	The execution of an IMS Interaction sends a request to IMS and receives a response synchronously. A typical <code>SYNC_SEND_RECEIVE</code> interaction involves a non-conversational IMS transaction in which an input record (the IMS transaction input message) is sent to IMS, and an output record (the IMS transaction output message) is returned by IMS. <code>SYNC_SEND_RECEIVE</code> interactions are also used for the iterations of a conversational IMS transaction. A conversational transaction requires commit mode 1. A non-conversational transaction can run using either commit mode 1 or commit mode 0. If commit mode 0 is used on a dedicated persistent socket, a value for the <code>clientID</code> property of <code>IMSConnectionSpec</code> must be provided. If commit mode 0 is used on a shareable persistent socket, a value for the <code>clientID</code> property of <code>IMSConnectionSpec</code> must not be provided.
3	<code>SYNC_END_CONVERSATION</code>	If the application executes an interaction with the <code>interactionVerb</code> property set to <code>SYNC_END_CONVERSATION</code>, the IMS TM resource adapter sends a message to force the end of an IMS conversational transaction. The <code>commitMode</code> property of the <code>IMSInteractionSpec</code> class and the <code>clientID</code> property of the <code>IMSConnectionSpec</code> class do not apply when <code>SYNC_END_CONVERSATION</code> is specified as the interaction verb.
4	<code>SYNC_RECEIVE_ASYNCOUTPUT</code>	The interaction verb <code>SYNC_RECEIVE_ASYNCOUTPUT</code> has been replaced by the more specific <code>SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT</code>. <code>SYNC_RECEIVE_ASYNCOUTPUT</code> is supported for backward compatibility. New applications must use either <code>SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT</code> or <code>SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_WAIT</code>.
5	<code>SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT</code>	The interaction verb <code>SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT</code> is valid on both shareable and dedicated persistent socket connections. It is used to retrieve asynchronous output. A <code>SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT</code> interaction on a shareable persistent socket connection must be in the same application as the original <code>SYNC_SEND</code> or <code>SYNC_SEND_RECEIVE</code> interaction and must use the same shareable persistent connection. This primarily occurs following execution timeout. With this type of interaction, the Java client can receive only one single message. If no messages are in the IMS OTMA asynchronous hold queue for the client ID when the request is made, no further attempts are made to retrieve the message. No message is returned and a timeout occurs after the length of time specified in the <code>executionTimeout</code> property of the <code>SYNC_RECEIVE_ASYNCOUTPUT_SINGLE_NOWAIT</code> interaction passes.

Table 128. Mode of interaction between the Java application and IMS (continued)

Value	Named constant in IMSInteractionSpecProperties	Description
6	SYNC_RECEIVE_ ASYNCOUPTUT_ SINGLE_WAIT	The interaction verb SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_WAIT is used to retrieve asynchronous output. It is valid on both shareable and dedicated persistent socket connections. A SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_WAIT interaction on a shareable persistent socket connection must be in the same application as the original SYNC_SEND or SYNC_SEND_RECEIVE interaction and must use the same shareable persistent connection. This primarily occurs following execution timeout. With this type of interaction, the Java client can receive only one single message. If the IMS OTMA asynchronous hold queue holds no messages for the client ID when the request is made, IMS Connect waits for OTMA to return a message. IMS Connect waits the length of time specified in the executionTimeout property of the SYNC_RECEIVE_ASYNCOUPTUT_SINGLE_WAIT interaction before returning an exception.
7	SYNC_RECEIVE_ CALLOUT	The interaction verb SYNC_RECEIVE_CALLOUT is used to retrieve asynchronous or synchronous callout messages. Use the calloutRequestType property to specify whether only synchronous, only asynchronous, or both synchronous and asynchronous callout messages are to be retrieved. With this type of interaction, if the IMS OTMA asynchronous hold queue contains no messages for the specified client ID when the request is made, IMS Connect waits for OTMA to return a message. The time the IMS TM resource adapter waits for the response from IMS Connect and OTMA depends on the value for the executionTimeout property.

Restriction: The Java EE Connection Architecture (JCA) value SYNC_RECEIVE (2) is not supported.

Lterm name (ltermName)

The LTERM name that is used to override the value in the LTERM field of the IMS application program's I/O PCB.

The value of this property can be set if the client application needs to provide an LTERM override name. This name will be in the IMS application program's I/O PCB, with the intent that the IMS application will make logic decisions based on this override value.

To use the LTERM override name for calling IBM Workload Manager for input transactions, you need to specify the **WLMLTRM** parameter with YES option. This enables OTMA to pass the LTERM override name to IBM Workload Manager, instead of the tpipe name, which is set as default.

Map name (mapName)

The mapName field contains the name of a Message Format Service (MFS) control block. MFS is the component of IMS that performs online formatting of transaction input and output messages.

Because IMS Connect uses IMS OTMA to access IMS, MFS online formatting is bypassed. The mapName field is used by the IMS TM resource adapter to process requests from Java applications whose input and output messages are formatted based on MFS message definitions. The mapName field is not for use directly by Java applications. For reference purposes, the following is a brief description of how the mapName field is used by the IMS TM resource adapter.

For input messages, the value of the mapName property is the name of an MFS message output descriptor, or MOD. The MOD name is provided to the IMS application program in the I/O PCB as the default MODNAME. You can overwrite this default value in your IMS application program. The IMS application program can either use this default MODNAME from the I/O PCB or override that value and use a different MODNAME of its choice.

In general, for output messages, the value of the mapName property is the name of an MFS message output descriptor (MOD). Depending on how the IMS application is programmed, this MOD name can be

set in different ways. The IMS application program can specify a different MODNAME when it inserts the transaction output message to the I/O PCB. If the IMS application program does not specify a different MOD name for use in the output message, the value of the mapName property from the input message will be used as the value of the mapName property in the output message. If a mapName value was not provided in the input message and the IMS application does not set a value in the MODNAME field in the I/O PCB for use in the output message, IMS places the default IMS MOD name in the mapName field of the output message.

Purge async output (purgeAsyncOutput)

This input property determines whether or not IMS Connect purges undelivered output.

This property is valid only for interactions on shareable persistent socket connections that use IMS interaction verb SYNC_SEND_RECEIVE. It is not valid for any interactions on dedicated persistent socket connections. It applies only to commit mode 0 interactions, not to commit mode 1 interactions. However, if a commit mode 1 interaction executes a program-to-program switch, the spawned program will run in commit mode 0 and, therefore, the property will apply.

If the purgeAsyncOutput property is not specified on a SYNC_SEND_RECEIVE interaction on a shareable persistent socket connection, the default is true, and the following output messages are purged:

- Undelivered output messages inserted to the I/O PCB by the primary IMS application program
- Output messages inserted to the I/O PCB by secondary IMS application programs that are invoked by a program-to-program switch

Reroute (reRoute)

This input property determines whether undelivered output is to be rerouted to a named destination that is specified in the reRouteName field.

The reRoute property is valid for interactions on shareable persistent socket connections that use the IMS interaction verb SYNC_SEND_RECEIVE. It is not valid for any interactions on dedicated persistent socket connections. It applies to commit mode 0 interactions, not to commit mode 1 interactions.

However, if a commit mode 1 interaction executes a program-to-program switch, the spawned program will run in commit mode 0 and therefore the property will apply. This property determines whether undelivered output is to be rerouted to a named destination that is specified in the reRouteName field. If reRoute is true, the asynchronous output is not queued to the pipe of the generated client ID. Instead, the asynchronous output is queued to the destination that is specified in the reRouteName field. The default value for reRoute is false.

If both reRoute and purgeAsyncOutput are set to true, an exception is thrown.

Reroute name (reRouteName)

This property provides the name of the destination to which asynchronous output is queued.

If the reRoute property is set to true, this property provides the named destination. If the reRoute property is set to false, the reRouteName property is ignored. If the reRoute property is set to true, and no value for the reRouteName property is provided, the value for the reRouteName property is:

- The value specified in the IMS Connect configuration file
- If no value is specified in the IMS Connect configuration file, the value HWS\$DEF is used

A valid value must meet the following rules:

- Must be a string of 1 to 8 characters. Valid characters are A-Z, 0-9, @, #, and \$.
- Must not begin with the IMS TM resource adapter reserved prefix, HWS.
- Must not be an IMS Connect port number.

Important: If lowercase letters are provided, the letters are changed to uppercase.

The `reRouteName` property is valid only for `SYNC_SEND_RECEIVE` interactions on shareable persistent socket connections. It is not valid for any interactions on dedicated persistent socket connections.

Resume tppe network security credentials (`resumeTppeNSC`)

This property is used by a Java application to indicate whether network security credentials are to be included in asynchronous callout requests from IMS.

If you set the value of the `resumeTppeNSC` property to 1, IMS TM resource adapter sets a flag byte in the OTMA message prefix that is sent to IMS to indicate that network security credentials are to be included in the asynchronous callout message.

Socket timeout (`socketTimeout`)

The maximum amount of time the IMS TM resource adapter waits for a response from IMS Connect before disconnecting the socket and returning an exception to the client application.

The `socketTimeout` value is represented in milliseconds. The value must be greater than zero. If a socket timeout is not specified for an interaction or it is supplied with a socket timeout value of zero milliseconds, the IMS TM resource adapter waits indefinitely for IMS Connect to respond.

Related reference

[Execution timeout \(`executionTimeout`\)](#)

This property specifies the maximum amount of time that is allowed for IMS Connect to send a message to IMS and receive a response

Synchronous callout correlator token (`syncCalloutCorrelatorToken`)

This property contains the correlator token that is generated by IMS Connect and sent with the callout request message to associate the synchronous callout request with the response.

If you are using the managed callout programming model (by using a message-driven bean), this token is automatically managed by the IMS TM resource adapter. In the client-managed callout programming model (non-MDB application), the Java application must manage the correlation token and pass it back with the response by using the associated getter and setter methods. The managed callout programming model is, therefore, the simpler and preferred approach of the two.

This property is valid only when the interaction verb is set to `SYNC_RECEIVE_CALLOUT`.

Synchronous callout status code (`syncCalloutStatusCode`)

This property contains a user-specified status code when the Java application sends an error response message for synchronous callout messages from an IMS application.

The valid user-specified status code is from 500 to 1000.

Sync level (`syncLevel`)

This input property specifies the synchronization level of the interaction between the IMS TM resource adapter and IMS OTMA.

Valid sync level values are 0 (NONE) and 1 (CONFIRM). The `syncLevel` property applies only when the `interactionVerb` property is set to `SYNC_SEND_RECEIVE`, `SYNC_SEND`, and `SYNC_RECEIVE_CALLOUT`. The value of the `syncLevel` property applies to both conversational and non-conversational applications, and is used in combination with the `commitMode` property.

Commit mode 1

Sync levels 0 and 1 are valid. 0 is the default value. If the `interactionVerb` property is set to `SYNC_SEND_RECEIVE` and the `commitMode` property is set to 1, it is not necessary to set the `syncLevel` property. If any other value besides 0 or 1 is passed to the `setSyncLevel(int)` method, an exception is thrown.

Commit mode 0

Sync level 1 is the only valid value. It is not necessary to set the sync level with this commit mode. If any other value is passed to the `setSyncLevel(int)` method, an exception is thrown.

Programs triggered by a program-to-program switch are always processed as commit mode 0, regardless of the sync level of the originating program. For this reason, secondary output from these programs can be put onto the reroute queue or purged according to the setting of the `reRoute` and `purgeNotDeliverable` properties in the original interaction.

Related reference

Commit mode (`commitMode`)

This property indicates the type of commit mode processing to be performed for an IMS transaction.

Transaction expiration (`transExpiration`)

Indicates to IMS OTMA, when the execution timeout value is reached, whether the transaction is considered expired, and OTMA no longer needs to process it.

When this property is set to true, the IMS TM resource adapter client application indicates to OTMA that the transaction can be discarded after the execution times out. This function relieves OTMA from having to process unnecessary messages.

The default is `false`.

Related reference

Execution timeout (`executionTimeout`)

This property specifies the maximum amount of time that is allowed for IMS Connect to send a message to IMS and receive a response

Transaction tracking ID (`trckID`)

This input-only property is an up-to-480 byte token (an alphanumeric string) that uniquely identifies an IMS™ transaction. This tracking ID can be used to show the components which a transaction involves, the order in which the components communicate with each other, and the time it takes a transaction to traverse each component.

The tracking ID is set by the ZAPM™ Connect client and returned in the HTTP response message for each transaction. By default, transaction tracking is disabled. To use a tracking ID, transaction tracking must be enabled.

You can set a tracking ID in the following ways:

- By using a client such as Instana or AppDynamics on Websphere Application Server (WAS) that uses `IMSInteractionSpec` Class methods to directly set the tracking ID. This process is described in [Optional IMS TM resource adapter plug-in support](#). For more information, see the documentation on instrumenting the application environment for [AppDynamics](#) and [Instana](#).
- By using the Z APM Connect Interceptor on z/OS Connect EE to set the tracking ID for the IMS Service Provider. This process is described in [Z APM Connect Interceptor](#). For more information, see the documentation on instrumenting the application environment for [AppDynamics](#) and [Instana](#).
- By specifying the tracking ID manually in the transactions IMS IRM control block as described in [Optional IMS Connect user message exit for AppDynamics](#).

Use conversation ID (`useConvID`)

This property indicates whether an IMS conversation is using the same connection for all iterations of the conversation or different connections.

By default, a subsequent iteration of a conversation on a shareable persistent socket is rejected by IMS Connect (if the socket is different between iterations) because IMS Connect cannot find the relevant information for the socket that is used in the first iteration.

When this property is set to `true`, IMS assigns a unique conversation token that must be passed back and forth between the Java application and IMS Connect. IMS Connect no longer keeps track of the

conversational status of the client. As a result, the **VIEWHWS** IMS Connect command will not display the CONV status for the client unless the command is entered when IMS Connect is processing a conversational iteration. This approach to manage the iterations of a conversation is called the client-managed conversation state programming model, as opposed to IMS Connect-managed conversation state programming model.

Important: The default is `false` for backward compatibility. For applications that are migrating to use IMS TM Resource Adapter Version 12 or later, if you are not using shareable socket connections for different iterations of a conversational transaction, no change is required. For new application development, set this property to `true` to take advantage of using different shareable persistent socket connections for the different iterations of a conversational transaction.

Java API specifications

The IMS TM resource adapter Java API specification includes reference information about the Java classes, interfaces, and methods.

[Version 15.2 Java API documentation](#)

Notices

This information was developed for products and services offered in the US. This material might 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 grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan, Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

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 jurisdictions 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 provide 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 Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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 actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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_.

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.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <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.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

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.

To learn more, see [IBM Privacy Statement](#).

Bibliography

This bibliography lists all of the publications in the IMS 15.3 library.

Title	Acronym
<i>IMS Version 15.3 Application Programming</i>	APG
<i>IMS Version 15.3 Application Programming APIs</i>	APR
<i>IMS Version 15.3 Commands, Volume 1: IMS Commands A-M</i>	CR1
<i>IMS Version 15.3 Commands, Volume 2: IMS Commands N-V</i>	CR2
<i>IMS Version 15.3 Commands, Volume 3: IMS Component and z/OS Commands</i>	CR3
<i>IMS Version 15.3 Communications and Connections</i>	CCG
<i>IMS Version 15.3 Database Administration</i>	DAG
<i>IMS Version 15.3 Database Utilities</i>	DUR
<i>IMS Version 15.3 Diagnosis</i>	DGR
<i>IMS Version 15.3 Exit Routines</i>	ERR
<i>IMS Version 15.3 Installation</i>	INS
<i>IMS Version 15.3 Licensed Program Specifications</i>	LPS
<i>IMS Version 15.3 Messages and Codes, Volume 1: DFS Messages</i>	MC1
<i>IMS Version 15.3 Messages and Codes, Volume 2: Non-DFS Messages</i>	MC2
<i>IMS Version 15.3 Messages and Codes, Volume 3: IMS Abend Codes</i>	MC3
<i>IMS Version 15.3 Messages and Codes, Volume 4: IMS Component Codes</i>	MC4
<i>IMS Version 15.3 Operations and Automation</i>	OAG
<i>IMS Version 15.3 Release Planning</i>	RPG
<i>IMS Version 15.3 System Administration</i>	SAG
<i>IMS Version 15.3 System Definition</i>	SDG
<i>IMS Version 15.3 System Programming APIs</i>	SPR
<i>IMS Version 15.3 System Utilities</i>	SUR

Index

Numerics

- 274X
 - defining to operate with MFS [444](#)
- 3270P Printer
 - defining to operate with MFS [444](#)
- 3290 Display Panel
 - defining to operate with MFS [444](#)
- 3601 workstation
 - defining to operate with MFS [444](#)
- 3770 Data Communication System
 - defining to operate with MFS [444](#)
- 3790 Communication System
 - defining to operate with MFS [444](#)
- 6670 Printer
 - defining to operate with MFS [444](#)

A

- abend codes
 - pseudo- [52](#)
 - S201 [329](#)
 - U0069 [53](#)
 - U0711 [111](#)
 - U0777 [45](#)
 - U1008 [49](#)
 - U119 [111](#)
 - U2478 [45](#)
 - U2479 [45](#)
 - U261 [329](#)
 - U3301 [49](#)
 - U3303 [45](#)
 - U476 [329](#)
 - U711 [95](#)
- access methods
 - DEDB [73](#)
 - description [69](#)
 - GSAM [75](#)
 - HDAM [70](#)
 - HIDAM [72](#)
 - HISAM [74](#)
 - HSAM [74](#)
 - MSDB [73](#)
 - PHDAM [69, 70](#)
 - PHIDAM [69, 72](#)
 - SHISAM [75](#)
 - SHSAM [75](#)
- access of
 - IMS databases through z/OS [75](#)
 - segments through different paths [81](#)
- accessibility
 - features [xix](#)
 - keyboard shortcuts [xix](#)
- accessing data
 - from an application program [541](#)
- accessing databases with application programs [36](#)
- accessing GSAM databases [283](#)

