

IMS
Version 12

Operations and Automation



IMS
Version 12

Operations and Automation



Note

Before using this information and the product that it supports, be sure to read the general information under “Notices” on page 371.

This edition applies to IMS Version 12 (program number 5635-A03), IMS Database Value Unit Edition, V12.1 (program number 5655-DSQ), and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1974, 2013.**

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

Contents

About this information vii

Prerequisite knowledge	vii
IMS function names used in this information	vii
How new and changed information is identified	vii
How to read syntax diagrams	viii
Accessibility features for IMS Version 12	ix
How to send your comments	x

Chapter 1. Controlling IMS. 1

Issuing commands with the IMS Control Center	1
Controlling IMS with the TSO SPOC application	1
Starting and setting up the TSO SPOC.	2
Type-1 and type-2 command responses from the TSO SPOC	5
Displaying command status in the TSO SPOC	5
Command shortcuts in the TSO SPOC.	6
Defining groups of IMS systems in the TSO SPOC	7
Reissuing commands in the TSO SPOC	7
Managing IMS resources by using TSO SPOC	8
Viewing the OM audit trail log by using TSO SPOC.	9
Issuing Batch SPOC commands.	10
Modifying and controlling system resources	12
Modifying system resources online	12
List of commands with similar functions for multiple resources	15
Modifying dependent regions	24
Modifying telecommunication lines	24
How to modify terminals.	24
Modifying and controlling transactions	24
Database control.	25
Creating, updating, deleting, and querying resource definitions dynamically	25
Modifying ETO user IDs and assignments of ISC users	35
Modifying Multiple Systems Coupling resources	35
Modifying security options	35
Displaying and terminating conversations	36
Modifying and controlling subsystems	36
Controlling OTMA input messages	36
Recovery during the IMSRSC repository data set update process	36
Controlling log data set characteristics	40
Changing online log data set characteristics	41
Changing write-ahead data set characteristics	43
Changing system log data set characteristics	45
Changing RECON data set characteristics	45
Connecting and disconnecting subsystems	46
Commands for IMS operations tasks	47

Chapter 2. Starting or restarting IMS . . . 61

Starting an IMSplex	61
Starting the CSL	62
Starting the CSL SCI	63
Restarting the CSL SCI	63

Starting the CSL ODBM	63
Restarting the CSL ODBM	64
Starting the CSL OM	64
Restarting the CSL OM	64
Starting the CSL RM	64
Restarting the CSL RM	65
Starting and stopping the IMSRSC repository	65
Opening the IMSRSC repository	66
Starting the IMSRSC repository.	67
Stopping the IMSRSC repository	67
Starting and stopping the Repository Server	68
Starting the Repository Server	68
Stopping the Repository Server.	69
How to start the IMS control region	70
Setting the z/OS TOD clock	71
Setting or changing local time	73
Cold starting IMS in a non-shared queues environment	73
Cold starting IMS in a shared-queues environment	73
Cold starting IMS with dynamic resource definition	74
Warm starting IMS (or normal restart)	77
Emergency starting IMS	77
Emergency restarts and dynamic resource definition	78
SLDS input to warm start and emergency restart	79
Specifying security options at startup.	79
Starting the IRLM	80
Starting the CQS.	80
Starting dependent regions	81
Message processing regions	82
Batch message processing regions	82
Fast Path message-driven regions	82
Java dependent regions	82
DEDB online utility regions	82
CCTL regions.	82
Restarting CQS	83
CQS warm start	83
CQS cold start	84
CQS registration with the z/OS Automatic Restart Manager.	84
Restarting CQS after CQS resource cleanup failures	85
Restarting CQS queue structures	85
CQS structure allocation	85
CQS structure warm start.	85
CQS structure cold start	86
Restarting resource structures	87
CQS structure recovery for restarting.	87
Starting Transaction Manager	88
Connecting to the VTAM network.	88
Connecting to devices	89
Connecting to APPC/MVS	90
Connecting ISC sessions from CICS to IMS.	90
Starting IMS systems and global command status	91