- activation specification property
 - dataStoreName [770](#)
 - groupName [770](#)
 - hostName [770](#)
 - password [770](#)
 - portNumber [770](#)
 - queueNames [770](#)
 - retryLimit [770](#)
 - SSLEnabled [770](#)
 - SSLEncryptionType [770](#)
 - SSLKeyStoreName [770](#)
 - SSLKeyStorePassword [770](#)
 - SSLTrustStoreName [770](#)
 - SSLTrustStorePassword [770](#)
 - userName [770](#)
- adding
 - data [538](#)
- addressability to UIB, establishing [219](#)
- addressing mode (AMODE) [233, 375](#)
- Advanced Function Printing (AFP) [427](#)
- AFPDS and IMS Spool API [427](#)
- aged timeout property [789](#)
- aggregates, data [21](#)
- AIB (application interface block)
 - address return [227](#)
 - AIB identifier (AIBID) [367](#)
 - AIBERRXT (reason code)
 - AIB mask [367](#)
 - AIBID (AIB identifier) field, AIB mask [214](#)
 - AIBLEN (DFSAIB allocated length) field [367](#)
 - AIBLEN (DFSAIB allocated length) field, AIB mask [214](#)
 - AIBOALEN (maximum output area length)
 - AIB mask [367](#)
 - AIBOAUSE (used output area length)
 - AIB mask [367](#)
 - description [214](#)
 - AIBREASN (reason code) [214](#)
 - AIBRSA1 (resource address)
 - AIB mask [367](#)
 - AIBRSNM1 (resource name)
 - AIB mask [367](#)
 - description [214](#)
 - AIBSFUNC (subfunction code)
 - AIB mask [367](#)
 - description [214](#)
 - defining storage [226, 370](#)
 - description [226, 370](#)
 - DFSAIB allocated length (AIBLEN) [214, 367](#)
 - fields [214, 367](#)
 - mask [214, 367](#)
 - program entry statement [227](#)
 - specifying [214, 367](#)
- AIB (Application Interface Block)
 - supported commands [477](#)
- AIB identifier (AIBID)
 - description [214](#)
- AIB interface [11](#)

- AIB mask, specifying [214](#)
- AIBERRXT (reason code) [214](#)
- AIBOALEN (maximum output area length) field, AIB mask [214](#)
- AIBOAUSE (used output area length) field, AIB mask [214](#)
- AIBREASN (reason code) AIB mask, field [214](#)
- AIBREASN (reason code) field, AIB mask [214](#)
- AIBRSA1 (resource address) field, AIB mask [214](#)
- AIBRSNM1 (resource name) field, AIB mask [214](#)
- AIBSFUNC (subfunction code) field, AIB mask [214](#)
- AIBTDLI interface [53](#), [226](#), [370](#)
- AJ status code [496](#)
- AL_LEN call [321](#), [325](#)
- allocation, dynamic [54](#)
- alternate client IDs [872](#)
- alternate destinations
 - sending messages [391](#)
- alternate PCB
 - alternate destinations [391](#)
 - change call [391](#)
 - CHNG call
 - description [391](#)
 - using PURG with [391](#)
 - defining in ISRT call [391](#)
 - destination of modifiable alternate PCBs [391](#)
 - express [391](#)
 - modifiable
 - description [391](#)
 - use [391](#)
 - modifiable alternate PCBs
 - changing the destination [391](#)
 - CHNG call [391](#)
 - description [391](#)
 - modifiable PCBs [391](#)
 - PURG call
 - description [391](#)
 - using CHNG with [391](#)
 - replying to one alternate terminal [391](#)
 - response [403](#)
 - SAMETRM=YES [403](#)
 - sending messages
 - to several alternate destinations [391](#)
 - using alternate PCBs [391](#)
 - using the PURG call [391](#)
 - sending messages to other terminals [391](#)
 - types and uses [367](#)
 - use with program-to-program message switching [393](#)
 - using the CHNG call with [391](#)
- alternate PCB mask
 - description [367](#)
 - format [367](#)
- alternate PCBs [98](#)
- alternate response PCBs [98](#)
- alternate terminals
 - responding [391](#)
- AM status code [496](#)
- AMODE [233](#)
- AMODE(31) [476](#)
- analysis of
 - processing requirements [35](#)
 - required application data [17](#)
 - user requirements [15](#)
- anchor point, root [70](#)
- AND operators

- AND operators (*continued*)
 - dependent [272](#)
 - independent [272](#)
- API (application programming interface) for LU 6.2 devices
 - explicit API [110](#)
 - implicit API [110](#)
- APPC
 - application program types for LU 6.2 devices [101](#)
 - basic conversation [104](#)
 - description [101](#)
 - entering IMS transactions from LU 6.2 devices [101](#)
 - LU 6.2 partner program design
 - DFSAPPC message switch [133](#)
 - flow diagrams [111](#)
 - integrity after conversation completion [131](#)
 - mapped conversation [104](#)
 - RRS [329](#)
- APPC conversational program
 - CPI-C driven [409](#)
 - ending the conversation [408](#)
 - message switching [406](#)
 - modified application program
 - MSC [409](#)
 - remote execution, MSC [409](#)
 - modified IMS application [409](#)
- application data
 - analyzing required [17](#)
 - identifying [17](#)
- application design
 - analyzing
 - processing requirements [55](#)
 - the data a program must access [56](#)
 - user requirements [15](#), [17](#)
 - data dictionary, using [20](#)
 - DataAtlas [20](#)
 - DB/DC Data Dictionary [20](#)
 - debugging [427](#)
 - designing a local view [21](#)
 - documenting [15](#), [163](#)
 - IMS Spool API interface [425](#)
 - overview [15](#)
- application interface block (AIB) [214](#)
- Application Interface Block (AIB)
 - supported commands [477](#)
- application program
 - checking success of SQL statements [509](#)
 - coding SQL statements
 - dynamic SQL [511](#)
 - selecting rows using a cursor [545](#)
 - design
 - tasks overview [12](#)
 - documentation [163](#)
 - HALDB environment, scheduling [233](#)
 - hierarchy examples [5](#)
 - I/O areas, specifying [221](#), [370](#)
 - including queries [509](#)
 - preparation
 - preparing for running [548](#)
 - sync point [257](#)
 - test [135](#)
 - TSO [45](#)
- application programming
 - catalog [234](#)
 - database versioning [237](#)

- application programming (*continued*)
 - IMS catalog [234](#)
 - SQL [505](#)
- Application programming
 - IMS database [165](#)
- application programming interface (API) [110](#)
- application programs
 - accessing databases [36](#)
 - assembler language
 - DL/I call formats [195](#), [349](#)
 - CICS
 - testing [155](#)
 - DL/I calls
 - sample call formats [195](#), [349](#)
 - host structures [511](#)
 - host variables [510](#)
 - Pascal [204](#), [357](#)
 - PL/I [206](#), [360](#)
 - PL/I language [713](#), [715](#), [717](#), [719](#)
 - PL/I top-down development [711](#)
 - sample PL/I templates [715](#)
- application programs, IFP [483](#)
- applications
 - development tasks overview [12](#)
- APSB (allocate program specification block) [110](#), [329](#)
- area, I/O [11](#)
- AREALIST call [326](#)
- ArrayResultSet property [611](#)
- assembler language
 - DL/I call-level sample [178](#)
 - DL/I command-level sample [462](#)
 - DL/I program structure [177](#)
 - entry statement [227](#)
 - MPP coding [418](#)
 - program entry [227](#)
 - register 1, program entry [227](#)
 - return statement [227](#)
 - skeleton MPP [418](#)
 - SSA definition examples [223](#)
 - UIB, specifying [219](#)
 - UIBDLTR
 - introduction [219](#)
 - UIBFCTR
 - introduction [219](#)
 - UIBPCBAL
 - introduction [219](#)
- asynchronous callout
 - definition [760](#)
 - message flow [763](#)
 - message-driven beans (MDBs) [763](#)
 - non-MDB applications [780](#)
 - non-MDB sample Java application [781](#)
 - requirements [727](#)
- asynchronous conversation, description for LU 6.2 transactions [102](#)
- asynchronous output
 - availability [873](#)
 - purging [878](#)
- AUTH call [91](#)
- authorization
 - ID Db2 for z/OS [91](#)
 - security [91](#)
- availability enhancements, data [503](#)

- availability of data
 - sensitivity [66](#)

B

- back-out database changes [64](#)
- backing out
 - database changes [486](#)
- backout point
 - description [410](#)
 - intermediate [486](#)
 - intermediate (SETS/SETU) [414](#)
- backout point, intermediate [267](#)
- backout, dynamic [39](#)
- bank account database example [5](#)
- basic checkpoint [48](#), [62](#), [65](#)
- basic conversation, APPC [104](#)
- Basic edit
 - IMS TM [437](#)
- Basic Edit
 - input message [382](#)
 - output message [382](#)
 - translation to uppercase [382](#)
- basic edit, overview of [94](#)
- Batch Backout utility [39](#)
- batch environment [36](#)
- batch message processing (BMP) programs
 - issuing checkpoints [485](#)
 - PCBs [483](#)
- batch message processing program. [42](#)
- batch programs
 - assembler language [177](#)
 - C language [180](#)
 - COBOL [182](#)
 - command-level samples
 - assembler [462](#)
 - C [471](#)
 - COBOL [465](#)
 - PL/I [468](#)
 - converting to BMPs [42](#)
 - databases that can be accessed [36](#), [56](#)
 - DB batch processing [39](#)
 - differences from online [39](#)
 - IMS database [180](#)
 - issuing checkpoints [49](#), [62](#), [485](#)
 - maintaining integrity [264](#)
 - Pascal [188](#)
 - PL/I [190](#)
 - recovery [39](#), [62](#)
 - recovery of database [64](#)
- Batch Terminal Simulator (BTS) [136](#)
- batch-oriented BMPs. [42](#)
- BILLING segment [5](#)
- binder options with EXEC DLI [476](#)
- binding, reference [192](#)
- BKO execution parameter [39](#)
- block descriptor word (BDW)
 - IMS Spool API [426](#)
- BMP (batch message processing) program
 - batch-oriented
 - checkpoints in [49](#), [62](#)
 - databases that can be accessed [42](#)
 - description of [42](#)

- BMP (batch message processing) program (*continued*)
 - batch-oriented (*continued*)
 - limiting number of locks with LOCKMAX= parameter [49](#)
 - recovery [42](#)
 - databases that can be accessed [36](#), [56](#)
 - transaction-oriented
 - checkpoints in [49](#)
 - databases that can be accessed [44](#)
 - recovery [44](#)
- BMP (batch message processing) programs
 - basic checkpoint
 - issuing [485](#)
 - checkpoint (CHKP)
 - description [485](#)
 - EXEC DLI command [485](#)
 - checkpoint (CHKP) command
 - in a Batch or BMP Program [485](#)
 - issuing [485](#)
 - checkpoint (CHKP) EXEC DLI command
 - current position [485](#)
 - CHKP (Checkpoint)
 - description [485](#)
 - EXEC DLI command [485](#)
 - CHKP (Checkpoint) command
 - issuing in a Batch or BMP Program [485](#)
 - CHKP (Checkpoint) EXEC DLI command
 - current position [485](#)
 - committing your program's changes to a database [485](#)
 - EXEC DLI recovery commands
 - CHKP (Checkpoint) [485](#)
 - SYMCHKP (Symbolic Checkpoint) [485](#)
 - I/O area
 - symbolic CHKP [485](#)
 - issuing checkpoints [485](#)
 - PCBs [483](#)
 - planning for database recovery
 - CHKP command checkpoints [485](#)
 - taking checkpoints [485](#)
 - recovery EXEC DLI commands
 - basic CHKP [485](#)
 - SYMCHKP [485](#)
 - SYMCHKP (Symbolic Checkpoint) command
 - description [485](#)
- BMPs, transaction-oriented
 - ROLB [264](#)
- Boolean operators
 - dependent AND [272](#)
 - independent AND [272](#)
 - SSA, coding [222](#)
- BTS (Batch Terminal Simulator) [136](#)
- buffer pool, STAT call and OSAM [139](#)
- buffer subpool, statistics for debugging
 - enhanced STAT call and
 - OSAM [143](#)
 - VSAM [141](#), [148](#)
- business applications
 - composite [822](#)
 - developing [751](#)
- business process choreography
 - overview [820](#)
 - planning details [822](#)
- ByteBuffers [611](#)

C

- C language
 - __pcblst [227](#)
 - application programming
 - IMS database [198](#)
 - IMS Transaction Manager [352](#)
 - batch program, coding [180](#)
 - DL/I call formats [198](#), [352](#)
 - DL/I calls
 - sample call formats [198](#), [352](#)
 - DL/I program structure [180](#)
 - entry statement [227](#)
 - exit [227](#)
 - I/O area [198](#), [352](#)
 - PCBs, passing [227](#)
 - return statement [227](#)
 - skeleton MPP [418](#)
 - SSA definition examples [223](#)
 - system function [227](#)
- C program
 - DL/I command-level sample [471](#)
- C/MVS [53](#)
- CALL statement (DL/I test program) [135](#)
- call-level programs
 - comparing with command-level programs
 - command codes and options [501](#)
 - commands and calls [499](#)
 - DL/I calls available to IMS and CICS command-level [499](#)
- call-level programs, scheduling a PSB [59](#)
- callout IVP (installation verification program)
 - as EAR file [746](#)
- callout request
 - asynchronous callout
 - programming model [457](#)
 - asynchronous callout request [449](#)
 - comparison of synchronous and asynchronous callout requests [449](#)
 - IMS TM Resource Adapter
 - Enterprise JavaBeans (EJB) [450](#)
 - Java EE application [450](#)
 - message-driven bean (MDB) [450](#)
 - web service [450](#)
 - JMS (Java Message Service) implementation
 - IMSQueueConnectionFactory [692](#)
 - overview [449](#)
 - resume tpipe
 - IRM_TIMER field [452](#)
 - protocol [452](#)
 - security [452](#)
 - security exit routine (DFSYRTUX) [452](#)
- RESUME TPIPE call [450](#)
- SOAP Gateway
 - web service [450](#)
- synchronous callout
 - COBOL code example [453](#)
 - control data [456](#)
 - JBP (Java batch processing) regions [692](#)
 - JMP (Java message processing) regions [692](#)
 - programming model [453](#)
- synchronous callout request [449](#)
- synchronous program switch
 - JBP (Java batch processing) regions [697](#)

- callout request (*continued*)
 - synchronous program switch (*continued*)
 - JMP (Java message processing) regions [697](#)
 - synchronous program switch request [449](#)
 - user-written IMS Connect TCP/IP application [450](#)
- callout requests
 - diagnosis [837](#)
 - IMS callout applications [747](#)
 - IVP samples [742](#)
 - message-driven bean (MDB) sample [774](#)
 - retrieving from message-driven beans (MDB) [766](#)
 - types [873](#)
- calls, DL/I [11](#)
- calls, system service
 - calls, system service
 - SETS/SETU (set a backout point) [267](#)
 - PLITDLI [267](#)
 - SETS (Set a Backout Point) call
 - description [267](#)
 - SETS/SETU (set a backout point)
 - backing out to an intermediate backout point [267](#)
 - SETU (Set a Backout Point Unconditional) call
 - description [267](#)
 - SETU, call function [267](#)
- CCTL (coordinator controller)
 - restrictions
 - with BTS (Batch Terminal Simulator) [136](#)
 - with DL/I test program [135](#)
- CEETDLI
 - address return [227](#)
 - program entry statement [227](#)
- character host variable
 - COBOL [525](#)
- checkpoint
 - basic [48](#), [62](#)
 - calls, when to use [49](#)
 - frequency, specifying [51](#), [62](#)
 - IDs [62](#)
 - in batch programs [49](#), [62](#)
 - in batch-oriented BMPs [49](#), [62](#)
 - in MPPs [49](#)
 - in transaction-oriented BMPs [49](#)
 - issuing [39](#)
 - printing log records [62](#)
 - restart [49](#), [65](#)
 - summary of [48](#)
 - symbolic [48](#), [62](#)
- checkpoint (CHKP) calls
 - considerations [175](#)
 - description [263](#)
 - issuing [263](#)
- checkpoints
 - relationship to commit point and sync point [257](#)
- CHKP (checkpoint) [48](#)
- CHKP (checkpoint) call
 - considerations [175](#)
- CHKP (symbolic checkpoint) call
 - with GSAM [289](#)
- CHKPT=EOV [48](#)
- CHNG call
 - usage [429](#)
 - with directed routing [397](#)
- CHNG system service call [425](#)
- CICS
 - command language translator [475](#)
 - distributed transactions
 - accessing IMs [60](#)
 - CICS applications
 - unit testing [155](#)
 - CICS applicationsIMS batch regions
 - database logging [67](#)
 - ESTAE routine [67](#)
 - resource cleanup [67](#)
 - STAE routine [67](#)
 - CICS DL/I call program
 - compiling [176](#)
 - CICS online programs
 - assembler language
 - sample [178](#)
 - COBOL, sample [184](#)
 - PL/I, sample [192](#)
 - CIMS [329](#)
 - classes schedule, example [28](#)
 - client IDs [869](#)
 - CLOSE
 - statement
 - description [548](#)
 - WHENEVER NOT FOUND clause [512](#), [514](#)
 - closing a GSAM database explicitly [287](#)
 - CMOResponse
 - interaction property [874](#)
 - CMPAT option [483](#)
 - CMPAT=YES PSB specification [39](#)
 - COBOL
 - application programming [201](#), [355](#)
 - copybook types [640](#)
 - data types
 - mapped to COBOL [640](#)
 - DL/I call formats [201](#), [355](#)
 - DL/I call-level, sample [184](#)
 - DL/I command-level sample [465](#)
 - DL/I program structure [182](#)
 - entry statement [227](#)
 - mapping to IMS [640](#)
 - return statement [227](#)
 - skeleton MPP [419](#)
 - SQL aggregate function
 - ASC [536](#)
 - AVG [536](#)
 - COUNT [536](#)
 - DESC [536](#)
 - GROUP BY [536](#)
 - MAX [536](#)
 - MIN [536](#)
 - ORDER BY [536](#)
 - SUM [536](#)
 - SQL aggregate functions supported [536](#)
 - SSA definition examples [223](#)
 - types
 - data, mapped to COBOL [640](#)
 - UIB, specifying [219](#)
- COBOL application program
 - defining the SQLIMSDA [510](#), [524](#)
 - host structure [529](#)
 - host variable
 - use of hyphens [533](#)
 - host variable, declaring [525](#)

- COBOL application program (*continued*)
 - INCLUDE statement [533](#)
 - including SQLIMSCA [524](#)
 - naming convention [533](#)
 - options [533](#)
 - SQLIMSCODE host variable [524](#)
 - SQLIMSTATE host variable [524](#)
 - WHENEVER statement [533](#)
- code
 - course [17](#)
 - transaction [40](#)
- codes abend [45](#)
- codes, status
 - logical relationships [277](#)
- coding batch
 - PL/I programs [190](#)
- coding DC calls and data areas
 - in assembler language [418](#)
 - in C Language [418](#)
 - in COBOL [419](#)
 - in Pascal [420](#)
 - in PL/I [422](#)
 - skeleton MPP [417–420](#), [422](#)
- coding rules, SSA [222](#)
- coding SQL statements
 - dynamic [511](#)
- colon
 - preceding a host variable [513](#)
- column
 - labels, usage [514](#)
 - name, with UPDATE statement [539](#)
 - retrieving, with SELECT [541](#)
- columns
 - fields, compared to [555](#)
 - relational representation, in [555](#)
- command codes
 - D
 - examples [174](#)
 - DEDBs [174](#)
 - F
 - restrictions [320](#)
 - overview [174](#)
 - Q [270](#)
 - qualified SSAs [174](#)
 - restrictions [222](#)
 - subset pointers [174](#)
 - unqualified SSAs [174](#)
- command language translator, CICS [475](#)
- command-level program
 - DFHEIENT [462](#)
 - DFHEIRET [462](#)
 - DFHEISTG [462](#)
 - parameters
 - EIBREG [462](#)
 - RCREG [462](#)
 - RCREG parameters [462](#)
 - reentry [462](#)
 - SCHD PSB command [462](#)
- command-level programs
 - adjustable character string [481](#)
 - array, connected [481](#)
 - assembler language
 - I/O area [481](#)
 - automatic storage [468](#)

- command-level programs (*continued*)
 - C code standard header file [471](#)
 - character string
 - adjustable [481](#)
 - fixed-length [481](#)
 - COBOL
 - I/O area [481](#)
 - commands
 - SCHD PSB [468](#)
 - comparing with call-level programs
 - command codes and options [501](#)
 - commands and calls [499](#)
 - concatenated key, segment [480](#)
 - connected array [481](#)
 - DIB (DL/I interface block) [477](#)
 - DL/I calls available to IMS and CICS [499](#)
 - EIBREG parameter [462](#)
 - fixed-length character string [481](#)
 - GE status code [468](#), [471](#)
 - I/O area
 - assembler language [481](#)
 - COBOL [481](#)
 - coding [481](#)
 - PL/I [481](#)
 - restriction [481](#)
 - I/O area, defining [481](#)
 - key feedback area, defining [480](#)
 - major structure [481](#)
 - minor structure [481](#)
 - PL/I
 - I/O area [481](#)
 - preparing EXEC DL/I program for execution [475](#)
 - reentrance [462](#)
 - restrictions
 - I/O area [481](#)
 - I/O area, PL/I [481](#)
 - samples
 - assembler language [462](#)
 - C [471](#)
 - COBOL [465](#)
 - PL/I [468](#)
 - SCHD PSB command [468](#)
 - segment
 - concatenated key [480](#)
 - standard header file, C code [471](#)
 - status codes
 - GE [468](#), [471](#)
 - structure
 - major [481](#)
 - minor [481](#)
- commands, EXEC DLI [11](#)
- COMMENTS statement [135](#)
- commit
 - single-phase [260](#)
 - UOR [260](#)
- commit mode processing
 - client IDs [783](#)
 - interaction verbs [783](#)
 - transaction pipes [783](#)
- commit modes
 - interaction property [874](#)
 - processing [782](#)
 - supported interactions [782](#)

- commit point
 - process [257](#)
 - relationship to check point and sync point [257](#)
- commit point processing
 - DEDB [319](#)
 - MSDB [302](#)
- commit points [39](#), [45](#), [62](#)
- Common Client Interface (CCI)
 - IMS TM resource adapter [828](#)
 - IMS TM resource adapter samples [830](#)
 - passing input and output [828](#)
- Common Connector Framework
 - file names [734](#)
 - migration [732](#)
- communicating with other IMS TM systems [395](#)
- communication protocols
 - TCP/IP connections [792](#)
- COMPARE statement [135](#)
- comparing EXEC DLI
 - commands with DL/I calls [499](#)
 - options with command codes [501](#)
- comparison of symbolic CHKP and basic CHKP [48](#)
- comparison to ROLB and ROLS call [410](#)
- comparison to ROLL and ROLB call [410](#)
- comparison to ROLL and ROLS call [410](#)
- compiler, COBOL [465](#)
- compiling, options with EXEC DLI [476](#)
- component-managed transactions [826](#)
- composite business applications
 - business process choreography [822](#)
 - conversational programs [820](#)
 - requirements [727](#)
- concatenated data sets, GSAM [293](#)
- concatenated segments, logical relationships [274](#)
- concurrent access to full-function databases [39](#)
- connection factories
 - alternate resources [749](#)
 - configuration [791](#)
 - creating [736](#)
 - custom properties [749](#)
- connection factory
 - WebSphere Liberty servers [738](#)
- connection pooling [789](#)
- connections
 - managed [828](#)
 - non-managed [828](#)
 - persistent
 - configuring [791](#)
 - dedicated [784](#)
 - overview [807](#)
 - reconnection [789](#)
 - releasing sockets [789](#)
 - samples [830](#)
 - shareable persistent sockets [787](#)
- considerations in screen design [94](#)
- CONTINUE clause of WHENEVER statement [523](#)
- CONTINUE-WITH-TERMINATION indicator [53](#)
- continuing a conversation [95](#)
- control data
 - synchronous callout [456](#)
- control, passing processing [11](#)
- conventions, naming [15](#)
- conversation attributes
 - asynchronous [102](#)
 - conversation attributes (*continued*)
 - MSC synchronous and asynchronous [102](#)
 - synchronous [102](#)
- conversation ID
 - conversational [821](#)
- conversation IDs [822](#), [873](#), [880](#)
- conversation state, rules for APPC verbs [104](#)
- conversational mode
 - description [98](#)
 - LU 6.2 transactions [102](#)
- conversational processing
 - abnormal termination of, precautions for [96](#)
 - coding necessary information [408](#)
 - deferred program switch [95](#)
 - designing a conversation [95](#)
 - DFSCONE0 [96](#)
 - example [398](#)
 - for APPC/IMS [407](#)
 - gathering requirements [95](#)
 - how to continue the conversation [95](#)
 - how to end the conversation [95](#)
 - immediate program switch [95](#)
 - overview [95](#), [398](#)
 - passing control and continuing the conversation [404](#)
 - passing the conversation to another program [95](#)
 - recovery considerations [96](#)
 - replying to the terminal [403](#)
 - SPA [96](#)
 - structure [399](#)
 - use with alternate response PCBs [98](#)
 - using a deferred program switch to end the conversation [95](#)
 - using ROLB, ROLL and ROLS in [403](#)
 - what happens in a conversation [95](#)
- conversational program
 - definition [398](#)
- conversational programs
 - building business process choreography applications for [822](#)
 - client-managed conversation state programming model [823](#)
 - enabling Java clients [822](#)
 - IMS Connect-managed conversation state programming model [823](#)
 - managing conversation states [821](#)
 - orphaned [821](#)
 - overview [820](#), [822](#)
 - programming models [821](#)
- conversations
 - ended [873](#)
 - orphaned [821](#)
 - transaction processing [822](#)
- conversations, preventing abnormal termination [96](#)
- converting an existing application [15](#)
- coordinator controller. [135](#)
- coordinator, sync-point [105](#)
- copybook types [640](#)
- correlation tokens
 - interaction property [879](#)
 - managing [777](#)
 - message flow [762](#)
- course code [17](#)
- CPI Communications driven program
 - sync point [257](#)

- CPI Resource Recovery calls [257](#)
- creation of
 - a new hierarchy [81](#)
 - reports [20](#)
- crossing a unit of work (UOW) boundary when processing DEDBs [498](#)
- currency of data [3](#)
- current position
 - determining [241](#)
- current roster [17](#)
- cursor
 - closing [548](#)
 - description [545](#), [546](#)
 - non-scrollable [546](#)
 - OPEN statement [546](#)
 - result segment [545](#)
 - row-positioned
 - declaring [546](#)
 - description [545](#)
 - end-of-data condition [547](#)
 - steps in using [546](#)

D

- data
 - a program's view [9](#)
 - accessing from an application program [541](#)
 - adding [538](#)
 - aggregate [21](#)
 - associated with WHERE clause [541](#)
 - documentation [20](#)
 - elements, homonym [19](#)
 - elements, isolating repeating [22](#)
 - elements, naming [19](#)
 - hierarchical relationships [5](#)
 - integrity, how DL/I protects [61](#)
 - keys [22](#)
 - modifying [538](#)
 - recording its availability [20](#)
 - relationships, analyzing [22](#)
 - retrieval using SELECT *
 - variable-length segments [544](#)
 - retrieving a set of rows [547](#)
 - structuring [22](#)
 - unique identifier [19](#)
- data areas
 - coding [175](#)
- data availability
 - considerations [52](#), [65](#)
 - levels [9](#)
 - recording [20](#)
- data availability enhancements [503](#)
- data bindings
 - creating in MDBs [766](#)
 - input message formats [793](#)
 - output message formats [793](#)
- data capture [226](#), [370](#)
- data currency [3](#)
- data definition [15](#)
- data dictionary
 - DataAtlas [20](#), [163](#)
 - DB/DC Data Dictionary [20](#), [163](#)
 - documentation for other programmers [163](#)
 - in application design [20](#)

- data element
 - description [17](#)
 - homonym [19](#)
 - isolating repeating [22](#)
 - listing [17](#)
 - naming [19](#)
 - synonym [19](#)
- data elements, grouping into hierarchies [22](#)
- data entity [17](#)
- data entry database [36](#)
- data mask [11](#)
- data propagation
 - sync point [261](#)
- data redundancy [3](#)
- data redundancy, reducing [274](#)
- data sensitivity [9](#)
- data sensitivity, defined [84](#)
- data storage methods
 - combined file [3](#)
 - in a database [3](#)
 - separate files [3](#)
- data store names [870](#)
- data structure [9](#)
- data structure conflicts, resolving [76](#)
- data structures [175](#)
- data type
 - comparisons [513](#)
 - compatibility
 - COBOL and SQL [531](#)
- DataAtlas [20](#), [163](#)
- database
 - access to [56](#)
 - administrator [5](#)
 - availability
 - obtaining information [503](#)
 - status codes, accepting [503](#)
 - calls
 - Fast Path [327](#)
 - changes, backing out [64](#)
 - DBCTL facilities
 - REFRESH command [503](#)
 - description (DBD) [9](#)
 - example, medical hierarchy [5](#)
 - hierarchy [5](#)
 - integrity, maintaining [485](#)
 - options [69](#)
 - planning for recovery
 - backing out database changes [486](#)
 - position
 - determining [241](#)
 - record, processing [11](#)
 - recovery [485](#)
 - recovery with ROLL call [264](#)
 - recovery, back out changes [264](#)
 - REFRESH command [503](#)
 - unavailability [52](#), [65](#)
- database and data communications security [15](#)
- DATABASE macro [89](#)
- database record [5](#)
- database recovery
 - backing out [412](#)
 - backing out database changes [64](#)
 - checkpoints, description [62](#)
 - planning

- database recovery (*continued*)
 - planning (*continued*)
 - XRST command [486](#)
 - restarting your program, description [65](#)
- Database Resource Adapter (DRA) [329](#)
- database statistics, retrieving [139](#)
- database types
 - areas [36](#)
 - Db2 for z/OS [36](#), [56](#)
 - DEDB [36](#), [73](#)
 - description [36](#)
 - full-function [36](#)
 - GSAM [36](#), [75](#)
 - HDAM [70](#), [72](#)
 - HISAM [74](#)
 - HSAM [74](#)
 - MSDB [36](#), [73](#)
 - PHDAM [69](#), [70](#)
 - PHIDAM [69](#), [72](#)
 - relational [36](#)
 - root-segment-only [36](#)
 - SHISAM [75](#)
 - SHSAM [75](#)
- databases
 - accessing with application programs [36](#)
 - versioning
 - application programming [237](#)
- DB batch processing [39](#)
- DB Control
 - DRA (Database Resource Adapter) [329](#)
- DB Control DRA (Database Resource Adapter) [329](#)
- DB PCB
 - definition [483](#)
- DB PCB (database program communication block)
 - concatenated key and PCB mask [283](#)
 - database
 - DB PCB, name [283](#)
 - database name [283](#)
 - entry statement, pointer [283](#)
 - fields [212](#)
 - fields in a DB PCB [283](#)
 - key feedback area
 - length field in DB PCB [283](#)
 - key feedback area length field [283](#)
 - length of key feedback area [283](#)
 - mask
 - fields [283](#)
 - fields, GSAM [283](#)
 - general description [212](#)
 - name [283](#)
 - relation [212](#)
 - specifying [212](#)
 - masks
 - DB PCB [212](#)
 - multiple DB PCBs [255](#)
 - processing options
 - field in DB PCB [283](#)
 - processing options field [283](#)
 - relation to DB PCB [212](#)
 - RSA (record search argument)
 - overview [283](#)
 - secondary indexing, contents [274](#)
 - status code field [283](#)
 - status codes
- DB PCB (database program communication block) (*continued*)
 - status codes (*continued*)
 - field in DB PCB [283](#)
 - undefined-length records [283](#)
 - variable-length records [283](#)
- DB/DC
 - Data Dictionary [20](#), [163](#)
 - environment [36](#)
- DB2 (DATABASE 2)
 - with IMS TM [390](#), [423](#)
- Db2 for z/OS
 - databases [36](#), [56](#)
- Db2 for z/OS access
 - application programming [703](#)
 - committing work [703](#)
 - drivers [703](#)
 - IMS databases, compared to [703](#)
 - rolling back work [703](#)
- Db2 for z/OS stored procedures
 - ASUTIME, recommendation [334](#)
 - best practices, design [334](#)
 - best practices, ODBA [334](#)
 - best practices, writing [336](#)
 - design best practices [334](#)
 - ODBA best practices [334](#)
 - stopping threads [336](#)
 - threads, stopping [336](#)
 - writing [336](#)
 - writing best practices [336](#)
- DBA 5
- DBASF, formatted OSAM buffer pool statistics [139](#)
- DBASS, formatted summary of OSAM buffer pool statistics [139](#)
- DBASU, unformatted OSAM buffer pool statistics [139](#)
- DBCTL
 - environment [36](#)
- DBCTL (Database Control)
 - single-phase commit [260](#)
 - two-phase commit [258](#)
- DBCTL environment [56](#)
- DBCTL facilities
 - ACCEPT command [503](#)
 - data availability [503](#)
 - QUERY command [503](#)
 - ROLS (Roll Back to SETS or SETU) command [486](#)
 - SETS (Set a Backout Point) command [486](#)
- DBCTLID parameter [329](#)
- DBD (database description) [9](#)
- DBESF, formatted OSAM subpool statistics [143](#)
- DBESO, formatted OSAM pool online statistics [143](#)
- DBESS, formatted summary of OSAM pool statistics [143](#)
- DBESU, unformatted OSAM subpool statistics [143](#)
- DCCTL
 - environment [36](#)
- DDL [618](#)
- DDM (distributed data management) [560](#)
- deadlock, program [39](#)
- debug a program, How to [159](#)
- DECLARE CURSOR statement
 - description, row-positioned [546](#)
 - prepared statement [512](#), [514](#)
- DEDB
 - DL/I calls
 - AL_LEN call [325](#)

- DEDB (*continued*)
 - DL/I calls (*continued*)
 - AREALIST call [326](#)
 - DEDBINFO call [326](#)
 - DEDSTR call [327](#)
 - DI_LEN call [325](#)
 - DS_LEN call [325](#)
 - summary [321](#)
- DEDB (data entry database)
 - call restrictions [320](#)
 - crossing a unit of work (UOW) boundary when processing [498](#)
 - data entry database [487](#)
 - database
 - processing, Fast Path [487](#)
 - dependent segments
 - sequential [487](#)
 - direct dependent segments, in DEDBs [487](#)
 - DL/I calls [320](#)
 - Fast Path
 - database, processing [487](#)
 - multiple qualification statements [173](#)
 - processing
 - commit point [319](#)
 - DEDBs [487](#)
 - fast path [295](#)
 - H option [319](#)
 - overview [487](#)
 - P option [319](#)
 - POS (Position) command [496](#)
 - POS call [316](#)
 - secondary index [308](#)
 - subset pointers [304](#), [487](#)
 - segment
 - sequential dependent [487](#)
 - sequential dependent segments
 - in DEDBs [487](#)
 - updating segments [298](#)
 - updating with secondary index [308](#)
 - updating with subset pointers [304](#)
- DEDB (data entry database) and the PROCOPT operand [87](#)
- DEDBINFO call [326](#)
- dedicated persistent socket connections
 - establishing [792](#)
 - overview [792](#)
 - processing models [752](#), [754](#)
 - programming models [756](#), [757](#)
 - TCP/IP connections [792](#)
- DEDSTR call [327](#)
- deferred program switch [95](#)
- defining application program elements to IMS
 - AIB [477](#)
 - AIB (Application Interface Block)
 - AIB mask [477](#)
 - restrictions [477](#)
 - Application Interface Block (AIB)
 - AIB mask [477](#)
 - restrictions [477](#)
 - DIB [477](#)
 - execution diagnostic facility [477](#)
 - I/O area [481](#)
 - key feedback area [480](#)
 - restrictions
 - AIB [477](#)
- defining application program elements to IMS (*continued*)
 - Transaction Server, CICS [477](#)
- defining subset pointers [489](#)
- definition
 - data [15](#)
 - dependent segment [5](#)
 - root segment [5](#)
- Delete (DLET) call
 - with MSDB, DEDB or VSO DEDB [298](#)
- DELETE statement
 - description [540](#)
- deleting
 - current rows [547](#)
 - data [540](#)
 - every row from a table [540](#)
 - rows from a table [540](#)
- dependent AND operator [272](#)
- dependent segment [5](#)
- dependent segments
 - retrieving [80](#)
 - sequential
 - identifying free space [498](#)
 - locating a specific dependent [497](#)
 - locating the last inserted dependent [497](#)
- dependents
 - direct [36](#)
 - sequential [36](#)
- dependents, direct [298](#)
- deploying
 - callout IVP EAR file [744](#)
 - callout IVP EAR file in WebSphere Liberty servers [746](#)
 - generic application server [735](#)
 - IVP EAR file [740](#)
 - IVP EAR file in WebSphere Application Server V7 [740](#)
 - IVP EAR file in WebSphere Application Server V8 [740](#)
 - IVP EAR file in WebSphere Liberty servers [741](#)
 - WebSphere Application Server [735](#)
 - WebSphere Liberty servers [737](#)
- dequeue count [759](#)
- DESCRIBE statement
 - column labels [514](#)
 - INTO clauses [514](#)
- description, segment [5](#)
- design efficiency, programs [461](#)
- design of
 - an application [15](#)
 - conversation [95](#)
 - local view [21](#)
- designing
 - terminal screen [94](#)
- determination of mappings [27](#)
- device input format (DIF), control block [93](#)
- device output format (DOF), control block [93](#)
- devices supported by MFS [444](#)
- DFSAPPC
 - DFSAPPC
 - format [406](#)
 - option keywords [406](#)
 - message switching [406](#)
 - DFSAPPC message switch [133](#)
 - DFSCONE0 (Conversational Abnormal Termination exit routine) [96](#)
 - DFSDDLTO (DL/I test program) [135](#)
 - DFSDLTRO (DL/I image capture). [156](#)

- DFSHALDB ddname
 - selective partition processing [279](#)
- DFSLI000 (language interface module)
 - binding COBOL code to [184](#)
- DFSMMA macro [54](#)
- DI_LEN call [325](#)
- diagnosing multiple parsing error return codes [429](#)
- DIB (DL/I interface block)
 - accessing information [477](#)
 - assembler language program
 - DIB fields [477](#)
 - variable names, mandatory [477](#)
 - BA status code [477](#)
 - BC status code [477](#)
 - C program
 - DIB fields [477](#)
 - variable names, mandatory [477](#)
 - CICS
 - HANDLE ABEND command [477](#)
 - COBOL program
 - DIB fields [477](#)
 - variable names, mandatory [477](#)
 - FH status code [477](#)
 - fields [477](#)
 - FW status code [477](#)
 - GA status code [477](#)
 - GB status code [477](#)
 - GD status code [477](#)
 - GE status code [477](#)
 - GG status code [477](#)
 - GK status code [477](#)
 - II status code [477](#)
 - label restriction [477](#)
 - labels [477](#)
 - LB status code [477](#)
 - NI status code [477](#)
 - PL/I
 - program variable names, mandatory [477](#)
 - restrictions
 - DIB label [477](#)
 - status codes
 - BA [477](#)
 - BC [477](#)
 - FH [477](#)
 - FW [477](#)
 - GA [477](#)
 - GB [477](#)
 - GD [477](#)
 - GE [477](#)
 - GG [477](#)
 - GK [477](#)
 - II [477](#)
 - LB [477](#)
 - NI [477](#)
 - TG [477](#)
 - structure [477](#)
 - translator version [477](#)
- DIB (DLI interface block) [11](#)
- dictionary, data [20](#)
- DIF (device input format), control block [93](#)
- direct access methods
 - characteristics [70](#)
 - HDAM [70](#)
 - HIDAM [72](#)
- direct access methods (*continued*)
 - PHDAM [69](#), [70](#)
 - PHIDAM [69](#), [72](#)
 - types of [70](#)
- direct dependents [36](#)
- directed routing [395](#)
- distributed data management (DDM) [560](#)
- distributed presentation management [446](#)
- Distributed Relational Database Access (DRDA) [559](#), [560](#)
- distributed relational database architecture (DRDA)
 - DDM commands [344](#)
 - overview [343](#)
- distributed security credential propagation
 - enabling in IMS TM resource adapter [810](#), [812](#)
 - propagating
 - in asynchronous callout messages [814](#)
 - in synchronous callout messages [814](#)
 - WebSphere Application Server [810](#)
 - WebSphere Liberty [812](#)
- Distributed Sync Point [109](#)
- DL/I
 - calls
 - for CICS and IMS programs [499](#)
- DL/I access methods
 - considerations in choosing [69](#)
 - DEDB [73](#)
 - direct access [70](#)
 - GSAM [75](#)
 - HDAM [70](#)
 - HIDAM [72](#)
 - HISAM [74](#)
 - HSAM [74](#)
 - MSDB [73](#)
 - PHDAM [69](#), [70](#)
 - PHIDAM [69](#), [72](#)
 - sequential access [74](#)
 - SHISAM [75](#)
 - SHSAM [75](#)
- DL/I call trace [135](#)
- DL/I calls
 - general information
 - coding [175](#)
 - image capture
 - batch job [138](#)
 - JDBC driver [628](#)
 - log data set
 - DFSERA50 call trace exit routine [138](#)
 - relationships to PCB types
 - I/O PCBs [362](#)
 - sample call formats [201](#), [204](#), [355](#), [357](#)
 - tracing
 - DLITRACE control statement [138](#)
 - image capture [136–138](#), [339](#)
 - IMS TRACE command [339](#)
 - TRACE command [137](#)
 - Universal JDBC driver [628](#)
- DL/I calls (general information)
 - qualification statements
 - overview [168](#)
 - qualified calls [168](#)
 - qualified SSAs (segment search arguments)
 - structure [168](#)
 - qualifying calls
 - command codes [174](#)

- DL/I calls (general information) *(continued)*
 - qualifying calls *(continued)*
 - field [168](#)
 - segment type [168](#)
 - segment search arguments (SSAs) [168](#)
 - SSAs (segment search arguments)
 - qualified [168](#)
 - unqualified calls [168](#)
- DL/I calls, system service
 - ROLB [264](#)
 - ROLL [264](#)
- DL/I calls, testing DL/I call sequences [135](#), [156](#)
- DL/I database
 - access to [56](#)
 - description [56](#)
- DL/I image capture (DFSDDLTR0) programs [156](#)
- DL/I interface block [477](#)
- DL/I language interfaces
 - overview [195](#), [349](#)
 - supported interfaces [195](#), [349](#)
- DL/I options
 - field level sensitivity [76](#)
 - logical relationships [81](#), [274](#)
 - secondary indexing [77](#), [271](#)
- DL/I program ROLB scenario [111](#)
- DL/I test program (DFSDDLTR0)
 - call statements [135](#)
 - checking program performance [135](#)
 - comments statements [135](#)
 - compare statements [135](#)
 - control statements [135](#)
 - description [135](#)
 - status statements [135](#)
 - testing DL/I call sequences [135](#), [156](#)
- DLET (Delete) call
 - with MSDB, DEDB or VSO DEDB [298](#)
- DLI
 - GUR call [234](#)
- DLITPLI [229](#)
- documentation for users [163](#)
- documentation of
 - data [20](#)
 - the application design process [15](#)
- DOF (device output format), control block [93](#)
- DPM (distributed presentation management)
 - using [446](#)
 - with ISC [446](#)
- DRA (Database Resource Adapter)
 - description [329](#)
 - startup table [329](#)
- DRDA (Distributed Relational Database Access) [559](#), [560](#)
- DRDA (distributed relational database architecture)
 - DDM commands [344](#)
 - overview [343](#)
- DS_LEN call [325](#)
- duplicate values, isolating [22](#)
- DYNAM option of COBOL [533](#)
- dynamic allocation [54](#), [67](#)
- dynamic backout [39](#), [486](#)
- dynamic MSDBs (main storage databases) [5](#)
- dynamic SQL
 - COBOL application program [533](#)
 - description [511](#)
 - EXECUTE [520](#)

- dynamic SQL *(continued)*
 - fixed-list SELECT statements [512](#)
 - non-SELECT statements [518](#), [520](#)
 - parameter marker [520](#)
 - PREPARE [520](#)
 - programming [511](#)
 - SELECT [519](#)
 - varying-list SELECT statements [514](#)

E

- EBCDIC [62](#)
- editing
 - considerations in your application [94](#)
 - messages
 - considerations in message and screen design [94](#)
 - overview [93](#)
- editing messages
 - edit routines
 - Basic Edit [381](#)
 - Intersystem Communication (ISC) Edit [381](#)
 - Message Format Service (MFS) [381](#)
- efficient program design [461](#)
- EIS
 - component-managed signon [802](#)
 - container-managed signon [799](#)
- EJB
 - process flow for synchronous callout [777](#)
 - synchronous callout sample [779](#)
 - transaction timeouts [819](#)
- elements
 - data, description [17](#)
 - data, naming [19](#)
- emergency restart [427](#)
- EMH (expedited message handler) [41](#)
- end a conversation, how to [95](#)
- end conversation request
 - conversational programs [820](#)
- end-of-data condition [547](#)
- enhanced STAT call formats for statistics
 - OSAM buffer subpool [143](#)
 - VSAM buffer subpool [148](#)
- enhancements
 - INQY DL/I call [661](#)
- enqueue count [759](#)
- Enterprise Information System (EIS)
 - component-managed signon [802](#)
 - container-managed signon [799](#)
- entity, data [17](#)
- entry and return conventions [227](#)
- entry point
 - AIB (application interface block)
 - address return [371](#)
 - and program entry statement [371](#)
 - assembler language
 - program entry [371](#)
 - register 1 at program entry [371](#)
- C language
 - __pcblist [371](#)
 - entry statement [371](#)
 - exit [371](#)
 - longjmp [371](#)
 - passing PCBs [371](#)
 - return [371](#)

- entry point (*continued*)
 - C language (*continued*)
 - system function [371](#)
 - CEETDLI
 - address return [371](#)
 - program entry statement [371](#)
 - COBOL
 - DLITCBL [371](#)
 - entry point
 - assembler language [371](#)
 - Java [371](#)
 - overview [371](#)
 - Pascal
 - entry statement [371](#)
 - passing PCBs [371](#)
 - PL/I
 - passing PCBs [371](#)
 - pointers in entry statement [371](#)
- environments
 - DB/DC [36](#)
 - DBCTL [36](#)
 - DCCTL [36](#)
 - options in [36](#), [56](#)
 - program and database types [36](#)
- equal-to relational operator [168](#)
- ERASE parameter [87](#)
- error
 - execution [159](#)
 - handling [523](#)
 - initialization [159](#)
 - return codes [522](#)
- ESTAE routines [53](#)
- example
 - current roster [17](#)
 - field level sensitivity [76](#)
 - instructor schedules [28](#)
 - instructor skills report [28](#)
 - local view [28](#)
 - logical relationships [81](#)
 - schedule of classes [28](#)
- examples
 - bank account database [5](#)
 - Boolean operators [173](#)
 - conversational processing [398](#)
 - D command code [174](#)
 - FLD/CHANGE [302](#)
 - FLD/VERIFY [302](#)
 - medical database [5](#)
 - multiple qualification statements [173](#)
 - path call [174](#)
 - UIB, defining [219](#)
- exception condition handling [523](#)
- exceptions
 - IMS TM resource adapter
 - HWSP1445E [867](#)
 - HWSSSL00E [868](#)
 - J2C applications [842](#)
 - J2C messages [842](#)
 - J2CA0056I [866](#)
 - other errors [866](#)
 - WLTC0017E [867](#)
 - output messages [837](#)
 - socket timeouts [817](#)
- EXEC DLI

- EXEC DLI (*continued*)
 - binder options, required [476](#)
 - compiler options, required [476](#)
 - DLI option [476](#)
 - preparing program for execution [475](#)
 - PROCESS statement overrides [476](#)
 - recovery commands
 - XRST (Extended Restart) [486](#)
 - translator options, required [476](#)
 - z/OS & VM [476](#)
 - z/OS & VM translator [476](#)
- EXEC DLI commands [11](#)
- EXEC DLI program translating [476](#)
- EXECUTE statement
 - dynamic execution [520](#)
 - parameter types [514](#)
- execution errors [159](#)
- execution timeouts
 - conversational transactions [815](#)
 - definition [815](#)
 - exceptions [815](#)
 - IMS Connect [815](#)
 - interaction properties [874](#)
 - specification [817](#)
 - valid values [815](#)
- existing application, converting an [15](#)
- explicit API for LU 6.2 devices [110](#)
- explicitly opening and closing a GSAM database [287](#)
- express alternate PCB [391](#)
- express PCBs [98](#)
- Extended Restart [48](#), [65](#)

F

- F command code
 - restrictions [320](#)
- failover [749](#)
- Fast Path
 - database calls [295](#), [296](#)
 - databases [36](#)
 - databases, processing [295](#)
 - DEDB (data entry database)
 - processing [295](#)
 - DEDB and the PROCOPT operand [87](#)
 - IFPs [41](#)
 - MSDB (main storage database)
 - processing [295](#)
 - P (position) processing option [498](#)
 - secondary index, using with DEDBs [308](#)
 - subset pointers with DEDBs [487](#)
 - subset pointers, using with DEDBs [304](#)
 - types of databases [295](#)
- FETCH statement
 - description, single row [547](#)
 - host variables [512](#)
- field
 - changing contents [301](#)
 - checking contents: FLD/VERIFY [299](#)
- Field (FLD) call [298](#)
- field level sensitivity
 - as a security mechanism [85](#)
 - defining [9](#)
 - description [76](#)
 - example [76](#)

- field level sensitivity (*continued*)
 - specifying [76](#)
 - uses [76](#)
- field name
 - FSA [300](#)
 - SSA
 - qualification statement [168](#)
- field search argument (FSA)
 - description [299](#)
 - with DL/I calls [299](#)
- field value
 - FSA [300](#)
 - SSA qualification statement [168](#)
- fields
 - columns, compared to [555](#)
 - in SQL queries [555](#)
- File Select and Formatting Print Program (DFSERA10) [48](#)
- FIN (Finance Communication System)
 - defining to operate with MFS [444](#)
- fixed, MSDBs (main storage databases) [5](#)
- FLD (Field) call
 - description [298](#)
 - FLD/CHANGE
 - [301](#)
 - FLD/VERIFY [299](#)
- flow diagrams, LU 6.2
 - CPI-C driven commit scenario [111](#)
 - DFSAPPC, synchronous SL=none [111](#)
 - DL/I program backout scenario [111](#)
 - DL/I program commit scenario [111](#)
 - local CPI communications driven program, SL=none [111](#)
 - local IMS Command
 - asynchronous SL=confirm [111](#)
 - local IMS command, SL=none [111](#)
 - local IMS conversational transaction, SL=none [111](#)
 - local IMS transaction
 - asynchronous SL=confirm [111](#)
 - asynchronous SL=none [111](#)
 - synchronous SL=confirm [111](#)
 - synchronous SL=none [111](#)
 - multiple transactions in same commit [111](#)
 - remote MSC conversation
 - asynchronous SL=confirm [111](#)
 - asynchronous SL=none [111](#)
 - synchronous SL=confirm [111](#)
 - synchronous SL=none [111](#)
- formats
 - PSB [483](#)
- formatting
 - result tables [543](#)
- frequency, checkpoint [51](#)
- FROM clause
 - SELECT statement [541](#)
- FSA (field search argument)
 - description [299](#)
 - with DL/I calls [299](#)
- full-function databases
 - and the PROCOPT operand [87](#)
 - how accessed, CICS [56](#)
 - how accessed, IMS [36](#)

G

- gather requirements

- gather requirements (*continued*)
 - for conversational processing [95](#)
- gathering requirements
 - for database options [69](#)
 - for message processing options [91](#)
- GC status code [498](#)
- general programming guidelines [461](#)
- Generalized Sequential Access Method (GSAM)
 - program access [283](#)
- generalized sequential access method (GSAM))
 - DB PCB (database program communication block)
 - mask [226](#)
- global transaction support
 - TCP/IP [826](#)
- GO processing option [49](#)
- GO TO clause of WHENEVER statement [523](#)
- GPSB (generated program specification block)
 - format [373](#)
- greater-than relational operator [168](#)
- greater-than-or-equal-to relational operator [168](#)
- group data elements
 - into hierarchies [22](#)
 - with their keys [22](#)
- group names [870](#)
- GSAM (generalized sequential access method)
 - GSAM (generalized sequential access method)
 - RSA [285](#)
 - record search argument [285](#)
- GSAM (generalized sequential access method)
 - accessing databases [283](#)
 - BMP region type [294](#)
 - call summary [289](#)
 - CHKP [289](#)
 - coding considerations [289](#)
 - data areas [226](#)
 - data set
 - attributes, specifying [293](#)
 - characteristics, origin [290](#)
 - concatenated [293](#)
 - DD statement DISP parameter [291](#)
 - extended checkpoint restart [292](#)
 - database, explicitly opening and closing [287](#)
 - DB PCB masks [226](#)
 - DBB region type [294](#)
 - description [283](#)
 - designing a program [283](#)
 - DLI region types [294](#)
 - fixed-length records [287](#)
 - I/O areas [288](#)
 - record formats [287](#)
 - records, retrieving and inserting [285](#)
 - restrictions on CHKP and XRST [289](#)
 - RSA [226](#)
 - RSA (record search argument)
 - description [285](#)
 - status codes [288](#)
 - undefined-length records [287](#)
 - variable-length records [287](#)
 - XRST [289](#)
- GSAM (Generalized Sequential Access Method)
 - accessing GSAM databases [56](#)
 - database type [36](#)
 - description [75](#)
- GSAM PCB [483](#)

guidelines, general programming [461](#)
guidelines, programming [167](#)
GUR call [234](#)

H

H processing option [319](#)
HALDB
 DFSHALDB ddname [279](#)
 HALDB control statement [279](#)
 partitions
 selective partition processing [279](#)
 selective partition processing [279](#)
HALDB (High Availability Large Database)
 application programs
 scheduling against [233](#)
 initial load [233](#)
HALDB control statement [279](#)
HDAM
 multiple qualification statements [173](#)
HDAM (Hierarchical Direct Access Method) [70](#)
HIDAM (Hierarchical Indexed Direct Access Method) [72](#)
hierarchical database
 example [555](#)
 relational database, compared to [555](#)
hierarchical database example, medical [5](#)
Hierarchical Direct Access Method (HDAM) [70](#)
Hierarchical Indexed Direct Access Method (HIDAM) [72](#)
Hierarchical Indexed Sequential Access Method (HISAM) [74](#)
Hierarchical Sequential Access Method (HSAM) [74](#)
hierarchy
 bank account database [5](#)
 data structures [175](#)
 description [5](#)
 grouping data elements [22](#)
 medical database [5](#)
hierarchy examples [5](#)
High Availability Large Database (HALDB)
 application programs
 scheduling against [233](#)
 initial load [233](#)
HISAM (Hierarchical Indexed Sequential Access Method) [74](#)
homonym, data element [19](#)
Hospital database example [568](#)
host names [870](#)
host structure
 COBOL [529](#)
 description [511](#)
 indicator structure [511](#)
host variable
 COBOL [525](#)
 description [510](#)
 FETCH statement [512](#)
 indicator variable [511](#)
 PREPARE statement [512](#)
 using [513](#)
host variables [510](#)
HOUSHOLD segment [5](#)
HSAM (Hierarchical Sequential Access Method) [74](#)

I

I/O area

I/O area (*continued*)
 command-level program [481](#)
 specifying [221](#), [370](#)
 XRST [486](#)
I/O PCB
 in different environments [58](#)
 mask
 12-byte time stamp [208](#)
 general description [208](#)
 group name field [208](#)
 input message sequence number [208](#)
 logical terminal name field [208](#)
 message output descriptor name [208](#)
 specifying [208](#)
 status code field [208](#)
 userid field [208](#)
 userid indicator field [208](#)
I/O PCB mask
 general description [363](#)
 specifying [363](#)
IBM Enterprise COBOL for z/OS
 Java dependent region
 interoperability [702](#)
 Java dependent regions
 backend application for Java applications [703](#)
 frontend application for Java applications [703](#)
 issuing DL/I calls in COBOL [703](#)
ICAL callout with control data
 IMS Java dependent region resource adapter [690](#)
identification of
 recovery requirements [49](#)
identifying
 application data [17](#)
 online security requirements [91](#)
 output message destinations [97](#)
 security requirements [84](#)
IDEs (integrated development workbench)
 IMS TM resource adapter applications [751](#)
IDs, checkpoint [62](#)
IFP (IMS Fast Path) program
 databases that can be accessed [36](#)
 differences from an MPP [41](#)
 recovery [41](#)
 restrictions [41](#)
IFP application programs [483](#)
ILLNESS segment [5](#)
image capture program
 CICS application program [156](#)
 IMS application program [136](#)
immediate program switch [95](#)
implicit API for LU 6.2 devices [110](#)
IMS application
 diagnosing
 abnormal termination (abend) [151](#)
 program execution errors [151](#)
 program initialization errors [151](#)
IMS application programs, standard [408](#)
IMS catalog
 application programming [234](#)
 PSBs [234](#)
IMS commands
 submitting [791](#)
IMS connection factories
 configuration [791](#)

IMS connection factories (*continued*)

overview [793](#), [828](#)

properties

clientID [869](#)

CM0Dedicated [869](#)

dataStoreName [870](#)

groupName [870](#)

hostName [870](#)

password [870](#)

passwordPhrase [870](#)

portNumber [870](#)

SSLEnabled [870](#)

SSLEncryptionType [870](#)

SSLKeyStoreName [871](#)

SSLKeyStorePassword [871](#)

SSLProtocol [872](#)

SSLTrustStoreName [872](#)

SSLTrustStorePassword [872](#)

userName [872](#)

IMS conversations

conversational program [398](#)

nonconversational program [398](#)

IMS coprocessor

processing SQL statements [548](#)

IMS database

batch programming [180](#)

database design

logical relationships [83](#)

IMS DB

Application programming [165](#)

IMS Explorer for Development

com.ibm.ims.db.DLIDatabaseView class

generating [567](#)

DLIDatabaseView class

generating [567](#)

Java metadata class

generating [567](#)

IMS Fast Path (IFP) programs, description of [41](#)

IMS interaction specification

asynchronous callout sample [781](#)

client-managed conversation state [823](#)

configuration [790](#)

IMS Connect-managed conversation state [823](#)

properties

altClientID [872](#)

asyncOutputAvailable [873](#)

calloutRequestType [873](#)

CM0Response [874](#)

commitMode [874](#)

convEnded [873](#)

convID [873](#)

executionTimeout [874](#)

ignorePURGCall [875](#)

imsRequestType [875](#)

interactionVerb [876](#)

ltermName [877](#)

mapName [877](#)

purgeAsyncOutput [878](#)

reRoute [878](#)

reRouteName [878](#)

resumeTpPipeNSC [879](#)

socketTimeout [879](#)

syncCalloutCorrelationToken [879](#)

syncCalloutStatusCode [879](#)

IMS interaction specification (*continued*)

properties (*continued*)

syncLevel [879](#)

transExpiration [880](#)

useConvID [880](#)

synchronous callout [760](#)

synchronous callout sample [778](#)

IMS Java dependent region resource adapter

Java batch processing (JBP) regions [674](#)

Java message processing (JMP) regions [674](#)

support for ICAL callout with control data [690](#)

IMS request types [875](#)

IMS solutions for Java development

overview [553](#), [566](#)

IMS Spool API

dynamic allocation

error messages [432](#)

print data sets

CHNG call [433](#)

z/OS services for Dynamic Output (SVC109)

[433](#)

IMS Spool API application design [425](#)

IMS support for DRDA [559](#)

IMS TM

application program

message Type [377](#)

DB2 considerations [390](#), [423](#)

IMS TM resource adapter

components [723](#)

connection factories [791](#)

conversational programs [820](#), [821](#)

diagnosis

callout requests [837](#)

installation verification program [835](#)

Java applications access to IMS [836](#)

distributed security credentials

in asynchronous callout messages [814](#)

in synchronous callout messages [814](#)

enabling distributed security credentials [810](#), [812](#)

enabling network security credentials [810](#), [812](#)

functions [723](#)

IMS Connect interactions [791](#)

logging [838–840](#)

migration [732](#)

network security credentials

in asynchronous callout messages [814](#)

in synchronous callout messages [814](#)

overview [723](#)

runtime process [724](#)

supported platforms [726](#)

supported software configurations [726](#)

supported versions [726](#)

tracing [838–840](#)

update installation [748](#)

IMS Transaction Manager (IMS TM)

application development [751](#)

interactions with [751](#)

IMS Universal Database resource adapter

CCI programming interface [584](#)

configuring SSL support

container-managed environment [669](#)

WebSphere Application Server [669](#)

connectivity

RRSLocalOption [573](#)

IMS Universal Database resource adapter (*continued*)

- connectivity (*continued*)
 - type-2 [573](#)
 - type-4 [573](#)
- creating a CCI Connection
 - managed environment [573](#)
- creating a CCI ConnectionFactory
 - managed environment [573](#)
- DLIInteractionSpec class
 - deleting data [585](#)
 - inserting data [585](#)
 - retrieving data [585](#)
 - updating data [585](#)
- JNDI lookup
 - connecting to IMS [573](#)
- logging [670](#)
- overview [571](#)
- sample application [584](#)
- specifying IMSConnectionSpec properties
 - managed environment [573](#)
- SQLInteractionSpec class
 - deleting data [589](#)
 - inserting data [589](#)
 - retrieving data [589](#)
 - updating data [589](#)
- supported software configurations [572](#)
- tracing [670](#)
- transaction management
 - bean managed [571](#)
 - container managed [571](#)
 - local transaction support [571](#)
 - LocalTransaction interface [571](#)
 - UserTransaction interface [571](#)
 - XA transaction support [571](#)
- WebSphere Application Server for distributed platforms
 - support [573](#)
- WebSphere Application Server for z/OS support [573](#)

IMS Universal DL/I driver

- Accessing
 - dynamic arrays [667](#)
- adding segments
 - example [656](#)
- AIB
 - example [664](#)
- AIB interface [645](#)
- application programming [642](#)
- batchDelete
 - example [660](#)
- batchRetrieve
 - example [651](#)
- batchUpdate
 - example [658](#)
- ByteBuffers
 - using [663](#)
- com.ibm.ims.base [642](#)
- com.ibm.ims.dli [642](#)
- commit
 - example [665](#)
- configuring SSL support
 - stand-alone environment [669](#)
- connecting
 - IMS database [643](#)
- connections
 - IMS database [643](#)

IMS Universal DL/I driver (*continued*)

- connections (*continued*)
 - properties [643](#)
- create
 - example [656](#)
- creating segments
 - example [656](#)
- DBArrayElementSet class [667](#)
- DBPCB
 - example [664](#)
- DBPCB interface [645](#)
- delete
 - example [659](#)
- deleting multiple segments
 - example [660](#)
- deleting segments
 - batch [660](#)
 - example [659](#), [660](#)
- DL/I DLET [659](#)
- DL/I ISRT [656](#)
- DL/I REPL [658](#)
- Dynamic arrays in
 - IMS Universal DL/I driver [667](#)
- fetch size
 - setting [652](#)
- getNext
 - example [649](#)
- getNextWithinParent [649](#)
- getPathForBatchUpdate
 - example [658](#)
- getPathForInsert
 - example [656](#)
- getPathForRetrieveReplace
 - example [649](#), [658](#), [659](#)
- getUnique
 - example [649](#)
- GSAMPCB interface [645](#)
- IMSConnectionSpec
 - creating [643](#)
 - example [643](#)
- IMSConnectionSpecFactory
 - example [643](#)
- IMSStatusCodes [664](#)
- insert
 - example [656](#)
- inserting segments
 - example [656](#)
- interfaces [645](#)
- Java packages [642](#)
- logging [670](#)
- overview [641](#)
- Path
 - data transformation [636](#), [653](#)
 - retrieving java.sql data types [636](#), [653](#)
- Path interface [645](#)
- PathSet interface [645](#)
- PCB interface [645](#)
- programming model [642](#)
- PSB
 - creating [643](#)
 - example [643](#)
- PSB interface [645](#)
- PSBFactory
 - example [643](#)

IMS Universal DL/I driver (*continued*)

- query performance
 - improving [652](#)
- replace
 - example [658](#)
- retrieving
 - error code extension [664](#)
 - reason code [664](#)
 - return code [664](#)
 - status code [664](#)
- retrieving data
 - example [649](#)
- retrieving multiple segments
 - example [651](#)
- retrieving segments
 - batch [651](#)
 - example [649](#), [651](#)
- rollback
 - example [665](#)
- segment search arguments
 - specifying [646](#)
- setFetchSize [652](#)
- SSAList
 - adding initial qualification [646](#)
 - appending additional qualifications [646](#)
 - creating [646](#)
 - debugging [646](#)
 - qualified [646](#)
 - setting command codes [646](#)
 - setting lock classes [646](#)
 - unqualified [646](#)
- SSAList interface [645](#)
- SSAs
 - specifying [646](#)
- tracing [670](#)
- transactions
 - example [665](#)
 - local [665](#)
 - one-phase commit [665](#)
 - processing [665](#)
 - scope [665](#)
 - unit of recovery [665](#)
 - unit of work [665](#)
- updating multiple segments
 - example [658](#)
- updating segments
 - batch [658](#)
 - example [658](#)
- Using Byte Buffers [663](#)

IMS Universal drivers

- application platforms [564](#)
- architecture [560](#)
- CICS support [560](#)
- com.ibm.ims.db.DLIDatabaseView class
 - generating [567](#)
- configuring JAXB [563](#)
- configuring SSL support
 - container-managed environment [669](#)
 - stand-alone environment [669](#)
 - WebSphere Application Server [669](#)
- connectivity
 - distributed (type-4) [560](#)
 - local (type-2) [560](#)
- data access methods [564](#)

IMS Universal drivers (*continued*)

- Db2 for z/OS stored procedures support [560](#)
 - DLIDatabaseView class
 - generating [567](#)
 - Java dependent region support [560](#)
 - Java metadata class
 - generating [567](#)
 - JBP region support [560](#)
 - JMP region support [560](#)
 - overview [559](#)
 - programming approaches [564](#)
 - SSL support
 - container-managed environment [668](#)
 - stand-alone environment [668](#)
 - transaction processing options [564](#)
 - variable length database segments [565](#)
 - WebSphere Application Server for distributed platforms support [560](#)
 - WebSphere Application Server for z/OS support [560](#)
- ## IMS Universal JCA/JDBC driver
- connecting [579](#)
 - data operations
 - DELETE [592](#)
 - INSERT [592](#)
 - PreparedStatement [592](#)
 - SELECT [592](#)
 - Statement [592](#)
 - syntax [592](#)
 - UPDATE [592](#)
 - deploying [579](#)
- ## IMS Universal JDBC driver
- columns compared to fields [613](#)
 - configuring SSL support
 - stand-alone environment [669](#)
 - connecting to IMS
 - DataSource [595](#)
 - DriverManager [601](#)
 - DDL-specific SQL usage
 - ALTER statement [622](#)
 - CREATE statement [621](#)
 - DROP statement [622](#)
 - foreign key fields
 - example [619](#)
 - SQL statement usage [619](#)
 - hierarchical databases compared to relational databases [613](#)
 - IMS-specific SQL usage
 - AS clause [622](#)
 - DELETE statement [625](#)
 - DISTINCT clause [622](#)
 - FROM clause [622](#)
 - GROUP BY clause [622](#)
 - INSERT statement [624](#)
 - ORDER BY clause [622](#)
 - SELECT statement [622](#)
 - UPDATE statement [625](#)
 - WHERE clause [626](#)
- ## IMSCALLOUTAPI [692](#)
- interfaces
 - DataSource [595](#)
 - DriverManager [601](#)
 - JDBC programming interface [607](#)
 - logging [670](#)
 - programming with the IMSCALLOUTAPI [692](#)

- IMS Universal JDBC driver (*continued*)
 - rows compared to segment instances [613](#)
 - sample application [607](#)
 - supported drivers [594](#)
 - tables compared to segments [613](#)
 - tracing [670](#)
 - using byte buffers with IMS Universal JDBC Driver [608](#)
 - using the `ArrayResultSet` property [611](#)
 - using the `removeInvalidCaseFields` property [610](#)
- IMSActivationSpec properties
 - configuring [745](#), [768](#)
 - list [770](#)
- IMSInteractionSpec properties
 - list [872](#)
- in-doubt UOR
 - definition [260](#)
- in-flight UOR
 - definition [260](#)
- independent AND operator [272](#)
- indexed field in SSA [271](#)
- indexing, secondary
 - DL/I Returns [274](#)
 - effect on program [271](#)
 - multiple qualification statements [272](#)
 - status codes [274](#)
- indicator structure
 - description [511](#)
- indicator variable
 - description [511](#)
- INIT system service call [52](#)
- initialization errors [159](#)
- input for a DL/I program [175](#)
- input message
 - format [378](#)
 - MFS [384](#)
- INQY DL/I call name enhancement [661](#)
- INQY system service call [52](#)
- INSERT statement
 - description [538](#)
 - single row [538](#)
 - VALUES clause [538](#)
- inserting a segment
 - GSAM records [285](#)
- installation
 - distributed platforms [733](#)
 - IMS TM adapter runtime component [731](#)
 - preparation [731](#), [732](#)
 - verification [739](#)
 - z/OS platform [734](#)
- instructor
 - schedules [28](#)
 - skills report [28](#)
- integrity
 - batch programs [264](#)
 - how DL/I protects data [61](#)
 - maintaining database [264](#)
 - read without [89](#)
 - using ROLB
 - MPPs and transaction-oriented BMPs [264](#)
 - using ROLL [264](#)
 - using ROLS [264](#)
- interaction verbs [876](#)
- interface, AIB [11](#)
- intermediate backout point

- intermediate backout point (*continued*)
 - backing out [267](#)
- intermediate backout points [486](#)
- interoperability
 - 31-bit COBOL and 64-bit Java [705](#)
- Introduction to Resource Recovery [105](#)
- invalid processing and ROLB/SETS/ROLLS calls [96](#)
- IPDS and IMS Spool API [427](#)
- ISC (intersystem communication)
 - defining to operate with MFS [446](#)
- ISC (Intersystem Communication) [41](#)
- isolation of
 - duplicate values [22](#)
 - repeating data elements [22](#)
- ISRT (Insert) call
 - with MSDB, DEDB or VSO DEDB [298](#)
- ISRT call
 - issuing to other terminals [391](#)
 - message call
 - in conversational programs [399](#)
 - referencing alternate PCBs [391](#)
 - usage [391](#)
- ISRT system service call [425](#)
- issue checkpoints [39](#)
- issuing
 - checkpoints in batch or BMP programs [485](#)
- IVP (installation verification program)
 - as EAR file [740](#), [741](#)
 - prerequisites to running [739](#)
 - running [742](#)
 - troubleshooting [835](#)
- IVP EAR file [747](#)

J

- J2C applications
 - adapters
 - activation specifications [767](#)
 - Common Client Interface (CCI) [751](#)
- Java API reference [881](#)
- Java Authentication and Authorization Service (JAAS) [797](#)
- Java Batch Processing (JBP)
 - applications [45](#)
 - databases that can be accessed [36](#)
- Java batch processing (JBP) application
 - accessing GSAM data [686](#)
 - program switching
 - immediate [694](#)
 - programming models [683](#)
 - restart [683](#)
 - symbolic checkpoint [683](#)
- Java batch processing (JBP) regions
 - Db2 for z/OS access
 - application programming [703](#)
 - IMS Java dependent region resource adapter [674](#)
 - JBP applications [673](#)
 - overview [673](#)
- Java data bindings
 - input message formats [793](#)
 - output message formats [793](#)
- Java Database Connectivity (JDBC) [559](#)
- Java EE Connector Architecture (JCA) [559](#)
- Java Message Processing (JMP)
 - applications [45](#)

Java Message Processing (JMP) *(continued)*

- databases that can be accessed [36](#)

Java message processing (JMP) application

- Db2 for z/OS data access [677](#)

- IMS data access [677](#)

- IMSFieldMessage class

- subclassing [675](#)

- input messages

- defining [675](#)

- message handling

- input messages [678](#)

- output messages

- defining [675](#)

- program switching

- deferred [696](#)

- immediate [694](#)

- programming models [677](#)

- transactions

- commit [678](#)

- rollback [678](#)

Java message processing (JMP) applications

- message handling

- conversational transactions [679](#)

- multi-segment messages [680](#)

- multiple input messages [681](#)

- repeating structures [681](#)

- scratchpad area (SPA) [679](#)

Java message processing (JMP) regions

- Db2 for z/OS access

- application programming [703](#)

- IMS Java dependent region resource adapter [674](#)

- JMP applications [673](#)

- overview [673](#)

Java metadata class

- IMS Explorer for Development

- generating [567](#)

- IMS Universal drivers [567](#), [632](#)

JBP (Java Batch Processing)

- applications [45](#)

- databases that can be accessed [36](#)

JBP (Java batch processing) regions

- Db2 for z/OS access

- application programming [703](#)

- IMS Java dependent region resource adapter [674](#)

- synchronous callout support [692](#)

- synchronous program switch support [697](#)

JCA (Java EE Connector Architecture) [559](#)

JDBC

- ARRAY [635](#)

- BIGINT [635](#)

- BINARY [635](#)

- BIT [635](#)

- CHAR [635](#)

- CLOB [635](#)

- DATE [635](#)

- DOUBLE [635](#)

- FLOAT [635](#)

- IMS-specific SQL usage

- WHERE clause subfield support [618](#), [627](#)

- INTEGER [635](#)

- mapping SQL data types to Java [635](#)

- overview [594](#)

- PACKEDDECIMAL [635](#)

- ResultSet

JDBC *(continued)*

- ResultSet *(continued)*

- data transformation [636](#), [653](#)

- retrieving java.sql data types [636](#), [653](#)

- SMALLINT [635](#)

- SQL aggregate function

- AS [615](#)

- ASC [615](#)

- AVG [615](#)

- COUNT [615](#)

- DESC [615](#)

- GROUP BY [615](#)

- MAX [615](#)

- MIN [615](#)

- ORDER BY [615](#)

- SUM [615](#)

- SQL aggregate functions supported [615](#)

- SQL keywords supported [613](#)

- STRUCT [635](#)

- TIME [635](#)

- TIMESTAMP [635](#)

- TINYINT [635](#)

- Universal drivers

- portable SQL keywords restrictions [616](#)

- XML support

- Java metadata class [632](#)

- overview [631](#)

- retrieval [634](#)

- SQL INSERT [632](#)

- SQL SELECT [634](#)

- storage [632](#)

- type-4 connectivity [631](#)

- ZONEDDECIMAL [635](#)

JDBC (Java Database Connectivity) [559](#)

JDBC driver

- DL/I calls [628](#)

JES Spool/Print server [427](#)

JMP (Java Message Processing)

- applications [45](#)

- databases that can be accessed [36](#)

JMP (Java message processing) regions

- Db2 for z/OS access

- application programming [703](#)

- IMS Java dependent region resource adapter [674](#)

- synchronous callout support [692](#)

- synchronous program switch support [697](#)

JNDI

- configuration in applications [830](#)

- configuration in WebSphere Application Server [745](#), [766](#), [768](#)

JOURNAL parameter [426](#)

K

- key feedback area

- command-level program [480](#)

- key sensitivity [85](#)

- keyboard shortcuts [xix](#)

- keys, data [22](#)

- keystores

- configuring for IMS TM resource adapter [807](#)

- creating [807](#)

L

- label, column [514](#)
- Language Environment
 - characteristics of CEETDLI [232, 374](#)
 - interoperability, COBOL and Java [705](#)
 - LANG= Option on PSBGEN for PL/I Compatibility with Language Environment [232, 374](#)
 - Language Environment
 - LANG = option for PL/I compatibility [232, 374](#)
 - supported languages [232, 374](#)
- Language Environment, with IMS [232, 374](#)
- Large Data Sets [293](#)
- legal notices
 - notices [883](#)
 - trademarks [883, 884](#)
- less-than relational operator [168](#)
- less-than-or-equal-to relational operator [168](#)
- limit access with signon security [91](#)
- link editing, EXEC DLI [476](#)
- link to another online program [60](#)
- LIST parameter [137](#)
- listing data elements [17](#)
- LL field
 - in input message [378](#)
 - in output message [379](#)
- Local Option
 - 64-bit support [724](#)
 - container-managed security [799](#)
 - restrictions [727](#)
 - shareable persistent connections [791](#)
 - supported communication protocol [728](#)
 - tpipe names [759](#)
 - WebSphere Application Server on z/OS [728](#)
- local view
 - designing [21](#)
 - examples [28](#)
- locating
 - a specific sequential dependent [497](#)
 - last inserted sequential dependent [497](#)
- locating dependents in DEDBs
 - last-inserted sequential dependent, POS call [316](#)
 - POS call [316](#)
 - specific sequential dependent, POS call [316](#)
- lock management [270](#)
- locking protocol [87](#)
- LOCKMAX= parameter, BMP programs [49](#)
- log
 - records
 - sync points [261](#)
- LOG call
 - description [151](#)
 - use in monitoring [159](#)
- log records
 - type 18 [62](#)
 - X'18' [48](#)
- LOG system service call [340](#)
- log, system [39](#)
- logging
 - for IMS TM resource adapter [838, 840](#)
 - in WebSphere Application Server [838](#)
 - in WebSphere Liberty servers [839](#)
 - standalone logger [839](#)

- logical child [274](#)
- logical parent [274](#)
- logical relationships
 - defining [81](#)
 - description [81](#)
 - effect on programming [276](#)
 - example [81](#)
 - introduction [274](#)
 - logical child [274](#)
 - logical parent [274](#)
 - physical parent [274](#)
 - processing segments [274](#)
 - programming, effect [274](#)
 - status codes [277](#)
- logical structure [274](#)
- lterm names [877](#)
- LTERM, local and remote [133](#)
- LU 6.2
 - conversations [407](#)
 - support for APPC [101](#)
- LU 6.2 devices, signon security [91](#)
- LU 6.2 partner program design
 - DFSAPPC message switch [133](#)
 - flow diagrams [111](#)
 - integrity after conversation completion [131](#)
 - scenarios [111](#)
- LU 6.2 User Edit Exit
 - using [390](#)

M

- macros
 - DATABASE [89](#)
 - DFSMDA [54](#)
 - TRANSACT [44](#)
- main storage database (MSDB) [73](#)
- main storage database (MSDBs)
 - types
 - nonrelated [5](#)
- main storage databases (MSDBs)
 - dynamic [5](#)
 - types
 - related [5](#)
- maintaining database integrity [485](#)
- managed connections
 - connection factories [791](#)
 - JNDI lookups [791](#)
- managing subset pointers in DEDBs with command codes [296](#)
- many-to-many mapping [27](#)
- map names [877](#)
- mapped conversation, APPC [104](#)
- mappings, determining [27](#)
- mask
 - AIB [214](#)
 - DB PCB [212](#)
- mask, data [11](#)
- master terminal
 - issuing timeout [395](#)
- MDBs (message-driven beans)
 - callout sample [774](#)
 - configuring IMSActivationSpec property [745, 768](#)
 - J2C activation specification configuration [767](#)
 - retrieving callout messages [766](#)

- MDBs (message-driven beans) (*continued*)
 - thread pool [772](#)
- medical database example
 - description [5](#)
 - segments [5](#)
- message
 - editing
 - description [381](#)
 - input message [382](#), [384](#)
 - output [382](#)
 - output message [390](#)
 - skipping line [382](#)
 - using Basic Edit [382](#)
 - using ISC Edit [383](#)
 - using LU 6.2 User Edit Exit [390](#)
 - using MFS Edit [383](#)
 - from terminals [377](#)
 - I/O PCB [381](#)
 - input [378](#), [384](#)
 - input descriptor (MID), control block [93](#)
 - input fields
 - contents [378](#)
 - ISC (intersystem communication)
 - editing output messages [383](#)
 - ISC (intersystem communication) edit
 - output message [383](#)
 - message formatting service [382](#)
 - MFS (Message Format Service)
 - editing message [382](#)
 - obtaining text
 - COBOL [533](#)
 - output [97](#), [379](#), [390](#)
 - output descriptor (MOD), control block [93](#)
 - output fields
 - contents [379](#)
 - printing [382](#)
 - processing of
 - summary [380](#)
 - processing options [91](#)
 - receiving by program [377](#)
 - result [381](#)
 - sending to other application programs [393](#)
 - type
 - message switch service [377](#)
 - types
 - another terminal [377](#)
- Message Format Service [437](#)
- Message Format Service (MFS)
 - control blocks
 - relationship with screen format [442](#)
 - LU 6.2 device restriction [383](#)
 - secondary logical unit (SLU) [383](#)
 - terminal
 - message processing program (MPP) [383](#)
- Message Input
 - Segment Format [378](#)
- message processing options
 - sending message to originating terminal [97](#)
- message processing program [418](#)
- message-driven beans (MDBs)
 - callout sample [774](#)
 - configuring IMSActivationSpec property [768](#)
 - J2C activation specification configuration [767](#)
 - retrieving callout messages [766](#)
- message-driven programs
 - definition [416](#)
 - supported message destinations [416](#)
 - usage restrictions [416](#)
- messages
 - IMS TM resource adapter [842](#)
 - WebSphere Application Server
 - HWSP1445E [867](#)
 - HWSSSL00E [868](#)
 - J2CA0056I [866](#)
 - other messages [866](#)
 - WLTC0017E [867](#)
- methods of data storage
 - combined file [3](#)
 - database [3](#)
 - separate files [3](#)
- MFS (Message Format Service)
 - components [444](#)
 - control blocks
 - relationship with screen format [442](#)
 - editing output messages [383](#)
 - example [441](#)
 - input message
 - formats [384](#)
 - MFS (message format service)
 - message editor [444](#)
 - online performance [437](#)
 - output message
 - formats [390](#)
 - overview [93](#), [437](#)
 - pool manager [444](#)
 - remote programs [444](#)
 - supported devices [444](#)
- MFS control blocks
 - DIF (device input format)
 - description [438](#)
 - DOF (device output format)
 - description [438](#)
 - MID (message input descriptor)
 - description [438](#)
 - MOD (message output descriptor)
 - description [438](#)
 - summary [438](#)
- MFS libraries
 - online change function [444](#)
- MFS message output descriptors (MODs) [877](#)
- MFS SOA support [724](#)
- MFSTEST procedure (language utility)
 - pool manager [444](#)
- MID (message input descriptor), control block [93](#)
- migration
 - IMS TM resource adapter [732](#)
- mixed-language programming [233](#), [375](#)
- MOD (message output descriptor), control block [93](#)
- mode
 - multiple [49](#)
 - processing [45](#)
 - response [98](#)
 - single [49](#)
- MODE parameter [45](#)
- modifying
 - data [538](#)
- MOVENEXT option
 - examples [489](#)

- MOVENEXT option (*continued*)
 - use when moving subset pointer forward [489](#)
- moving subset pointer forward [489](#)
- MPP (message processing program)
 - coding in assembler language [418](#)
 - coding in C language [418](#)
 - coding in COBOL [419](#)
 - coding in Pascal [420](#)
 - coding in PL/I [422](#)
 - coding necessary information [417](#)
 - databases that can be accessed [36](#), [40](#)
 - description [40](#)
 - executing [40](#)
 - input [417](#)
 - parmcount [422](#)
 - PL/I
 - entry statement restrictions [422](#)
 - MPP coding notes [422](#)
 - optimizing compiler [422](#)
- MPP regions
 - troubleshooting callout issues [837](#)
 - two-phase commit applications [827](#)
- MPPs
 - ROLB [264](#)
- MSC (multiple systems coupling)
 - description [395](#)
 - directed routing [395](#)
 - receiving messages from other IMS TM systems [396](#)
 - sending messages to other IMS TM systems [397](#)
- MSDB (main storage database)
 - call restrictions [297](#)
 - commit point processing [302](#)
 - updating segments [298](#)
- MSDBs (main storage database)
 - processing commit points [302](#)
- MSDBs (main storage databases)
 - nonrelated [297](#)
 - terminal related [297](#)
 - types
 - description [297](#)
 - nonrelated [5](#)
 - related [5](#)
- multi-segment messages [793](#)
- multiple
 - DB PCBs [255](#)
 - processing [249](#)
 - qualification statements
 - DEDB [173](#)
 - HDAM [173](#)
 - PHDAM [173](#)
- multiple ISRT calls [875](#)
- multiple mode [45](#), [49](#)
- multiple positioning
 - advantages of [253](#)
 - effecting your program [253](#)
 - resetting position [255](#)
- multiple PURG calls [875](#)
- multiple systems coupling [395](#)
- multiple-row FETCH statement
 - SQLIMSCODE +100 [522](#)
- MVS SJF (Scheduler JCL Facility) [425](#)
- MYLTERM [297](#)

N

- names of data elements [19](#)
- naming convention
 - COBOL [533](#)
- naming conventions [15](#)
- NDM (Non-Discardable Messages) routine [45](#)
- network security credentials
 - asynchronous callout [814](#)
 - IMS TM resource adapter [810](#), [812](#), [814](#)
 - propagating [810](#), [812](#), [814](#)
 - synchronous callout [814](#)
 - WebSphere Application Server [810](#)
 - WebSphere Liberty [812](#)
- network-qualified LU name [133](#)
- non-MDB Java applications
 - asynchronous callout sample [781](#)
- nonconversational program
 - definition [398](#)
- nonrelated (non-terminal-related) MSDBs [297](#)
- NOSTAE and NOSPIE [53](#)
- NOT FOUND clause of WHENEVER statement [523](#)
- not-equal-to relational operator [168](#)
- NTO (Network Terminal Option) [444](#)
- null encryption [797](#)
- null value
 - column value of UPDATE statement [539](#)

O

- ODBA
 - application execution environment
 - establishing [329](#), [333](#)
 - application programs
 - testing [337](#)
 - writing [329](#)
 - CIMS [329](#)
 - Db2 for z/OS stored procedures [333](#)
 - DRA (Database Resource Adapter) [329](#)
 - RRS [329](#)
 - server program [332](#)
- ODBA (Open Database Access)
 - best practices
 - Db2 for z/OS stored procedures [334](#)
 - Db2 for z/OS stored procedures, writing [336](#)
 - Db2 for z/OS stored
 - procedures
 - best practices [334](#)
 - design best practices [334](#)
 - stopping threads [336](#)
 - writing best practices [336](#)
 - threads, stopping
 - Db2 for z/OS stored procedures [336](#)
- one-to-many mapping [27](#)
- online performance [437](#)
- online processing
 - databases that can be accessed [56](#)
 - description [58](#)
 - linking and passing control to other applications [60](#)
 - performance, maximizing [61](#)
- online programs [40](#)
- online programs, command-level samples

- online programs, command-level samples (*continued*)
 - assembler [462](#)
 - C [471](#)
 - COBOL [465](#)
 - PL/I [468](#)
- online security
 - password security [91](#)
 - supplying information about your application [91](#)
 - terminal [91](#)
- OPEN
 - statement
 - opening a cursor [546](#)
 - prepared SELECT [512](#)
 - without parameter markers [514](#)
- Open Database Access (ODBA)
 - best practices
 - Db2 for z/OS stored procedures [334](#)
 - Db2 for z/OS stored procedures, writing [336](#)
 - Db2 for z/OS stored
 - procedures
 - best practices [334](#)
 - design best practices [334](#)
 - stopping threads [336](#)
 - writing best practices [336](#)
 - threads, stopping
 - Db2 for z/OS stored procedures [336](#)
- Open Database Access (ODBA) program
 - application interface block (AIB)
 - fields [216](#)
- Open Database Access (ODBA) programs
 - abnormal termination (abend)
 - diagnosing [340](#), [341](#)
 - initialization errors [341](#)
 - running errors [341](#)
 - application interface block (AIB)AERTDLI interface [231](#)
 - tracing
 - DFSDDLTO [338](#)
 - image capture [338](#)
- operator
 - FSA [300](#)
 - SSA [168](#)
- operators
 - AND operators
 - logical [171](#)
 - Boolean [171](#)
 - Boolean operators
 - logical AND operator [171](#)
 - logical OR [171](#)
 - OR operators
 - logical [171](#)
 - relational [171](#)
 - relational operators
 - independent AND [171](#)
 - logical AND [171](#)
 - logical OR [171](#)
- options
 - CMPAT [483](#)
 - MOVENEXT [489](#)
 - P processing [498](#)
- options for subset pointers
 - MOVENEXT [489](#)

- ORDER BY clause
 - SELECT statement [543](#)
- OSAM buffer pool, retrieving statistics [139](#)
- OTMA
 - destination descriptors [761](#)
 - resume tpipe security exit routine [762](#)
- OTMA, processing conversations with [410](#)
- output message
 - format [379](#)
 - printing [382](#)
 - sending [393](#)
 - to other application programs [393](#)
 - to other IMS TM systems [397](#)
 - using Basic Edit [382](#)
 - using MFS [390](#)
 - with directed routing [397](#)
- output message counts
 - displaying [759](#)
 - overview [759](#)
 - tpipes [759](#)
- output messages, identifying destinations for [97](#)
- overlap, storage [481](#)

P

- P processing option [319](#), [498](#)
- parameter marker
 - dynamic SQL [519](#), [520](#)
 - values provided by OPEN [512](#)
 - with arbitrary statements [514](#)
- parameters
 - BKO [39](#)
 - DBCTLID [329](#)
 - ERASE [87](#)
 - JOURNAL [426](#)
 - LIST [137](#)
 - LOCKMAX [49](#)
 - MODE [45](#)
 - PROCOPT [87](#)
 - TRANSACT [45](#)
 - XTTU [426](#)
 - WFI [44](#)
- parsing error return codes [429](#)
- Partitioned Hierarchical Direct Access Method (PHDAM) [69](#), [70](#)
- Partitioned Hierarchical Indexed Direct Access Method (PHIDAM) [69](#), [72](#)
- Partitioned Secondary Index (PSINDEX) [78](#)
- partitions
 - DFSHALDB ddname [279](#)
 - HALDB control statement [279](#)
 - selective partition processing [279](#)
- Pascal
 - application programming [204](#), [357](#)
 - batch program, coding [188](#)
 - DL/I call formats [204](#), [357](#)
 - DL/I program structure [188](#)
 - entry statement [227](#)
 - PCBs, passing [227](#)
 - skeleton MPP [420](#)
 - SSA definition examples [223](#)
 - syntax diagram, DL/I call format [204](#)
- pass control of processing [11](#)
- pass control to other applications [60](#)

- passing control
 - to a conversational program [404](#)
 - to another program in a conversation [404](#)
- password phrase [870](#)
- password security [91](#)
- passwords [870](#)
- path call
 - definition [174](#)
 - example [174](#)
 - overview [174](#)
- PATIENT segment [5](#)
- PAYMENT segment [5](#)
- PCB (program communication block)
 - 12-byte time stamp, field in I/O PCB [208](#)
 - address list, accessing [219](#)
 - alternate [483](#)
 - call [59](#)
 - description [9](#)
 - express [98](#)
 - group name, field in I/O PCB [208](#)
 - GSAM (generalized sequential access method)
 - DB PCB mask, fields [283](#)
 - I/O PCB mask
 - 12-byte time stamp [363](#)
 - group name field [363](#)
 - input message sequence number [363](#)
 - logical terminal name field [363](#)
 - message output descriptor name [363](#)
 - status code field [363](#)
 - userid field [363](#)
 - in application programs, summary [483](#)
 - input message sequence number, field in I/O PCB [208](#)
 - logical terminal name, field in I/O PCB [208](#)
 - masks
 - GSAM databases [283](#)
 - I/O PCB [208](#), [363](#)
 - message output descriptor name, field in I/O PCB [208](#)
 - modifiable alternate PCBs [263](#)
 - PCB (program communication block)
 - types [483](#)
 - RACF signon security [208](#), [363](#)
 - RACROUTE SAF [208](#)
 - signon security, RACF [208](#)
 - status codes, field in I/O PCB [208](#)
 - types [373](#)
 - userid, field in I/O PCB [208](#)
- PCB lists [373](#)
- PCB parameter list in assembler language MPPs [418](#)
- PCB, express alternate [391](#)
- performance
 - impact [426](#)
 - maximizing online [61](#)
- performance measurement [797](#)
- persistent socket connections
 - configuration [791](#)
 - dedicated [784](#)
 - overview [807](#)
 - reconnect [789](#)
 - release [789](#)
 - shareable [787](#)
- PHDAM
 - multiple qualification statements [173](#)
- PHDAM (Partitioned Hierarchical Direct Access Method) [69](#), [70](#)
- PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) [69](#), [72](#)
- physical parent [274](#)
- physical structure of a database [9](#)
- PL/I
 - application programming
 - DL/I call formats [206](#), [360](#)
 - DL/I calls [206](#), [360](#)
 - batch program, coding [190](#)
 - DL/I call-level sample [192](#)
 - DL/I command-level sample [468](#)
 - DL/I program, multitasking restriction [190](#)
 - entry statement [227](#)
 - PCBs, passing [227](#)
 - pointers in entry statement [227](#)
 - return statement [227](#)
 - skeleton MPP [422](#)
 - UIB, specifying [219](#)
- PL/I language [53](#)
- PL/I segmentation APIs
 - language structures [713](#)
 - limitations [719](#)
 - restrictions [719](#)
 - sample templates [715](#)
 - top-down development [711](#)
 - trace output [717](#)
- platform configuration
 - WebSphere Application Server [728](#)
- PLITDLI procedure
 - description [239](#)
- pool manager
 - MFS [444](#)
- port numbers [870](#)
- POS (Position) call
 - description [316](#)
- POS (Position) command
 - identifying free space [498](#)
 - locating a specific sequential dependent [497](#)
 - locating the last inserted sequential dependent [497](#)
 - using with DEDBs [496](#)
- Position (POS) command
 - identifying free space [498](#)
 - locating a specific sequential dependent [497](#)
 - locating the last inserted sequential dependent [497](#)
 - using with DEDBs [496](#)
- position, reestablishing with checkpoint calls [49](#), [62](#)
- positioning
 - after DLET [243](#)
 - after ISRT [245](#)
 - after REPL [245](#)
 - after retrieval calls [243](#)
 - after unsuccessful DLET or REPL call [246](#)
 - after unsuccessful retrieval or ISRT call [246](#)
 - CHKP, effect
 - modifiable alternate PCBs [263](#)
 - current position
 - unsuccessful calls [246](#)
 - database position
 - unsuccessful calls [246](#)
 - determining [241](#)
 - not-found status code
 - description [246](#)

- positioning (*continued*)
 - not-found status code (*continued*)
 - position after [246](#)
 - positioning
 - after unsuccessful calls [246](#)
 - understanding current [241](#)
- predicate
 - general rules [541](#)
- preloaded programs [233](#)
- PREPARE statement
 - dynamic execution [520](#)
 - host variable [512](#)
- preparing programs
 - for EXEC DLI [461](#)
 - for EXEC DLI execution [475](#)
- prerequisites
 - running IMS TM resource adapter client applications [833](#)
- primarily sequential processing [74](#)
- print checkpoint log records, how to [62](#)
- problem determination [159](#)
- procedures
 - CBLTDLI [239](#)
 - PLITDLI [239](#)
- process database records [11](#)
- process of requests [11](#)
- process of requirements, analyzing [35](#), [55](#)
- processing
 - commit-point in DEDB [319](#)
 - commit-point in MSDB [302](#)
 - current position
 - multiple positioning [249](#)
 - database position
 - multiple positioning [249](#)
 - database, several views [255](#)
 - DEDBs [304](#)
 - Fast Path
 - P (position) option [498](#)
 - Fast Path databases [295](#)
 - GSAM databases [283](#)
 - multiple
 - positioning [249](#)
 - options
 - H (position), for Fast Path [319](#)
 - P (position), for Fast Path [319](#)
 - segments in logical relationships [274](#)
 - single positioning [249](#)
- processing a message [380](#)
- processing mode [45](#)
- processing models
 - dedicated persistent socket connections
 - send-only [754](#)
 - send-receive [752](#), [754](#)
 - shareable persistent socket connections
 - send-only [754](#)
 - send-receive [752](#)
- processing options
 - A (all) [87](#)
 - D (delete) [87](#)
 - defined [84](#)
 - E (exclusive) [87](#)
 - G (get)
 - description and concurrent record access [87](#)
 - general description [87](#)

- processing options (*continued*)
 - GO (read only)
 - description [87](#)
 - invalid pointers and T and N options [87](#)
 - N option [87](#)
 - risks of GOx options [87](#)
 - T option [87](#)
 - I (insert) [87](#)
 - K (key) [85](#)
 - R (replace) [87](#)
- PROCOPT parameter [87](#)
- PROCOPT=GO [49](#)
- program
 - design [175](#)
 - design efficiency [461](#)
 - restarting [263](#)
- program communication block
 - I/O PCB mask
 - userid indicator field [363](#)
 - userid indicator, field in I/O PCB [208](#)
- program communication block (PCB) [9](#)
- program deadlock [39](#)
- program restarting
 - EXEC DLI XRST command [486](#)
- program sensitivity [52](#)
- program specification block (PSB) [9](#)
- program structure
 - conversational [399](#)
 - conversational processing
 - message formats [399](#)
 - restrictions [399](#)
 - ROLB call [399](#)
 - ROLL call [399](#)
 - ROLS call [399](#)
 - steps in a conversational program [399](#)
 - deferred program switch
 - passing control to another [399](#)
 - immediate program switch [399](#)
 - LL field [399](#)
 - message
 - in conversations [399](#)
 - ROLB call
 - use in conversations [399](#)
 - ROLL call
 - use in conversations [399](#)
 - ROLS call
 - use in conversations [399](#)
 - SPA (scratchpad area)
 - contents [399](#)
 - format [399](#)
 - inserting [399](#)
 - restrictions on using [399](#)
 - saving information [399](#)
 - system service calls
 - ROLB call [399](#)
 - ROLL call [399](#)
 - ROLS call [399](#)
- program switch
 - deferred [95](#)
 - immediate [95](#)
- program test [135](#)
- program types, environments and database types [36](#)
- program waits [49](#)

- program-to-program message switching
 - conversational [404](#)
 - conversational processing
 - by deferred switch [404](#)
 - by immediate switch [404](#)
 - ending the conversation and passing control [404](#)
 - passing control and continuing the conversation [404](#)
 - restrictions [404](#)
 - deferred program switch
 - in conversational programs [404](#)
 - ending a conversation and passing control to another program [404](#)
 - immediate program switch
 - in conversational programs [404](#)
 - MSC (multiple systems coupling)
 - conversational programming [404](#)
 - nonconversational [393](#)
 - passing a conversation to another IMS TM system [404](#)
 - passing control
 - restrictions [404](#)
 - restrictions [393](#)
 - security checks [393](#)
 - SPA (scratchpad area)
 - and program-to-program switches [404](#)
- programming
 - guidelines [167](#)
 - mixed language [233](#)
 - secondary indexing [271](#)
- programming guidelines, general [461](#)
- programming models
 - asynchronous callout requests
 - overview [760](#)
 - retrieving request messages [780](#)
 - asynchronous output requests [756–758](#)
 - client-managed callout requests [776](#)
 - client-managed conversation state [821](#)
 - conversational [821](#)
 - IMS Connect-managed conversation state [821](#)
 - message-driven callout requests [764](#)
 - non-managed callout requests [776](#)
 - non-managed synchronous callout [760](#)
 - synchronous callout requests
 - correlating responses [777](#)
 - managed [760](#)
 - request message retrieval [778](#)
 - retrieval from non-MDB applications [778](#)
- programs
 - DL/I image capture [156](#)
 - DL/I test [135](#)
 - online [40](#)
 - TM batch [40](#)
- programs, BMP [483](#)
- properties
 - trckID [880](#)
- protected resources [105](#)
- protocol, locking [87](#)
- protocols
 - selecting [728](#)
 - TCP/IP connections [792](#)
- PSB (program specification block)
 - APSB (allocate program specification block) [110](#)
 - CMPAT=YES [39](#)
 - defining subset pointers [489](#)

- PSB (program specification block) (*continued*)
 - description [9](#)
 - format
 - generated program specification block (GPSB), format [230](#)
 - GPSB (generated program specification block), format [230](#)
 - PCB (program communication block) [230](#)
 - PCB, types of [483](#)
 - scheduling in a call-level program [59](#)
- pseudo-abend [52](#)
- PSINDEX (Partitioned Secondary Index) [78](#)
- PURG system service call [425](#)

Q

- Q command code [270](#)
- QC status code [44](#)
- qualification statement
 - coding [222](#)
 - field name [168](#)
 - field value
 - SSA qualification statement [168](#)
 - multiple qualification statements
 - DEDB [173](#)
 - HDAM [173](#)
 - PHDAM [173](#)
 - qualification statement
 - field value [168](#)
 - randomizing routine
 - exit routine [173](#)
 - relational operator [168](#)
 - segment name [168](#)
 - structure [168](#)
- qualified calls
 - overview [168](#)
- qualified SSA
 - structure with command code [174](#)
- qualified SSAs (segment search arguments)
 - qualification statement [168](#)
- qualifying
 - DL/I calls with command codes [174](#)
 - SSAs [168](#)
- quantitative relationship between data aggregates [27](#)
- queries
 - in application programs [509](#)

R

- read access, specify with PROCOPT operand [87](#)
- read without integrity [89](#)
- read-only access, specify with PROCOPT operand [87](#)
- reading segments in MSDBs [297](#), [299](#)
- receiving messages
 - other IMS TM systems [396](#)
- record
 - database processing [11](#)
 - database, description of [5](#)
- record descriptor word (RDW)
 - IMS Spool API [426](#)
- recording
 - data availability [20](#)
 - information about your program [163](#)

- recoverable in-doubt structure. [260](#)
- recoverable resources [105](#)
- recovering databases [485](#)
- recovery
 - considerations in conversations [96](#)
 - identifying requirements [49](#)
 - in a batch-oriented BMP [42](#)
 - in batch programs [39](#)
 - RIS [260](#)
 - token
 - definition [260](#)
- recovery EXEC DLI commands
 - XRST [486](#)
- recovery of databases [64](#)
- Recovery process
 - distributed [105](#)
 - local [105](#)
- Recovery, Resource [105](#)
- redundant data [3](#)
- reestablish position in database [49](#)
- related (terminal related) MSDBs [297](#)
- relational database
 - hierarchical database, compared to [555](#)
- relational databases [36](#)
- relational operators
 - Boolean operators [171](#)
 - list [168](#)
 - overview [168](#)
 - SSA qualification statement [168](#)
 - SSA, coding [222](#)
- relationships
 - between data elements [22](#)
 - data, hierarchical [5](#)
 - defining logical [81](#)
 - mapping data [27](#)
- relationships between data aggregates [27](#)
- remote DL/I [56](#)
- removeInvalidCaseFields [610](#)
- repetitive data elements, isolating [22](#)
- REPL (Replace) call
 - with MSDB, DEDB or VSO DEDB [298](#)
- reply to the terminal in a conversation [95](#)
- replying to the terminal in a conversation [403](#)
- report of instructor schedules [28](#)
- reports, creating [20](#)
- requests, processing [11](#)
- required application data, analyzing [17](#)
- requirements
 - composite business applications [727](#)
 - global transactions [825](#)
 - IMS TM resource adapter [727](#)
- requirements, analyzing processing [35](#)
- reroute names [878](#)
- rerouting undelivered output [878](#)
- reserving
 - place for command codes [320](#)
 - segment
 - command code [270](#)
 - lock management [270](#)
- residency mode (RMODE) [233](#), [375](#)
- resolving data structure conflicts [76](#)
- resource managers [105](#)
- Resource Recovery
 - application program [105](#)

- Resource Recovery (*continued*)
 - Introduction to [105](#)
 - protected resources [105](#)
 - recoverable resources [105](#)
 - resource managers [105](#)
 - sync-point manager [105](#)
- Resource Recovery Services. [258](#)
- Resource Recovery Services/Multiple Virtual Storage (RRS)
 - introduction to [105](#)
- resource workload routing [749](#)
- resources
 - protected [105](#)
 - recoverable [105](#)
 - security [15](#)
- response mode, description [98](#)
- response time
 - measurement [797](#)
- restart your program
 - code for, description [65](#)
 - with basic CHKP [49](#)
 - with symbolic CHKP [49](#)
- Restart, Extended [48](#), [65](#)
- restarting your program, basic checkpoints [263](#)
- restrictions
 - CHKP and XRST with GSAM [289](#)
 - composite business applications [727](#)
 - database calls
 - to DEDBs [320](#)
 - to MSDBs [297](#)
 - IMS TM resource adapter [727](#)
 - XRST (Extended Restart) call with GSAM [289](#)
- result segment
 - description [543](#)
 - example [543](#)
 - of SELECT statement [543](#)
- result segments
 - formatting [543](#)
- retrieval of IMS database statistics [139](#)
- retrieving
 - data using SELECT *
 - variable-length segments [544](#)
 - dependent segments [80](#)
- RETRY option [53](#)
- return codes
 - UIB [219](#)
- RIS (recoverable in-doubt structure) [260](#)
- risks to security, combined files [3](#)
- RMODE [233](#)
- ROLB
 - in MPPs and transaction-oriented BMPs [264](#)
- ROLB (Roll Back) call
 - compared to ROLL call [264](#), [265](#)
 - description [264](#)
 - maintaining database integrity [264](#)
 - usage [264](#)
- ROLB call
 - description [412](#)
- ROLB command
 - compared to ROLL and ROLS [411](#)
- ROLB system service call [39](#), [64](#)
- ROLB, ROLL, ROLS [410](#)
- ROLL (Roll) call
 - compared to ROLB call [264](#), [265](#)
 - description [264](#)

- ROLL (Roll) call (*continued*)
 - maintaining database integrity [264](#)
- roll back point [410](#)
- ROLL call
 - description [412](#)
- ROLL command
 - compared to ROLB and ROLS [411](#)
- ROLL system service call [64](#)
- ROLS
 - backing out to an intermediate backout point [267](#)
- ROLS (Roll Back to SETS) call
 - maintaining database integrity [264](#)
 - TOKEN [264](#), [266](#)
- ROLS (Rollback to SETS or SETU) command
 - backout point, intermediate [486](#)
- ROLS call
 - with LU 6.2 [413](#)
 - with TOKEN [413](#)
 - without TOKEN [413](#)
- ROLS command
 - compared to ROLB and ROLL [411](#)
- ROLS system service call [39](#), [52](#), [67](#)
- root anchor point [70](#)
- root segment, definition [5](#)
- roster, current [17](#)
- routines
 - ESTAE [53](#)
 - STAE [53](#)
- row
 - selecting with WHERE clause [541](#)
 - updating [539](#)
- rows
 - relational representation, in [555](#)
 - segment instances, compared to [555](#)
- RRS (z/OS Resource Recovery Services)
 - summary of IMS support [108](#)
- RSA (record search argument)
 - GSAM, reference [226](#)
- rules
 - coding an SSA [222](#)
- runtime component
 - file contents [734](#)

S

- SAA resource recovery interface calls [257](#)
- SAMETRM=YES [403](#)
- sample programs
 - call-level assembler language
 - CICS online [178](#)
 - call-level COBOL, CICS online [184](#)
 - call-level PL/I, CICS online [192](#)
- sample programs, command level
 - assembler language [462](#)
 - C [471](#)
 - COBOL [465](#)
 - PL/I [468](#)
- samples
 - Common Client Interface (CCI) [828](#), [830](#)
 - IMS TM resource adapter [831](#)
 - JNDI lookups [830](#)
- schedule a PSB, in a call-level program, how to [59](#)
- schedule, classes example [28](#)

- screen design considerations [94](#)
- SCS1 devices
 - meaning of designation [444](#)
- SCS2 devices
 - meaning of designation [444](#)
- SDSF (Spool Display and Search Facility) [427](#)
- search condition
 - comparison operators [541](#)
 - WHERE clause [541](#)
- secondary index
 - description [308](#)
 - preparing to use [308](#)
 - using [308](#)
- secondary indexes
 - multiple qualification statements [272](#)
- secondary indexing
 - DB PCB contents [274](#)
 - description [77](#)
 - effect on programming [271](#)
 - examples of uses [78](#)
 - information returned by DL/I [274](#)
 - Partitioned Secondary Index (PSINDEX) [78](#)
 - specifying [78](#)
 - SSAs [271](#)
 - status codes [274](#)
- secondary logical unit [444](#)
- secondary processing sequence [271](#)
- Secure Sockets Layer (SSL)
 - client authentication [805](#)
 - configuration [807](#), [808](#)
 - enablement [870](#)
 - encryption types [870](#)
 - error messages
 - HWSP1445E [867](#)
 - HWSSSL00E [868](#)
 - keystore names [871](#)
 - keystore passwords [871](#)
 - null encryption [805](#)
 - overview [804](#)
 - process [805](#), [806](#)
 - server authentication [805](#)
 - SSL handshake [805](#)
 - supported certificates [804](#)
 - supported keystore type [804](#)
 - supported SSL protocol [804](#)
 - truststore names [872](#)
 - truststore passwords [872](#)
- security
 - and the PROCOPT= operand [87](#)
 - changing RACF password [808](#)
 - client authentication [805](#)
 - Common Client Interface (CCI) [802](#)
 - component-managed signon [797](#)
 - container-managed [799](#)
 - container-managed signon [797](#)
 - database [84](#), [85](#)
 - distributed security credential propagation [810](#), [812](#), [814](#)
 - EIS signon [797](#)
 - field level sensitivity [85](#)
 - identifying online requirements [91](#)
 - IMS TM resource adapter [810](#), [812](#), [814](#)
 - JAAS [797](#)
 - Java 2 Security Manager [797](#)

- security (*continued*)
 - key sensitivity [85](#)
 - Local Option connections [797](#)
 - network security credential propagation [810](#), [812](#), [814](#)
 - null encryption [805](#)
 - of databases and data communications [15](#)
 - of resources [15](#)
 - password security [91](#)
 - process [805](#)
 - RACF [797](#)
 - retrieving asynchronous output messages [809](#)
 - retrieving callout messages [809](#)
 - risks of combined files [3](#)
 - SAF [797](#)
 - Secure Sockets Layer (SSL) communications [797](#)
 - segment sensitivity [85](#)
 - server authentication [805](#)
 - signon [91](#)
 - SSL handshake [805](#)
 - SSL overview [804](#)
 - supplying information about your application [91](#)
 - supported certificates [804](#)
 - supported keystore type [804](#)
 - supported SSL protocol [804](#)
 - terminal [91](#)
 - TLS support [804](#)
- Security Authorization Facility (SAF) [797](#)
- security checks in program-to-program switching [393](#)
- segment
 - description [5](#)
 - preventing access to by other programs [61](#)
 - retrieving [545](#)
 - sensitivity [85](#)
 - sequential dependent
 - identifying free space [498](#)
 - locating a specific dependent [497](#)
 - locating the last inserted dependent [497](#)
- segment name
 - SSA qualification statement [168](#)
- segment search argument (SSA)
 - coding rules [222](#)
- segment search arguments (SSAs) [168](#)
- segment, information needed [175](#)
- segments
 - in medical database example [5](#)
 - in SQL queries [555](#)
 - medical database example [5](#)
 - tables, compared to [555](#)
- Segments
 - Message Input Format [378](#)
- SELECT keyword
 - example query [555](#)
- SELECT statement
 - AS clause
 - with ORDER BY clause [543](#)
 - clauses
 - FROM [541](#)
 - ORDER BY [543](#)
 - WHERE [541](#)
 - dynamic execution [519](#)
 - fixed-list [512](#)
 - named columns [541](#)
 - ORDER BY clause
 - derived columns [543](#)
- SELECT statement (*continued*)
 - ORDER BY clause (*continued*)
 - with AS clause [543](#)
 - parameter markers [514](#)
 - selecting a set of rows [545](#)
 - using with
 - * (to select all columns) [541](#)
 - column-name list [541](#)
 - DECLARE CURSOR statement [546](#)
 - varying-list [514](#)
 - selecting
 - all columns [541](#)
 - named columns [541](#)
 - rows [541](#)
 - some columns [541](#)
 - selective partition processing
 - DFSHALDB ddname [279](#)
 - HALDB control statement [279](#)
 - sending messages
 - defining alternate PCBs for [391](#)
 - other IMS TM systems [397](#)
 - overview [377](#)
 - to other application programs [393](#)
 - to other IMS TM systems [395](#)
 - using ISRT [391](#)
 - sensitivity
 - data [9](#)
 - field level [9](#), [85](#)
 - general description [84](#)
 - key [85](#)
 - program [52](#)
 - segment [85](#)
 - sequence field
 - virtual logical child, in [168](#)
 - sequence numbers
 - COBOL application program [533](#)
 - sequential access methods
 - characteristics of [74](#)
 - HISAM [74](#)
 - HSAM [74](#)
 - types of [74](#)
 - sequential dependent segments
 - how stored [298](#)
 - identifying free space [498](#)
 - locating a specific dependent [497](#)
 - locating the last inserted dependent [497](#)
 - POS (Position) command [496](#)
 - sequential dependents
 - overview [298](#)
 - sequential processing only [74](#)
 - SET clause of UPDATE statement [539](#)
 - SETO call
 - usage [429](#)
 - SETO system service call [425](#)
 - SETS
 - backing out to an intermediate backout point [267](#)
 - SETS call
 - description [414](#)
 - SETS system service call [39](#), [52](#), [67](#)
 - SETU
 - backing out to an intermediate backout point [267](#)
 - SETU system service call [67](#)
 - shareable persistent socket connections
 - establishing [792](#)

- shareable persistent socket connections (*continued*)
 - processing models [752](#)
 - programming models [756–758](#)
 - TCP/IP connections [792](#)
- shared queues option [91](#)
- SHISAM (Simple Hierarchical Indexed Sequential Access Method) [75](#)
- SHSAM (Simple Hierarchical Sequential Access Method) [75](#)
- signon security [91](#)
- simple HISAM (SHISAM) [75](#)
- simple HSAM (SHSAM) [75](#)
- single mode [41](#), [45](#), [49](#)
- skeleton programs
 - assembler language [177](#), [418](#)
 - C language [180](#), [418](#)
 - COBOL [182](#), [419](#)
 - Pascal [188](#), [420](#)
 - PL/I [190](#), [422](#)
- skills report, instructor [28](#)
- SLU
 - type 1
 - defining to operate with MFS [444](#)
 - type 2
 - defining to operate with MFS [444](#)
 - type 6.1
 - defining to operate with MFS [444](#)
 - type P
 - defining to operate with MFS [444](#)
- socket connections
 - dedicated persistent [784](#)
 - reconnect [789](#)
 - shareable persistent [787](#)
 - types [783](#)
- socket timeouts
 - definition [817](#)
 - exceptions [817](#)
 - interaction properties [879](#)
 - setting in a development environment [819](#)
 - setting in the CCI [819](#)
 - specifying values [818](#)
- sort key
 - ORDER BY clause [543](#)
 - ordering [543](#)
- SPA (scratchpad area) [96](#)
- specification of
 - field level sensitivity [76](#)
 - frequency, checkpoint [51](#)
- specifying
 - DB PCB mask [212](#)
 - GSAM data set attributes [293](#)
 - processing options for DEDBs [319](#)
- SPIE routine [53](#)
- Spool API
 - CHNG call, keywords [429](#)
 - code examples
 - Application PCB structure [434](#)
 - CHNG call to alternate PCB [434](#)
 - GU call to I/O PCB [434](#)
 - ISRT call to alternate PCB [434](#)
 - error codes
 - description [429](#)
 - diagnosis, examples [430](#)
 - parsing errors
 - diagnosis, examples [430](#)
- Spool API (*continued*)
 - parsing errors (*continued*)
 - error codes [429](#)
 - status codes [429](#)
 - print data set characteristics [429](#)
 - SETO call, keywords [429](#)
 - status codes [429](#)
- Spool Display and Search Facility (SDSF) [427](#)
- SQL
 - application programming [505](#)
 - writing application programs [509](#)
- SQL (Structured Query Language)
 - checking execution [521](#)
 - coding [511](#)
 - cursors [545](#)
 - dynamic
 - coding [511](#)
 - example query [555](#)
 - return codes
 - checking [522](#)
 - varying-list [514](#)
- SQL communication area (SQLIMSCA)
 - description [522](#)
- SQL statements
 - checking for successful execution [509](#)
 - CLOSE [512](#), [548](#)
 - COBOL program sections [533](#)
 - comments
 - COBOL [533](#)
 - continuation
 - COBOL [533](#)
 - DECLARE CURSOR
 - description [546](#)
 - example [512](#), [514](#)
 - DELETE
 - description [547](#)
 - example [540](#)
 - DESCRIBE [514](#)
 - EXECUTE [520](#)
 - FETCH
 - description [547](#)
 - example [512](#)
 - in application programs [509](#)
 - INSERT [538](#)
 - labels
 - COBOL [533](#)
 - margins
 - COBOL [533](#)
 - OPEN
 - description [546](#)
 - example [512](#)
 - PREPARE [520](#)
 - SELECT
 - description [541](#)
 - UPDATE
 - description [547](#)
 - example [539](#)
 - WHENEVER [523](#)
- SQLERROR clause of WHENEVER statement [523](#)
- SQLIMSCA (SQL communication area)
 - checking SQLIMSCODE [522](#)
 - checking SQLIMSERRD(3) [522](#)
 - checking SQLIMSTATE [522](#)
 - deciding whether to include [509](#)

- SQLIMSCA (SQL communication area) *(continued)*
 - description [522](#)
- SQLIMSCA (SQL communications area)
 - COBOL [524](#)
- SQLIMSCODE
 - +100 [523](#)
 - values [522](#)
- SQLIMSCODE host variable
 - deciding whether to declare [509](#)
- SQLIMSDA (SQL descriptor area)
 - allocating storage [514](#)
 - COBOL [524](#)
 - dynamic SELECT example [514](#)
 - no occurrences of SQLIMSVAR [514](#)
 - OPEN statement [512](#)
 - parameter markers [514](#)
 - requires storage addresses [514](#)
 - varying-list SELECT statement [514](#)
- SQLIMSN field of SQLIMSDA [514](#)
- SQLIMSSTATE
 - values [522](#)
- SQLIMSSTATE host variable
 - deciding whether to declare [509](#)
- SQLIMSVAR field of SQLIMSDA [514](#)
- SQLWARNING clause of WHENEVER statement [523](#)
- SSA (segment search argument)
 - coding
 - formats [223](#)
 - restrictions [222](#)
 - rules [222](#)
 - coding rules [222](#)
 - command codes [174](#)
 - qualification statement [222](#)
 - reference [222](#)
 - relational operators [168](#)
 - restrictions [222](#)
 - segment name field [222](#)
 - structure with command code [174](#)
 - usage
 - command codes [174](#)
 - guidelines [170](#)
 - multiple qualification statements [171](#)
 - virtual logical child [168](#)
- ssaOptimization [618](#)
- SSAs (segment search argument)
 - overview [168](#)
 - segment name field [168](#)
- SSAs (segment search arguments)
 - definition [168](#)
 - unqualified [168](#)
 - usage
 - secondary indexing [271](#)
- SSL (Secure Sockets Layer)
 - client authentication [805](#)
 - configuration [807](#), [808](#)
 - enablement [870](#)
 - encryption types [870](#)
 - error messages
 - HWSP1445E [867](#)
 - HWSSSL00E [868](#)
 - keystore names [871](#)
 - keystore passwords [871](#)
 - null encryption [805](#)
 - overview [804](#)

- SSL (Secure Sockets Layer) *(continued)*
 - process [805](#), [806](#)
 - protocol [872](#)
 - server authentication [805](#)
 - SSL handshake [805](#)
 - supported certificates [804](#)
 - supported keystore type [804](#)
 - supported SSL protocol [804](#)
 - truststore names [872](#)
 - truststore passwords [872](#)
- SSL concepts [805](#)
- SSLProtocol
 - Secure Sockets Layer (SSL)
 - protocol [872](#)
- STAE routines [53](#)
- standard application programs and MSC [408](#)
- STAT call
 - formats for statistics
 - OSAM buffer pool, STAT call [139](#)
 - OSAM buffer subpool, enhanced STAT call [143](#)
 - VSAM buffer subpool, enhanced STAT call [148](#)
 - VSAM buffer subpool, STAT call [141](#)
 - system service [340](#)
 - use in debugging [139](#), [159](#)
- static SQL
 - description [511](#)
- statistics, database [139](#)
- status code, QC [44](#)
- status codes
 - AJ [496](#)
 - AM [496](#)
 - FSA [300](#)
 - GSAM [288](#)
 - H processing option [319](#)
 - logical relationships [277](#)
 - P processing option [319](#), [498](#)
 - subset pointers [308](#), [496](#)
- STATUS statement [135](#)
- storage
 - acquiring
 - retrieved row [514](#)
 - SQLIMSDA [514](#)
 - addresses in SQLIMSDA [514](#)
- storage of data
 - in a combined file [3](#)
 - in a database [3](#)
 - in separate files [3](#)
- storage overlap [481](#)
- structure
 - data [9](#)
 - physical, of a database [9](#)
- structure of data, methods [22](#)
- Structured Query Language (SQL) [36](#)
- subset pointer command codes
 - restrictions [174](#)
- subset pointers
 - command codes
 - subset pointers [304](#)
- DEDB
 - managed by command codes [174](#)
- defining DBD [489](#)
- defining PSB [489](#)
- defining, DBD [304](#)
- defining, PCB [304](#)

- subset pointers (*continued*)
 - description [304, 487](#)
 - MOVENEXT option [489](#)
 - moving forward [489](#)
 - preparation for using [489](#)
 - preparing to use [304](#)
 - specifying
 - command codes for DEDBs [304](#)
 - status codes [308, 496](#)
 - using [304](#)
- summary of command codes [174](#)
- summary of symbolic CHKP and basic CHKP [48](#)
- supply security information, how to [91](#)
- supported features
 - IMS TM resource adapter all versions [724](#)
- supported platforms [726](#)
- supported software configurations [726](#)
- supported versions [726](#)
- symbolic checkpoint
 - description [48, 62](#)
 - IDs, specifying [62](#)
 - issuing [65](#)
 - restart [65](#)
 - restart with [49](#)
- Symbolic Checkpoint (SYMCHKP) command
 - restart [486](#)
 - XRST [486](#)
- SYMCHKP (Symbolic Checkpoint) command
 - restart [486](#)
 - XRST [486](#)
- sync levels
 - interaction property [879](#)
 - processing [782](#)
 - supported interactions [782](#)
- sync point
 - application program [257](#)
 - CPI Communications driven programs [257](#)
 - data propagation [261](#)
 - log records [261](#)
 - relationship to commit point and check point [257](#)
- sync_level values [104](#)
- sync-point manager (SPM) [105](#)
- synchronization point [410](#)
- synchronization point manager [104](#)
- synchronous callout
 - definition [760](#)
 - message flow [762](#)
 - non-MDB sample [779](#)
 - requirements [727](#)
 - status codes [879](#)
- synchronous callout requests
 - non-message-driven bean (MDB) applications [778](#)
- synchronous conversation, description for LU 6.2
- transactions [102](#)
- synchronous program switch request
 - JMS (Java Message Service) implementation
 - IMSQueueConnectionFactory [697](#)
- SYNCLVL [257](#)
- synonym, data element [19](#)
- syntax diagram
 - how to read [xvii](#)
- sysplex data-sharing [42](#)
- sysplex distributor [789](#)
- system log

- system log (*continued*)

- on tape [39](#)
 - storage [39](#)
- system service calls
 - CHNG [425](#)
 - INIT [52](#)
 - INQY [52](#)
 - ISRT [425](#)
 - LOG [151, 340](#)
 - PURG [425](#)
 - ROLB [39, 64](#)
 - ROLB call [412](#)
 - ROLL [64](#)
 - ROLL call [412](#)
 - ROLS [39, 52, 67](#)
 - SETO [425](#)
 - SETS [39, 52, 67](#)
 - SETU [67](#)
 - STAT [139, 340](#)

T

- table
 - deleting rows [540](#)
 - inserting single row [538](#)
 - updating rows [539](#)
- tables
 - relational representation, in [555](#)
 - segments, compared to [555](#)
- take checkpoints, how to [62](#)
- TCP/IP
 - global transaction support [826](#)
 - supported communication protocol [728](#)
- terminal screen, designing [94](#)
- terminal security [91](#)
- termination of a PSB, restrictions [59](#)
- termination, abnormal [45](#)
- test of application programs
 - using BTS [136](#)
 - using DFSDDLTO [156](#)
 - using DL/I test program [135](#)
 - what you need [135, 155](#)
- test of DL/I call sequences [135, 156](#)
- test, unit [135](#)
- testing
 - CICS programs
 - tools [155](#)
- thread pool
 - configuration [772](#)
- timeout
 - activating [395](#)
- timeouts
 - browsers [819](#)
 - Enterprise JavaBeans (EJB) sessions [819](#)
 - Enterprise JavaBeans (EJB) transactions [819](#)
 - execution [814](#)
 - HTTP sessions [819](#)
 - J2C connection factory [819](#)
 - socket [814](#)
- TM batch program [40](#)
- token, definition of [96](#)
- tpipes
 - dequeue count [759](#)
 - enqueue count [759](#)

- tracing
 - for IMS TM resource adapter [838, 840](#)
 - in WebSphere Application Server [838](#)
 - in WebSphere Liberty servers [839](#)
- trademarks [883, 884](#)
- TRANSACT macro [45](#)
- transaction code [40](#)
- transaction expiration
 - interaction property [880](#)
- transaction pipes [751](#)
- transaction response mode [41](#)
- Transaction tracking ID [880](#)
- transaction-oriented BMPs
 - ROLB [264](#)
- transaction-oriented BMPs. [49](#)
- transactions
 - conversational [820, 821, 827](#)
 - expiration [815](#)
 - global transaction processing [824](#)
 - global transaction support process [825](#)
 - global transaction with TCP/IP [826](#)
 - local transaction processing [827](#)
 - non-global transaction processing [827](#)
 - one-phase commit processing [827](#)
 - two-phase commit support process [825](#)
- translator
 - options required for EXEC DLI [476](#)
- Transport Layer Security (TLS) [804](#)
- TREATMNT segment [5](#)
- troubleshooting
 - IMS TM resource adapter
 - callout requests [837](#)
 - installation verification program [835](#)
 - Java applications access to IMS [836](#)
- truststores
 - configuring for IMS TM resource adapter [807](#)
 - connection property [872](#)
 - creating [807](#)
 - password [872](#)
- TSO application programs [45](#)
- tutorials
 - IMS TM resource adapter [831](#)
- two-phase [258](#)
- two-phase commit
 - overview [258](#)
 - single-phase [260](#)
 - UOR [260](#)
- two-phase commit process
 - UOR [105](#)
- two-phase commit processing
 - global transaction support [824, 825](#)
 - recommendations [827](#)
- two-phase commit protocol [105](#)
- TXTU parameter [426](#)
- type 18 log record [62](#)

U

- UIB (user interface block)
 - defining, in program [219](#)
 - field names [219](#)
 - PCB address list, accessing [219](#)
 - return codes, accessing [219](#)
- unavailability of data [52, 65](#)

- Unicode
 - data, retrieving from IMS [514](#)
- unique identifier, data [19](#)
- unit of recovery (UOR)
 - definition [260](#)
- unit of work [45](#)
- unit of work (UOW)
 - crossing a boundary when processing DEDBs [498](#)
- unit test [135](#)
- Universal JDBC driver
 - DL/I calls [628](#)
- unqualified calls
 - command codes
 - C [168](#)
 - SSAs (segment search arguments) [168](#)
 - definition [168](#)
 - DL/I calls (general information)
 - types [168](#)
 - overview [168](#)
 - qualified calls
 - definition [168](#)
 - SSAs (segment search arguments)
 - qualified [168](#)
 - unqualified [168](#)
- unqualified SSA
 - structure with command code [174](#)
 - usage with command codes [174](#)
- unqualified SSAs
 - segment name field [168](#)
- UOR (unit of recovery)
 - definition [260](#)
 - in-doubt
 - definition [260](#)
 - in-flight
 - definition [260](#)
- UOW (unit of work)
 - crossing a boundary when processing DEDBs [498](#)
- UOW boundary, processing DEDB [319](#)
- updatable cursor [546](#)
- update access, specify with PROCOPT operand [87](#)
- UPDATE statement
 - description [539](#)
 - SET clause [539](#)
- updating
 - segments in an MSDB, DEDB or VSO DEDB [298](#)
- uppercase, using Basic Edit [382](#)
- user interface block [219](#)
- user names [872](#)
- user requirements, analyzing [15](#)
- USER special register
 - value in UPDATE statement [539](#)
- Using
 - MFS [437](#)
- USING DESCRIPTOR clause
 - EXECUTE statement [514](#)
 - FETCH statement [514](#)
 - OPEN statement [514](#)
- utilities
 - Batch Backout [39](#)
 - DFSERA10 [62, 339](#)
 - File Select and Formatting Print program [48](#)

V

- VALUES clause, INSERT statement [538](#)
- values, isolating duplicate [22](#)
- variable
 - COBOL [525](#)
- variable length messages [793](#)
- variable-length database segments
 - IMS Universal drivers [565](#)
 - SQL support for [565](#)
- VBASF, formatted VSAM subpool statistics [141](#)
- VBASS, formatted summary of VSAM subpool statistics [141](#)
- VBASU, unformatted VSAM subpool statistics [141](#)
- VBESF, formatted VSAM subpool statistics [148](#)
- VBESS, formatted summary of VSAM subpool statistics [148](#)
- VBESU, unformatted VSAM subpool statistics [148](#)
- versioning
 - databases
 - application programming [237](#)
- view of data, a program's [9](#)
- view, local [28](#)
- virtual logical child [168](#)
- VisualGen [20](#)
- VSAM buffer subpool, retrieving
 - enhanced subpool statistics [148](#)
 - statistics [141](#), [148](#)
- VTAM I/O facility
 - effects on VTAM terminals [395](#)
- VTAM terminal
 - activating a timeout [395](#)

W

- wait-for-input (WFI)
 - transactions [41](#), [44](#)
- waits, program [49](#)
- WebSphere Application Server
 - aged timeout property [789](#)
 - distributed security credential propagation [810](#)
 - EAR file installation [833](#)
 - IMS TM resource adapter [810](#)
 - IMS TM resource adapter archive (RAR file) deployment [735](#)
 - network security credential propagation [810](#)
- WebSphere Liberty
 - distributed security credential propagation [812](#)
 - IMS TM resource adapter [812](#)
 - network security credential propagation [812](#)
- WebSphere Liberty servers
 - connection factory configuration [738](#)
 - IMS TM resource adapter archive (RAR file) deployment [737](#)
- WFI parameter [44](#)
- WHENEVER statement
 - COBOL [533](#)
 - CONTINUE clause [523](#)
 - GO TO clause [523](#)
 - NOT FOUND clause [523](#), [547](#)
 - specifying [523](#)
 - SQL error codes [523](#)
 - SQLERROR clause [523](#)
 - SQLWARNING clause [523](#)
- WHERE clause
 - SELECT statement

- WHERE clause (*continued*)
 - SELECT statement (*continued*)
 - description [541](#)
- writing
 - DDL [618](#)
- writing application programs
 - SQL [509](#)
- writing information to the system log [151](#)

X

- X'18' log record [48](#)
- XRST (Extended Restart) [48](#)

Z

- z/OS
 - extended addressing capabilities
 - addressing mode (AMODE) [375](#)
 - DCCTL environment [375](#)
 - preloaded program [375](#)
 - residency mode (RMODE) [375](#)
- z/OS files
 - access to [36](#), [56](#)
 - description [56](#)
- z/OS Resource Recovery Services [110](#), [257](#)
- z/OS Resource Recovery Services (RRS)
 - ODBA interface [329](#)
 - summary of IMS support [108](#)
- z/OS Scheduler JCL Facility (SJF) [425](#)
- Z1 field [379](#)
- Z2 field [379](#)
- ZZ field
 - in input message [378](#)
 - in output message [379](#)



Product Number: 5635-A06
5655-DS5
5655-TM4