Restarting a component of IMS	93
Restarting IMS	94
Cold start	95
Warm start	96
Emergency restart	97
Defer of BMP backout programs to speed emergency restart	99
IMS restart and global resource status in an IMSplex that uses RM	99
How to perform a subsystem restart with the z/OS Automatic Restart Manager	100
BMP restart after system failure	101
Restarting batch jobs	102
Reconnecting CCTLs or ODBA application programs	102

Chapter 3. Monitoring IMS 103

Monitoring the system	103
Monitoring IMS Connect connections	103
Checking port TCP/IP connection status	104
Checking client TCP/IP connection status	104
Checking remote IMS Connect TCP/IP connection status	105
Checking IMSplex member SCI connection status	105
Checking XCF data store connection status for IMS TM clients	106
IMS system log utilities	108
File Select and Formatting Print utility (DFSERA10)	108
Log Transaction Analysis utility (DFSILTA0)	108
Statistical Analysis utility (DFSISTS0)	109
Gathering performance-related data	109
Activating and controlling the DB Monitor	110
Logging the data from the DB Monitor	110
Data recording and the IMS Monitor	111

Chapter 4. Shutting down IMS 113

Stopping Transaction Manager	113
Stopping APPC	114
Stopping OTMA	114
Stopping dependent regions	114
Shutting down the IMS control region	114
Shutting down IMS using the /CHECKPOINT commands	116
Shutting down an IMS system that uses dynamic resource definition	118
Session termination	119
Shutting down an IMS network	120
Terminating an ISC session from CICS	120
Stopping the IRLM	121
Shutting down CQS	122
Shutting down an IMSplex	122
Shutting down the CSL	123
Shutting down the CSL using z/OS commands	124
Shutting down the CSL ODBM	124
Shutting down the CSL OM	125
Shutting down the CSL RM	125
Shutting down the CSL SCI	126
Forced termination of IMS	126

Offline dump formatter	127
Producing a dump using the z/OS MODIFY command	127
Producing a dump using the z/OS DUMP command	127
Producing a dump using the standalone dump (SADMP)	127
Keeping dump data sets available	128
Database resource adapter storage	129

Chapter 5. IMS failure recovery 131

z/OS system failures	131
Control region failures	132
Emergency restart failures	132
Re-establishing database integrity	132
System data set failures	133
Message queue data set failures	133
Other system data set failures	134
RECON data set recovery	134
Restoring RECON data sets if both are unusable	134
Log errors	135
Log error recovery	136
Log Recovery utility (DFSULTR0)	136
WADS or RDS log recovery	138
Dependent region failures	138
Application program failures	138
Region controller failures	139
Database failures	140
Database recovery	140
Recovering from a network failure when the remote terminal stops responding	141
CPI Communications failures	142
Session failure	142
System failure	142
Recovery processing for CPI Communications driven application programs	143
MSC VTAM message resynchronization and recovery	143
CQS failures and structure failures	144
CQS log recovery	144
CCTL failures	144
CCTL region failure	144
CCTL thread failure	145
CCTL thread looping	145
DBCTL failures	145
IRLM failures	146
Recovery with data sharing	146
DBRC and protecting data	147
IRLM and protecting data	147
Fast Database Recovery (FDBR) regions	148
User access problems	151
Recovery of IMS using Remote Site Recovery (RSR)	152
Preparing for tracking	152
Tracking subsystem and the RSR complex	157
Normal operations for an RSR active site	160
Normal operations for an RSR tracking site	164
RSR termination at the active site	172
Normal RSR termination at the tracking site	172
IMS restart without remote takeover (at the active site)	173

IMS restart without remote takeover (at the tracking site)	173	Secondary master terminal and auditing operational control	250
DRD and RSR	175	Operator control of conversational transactions	251
IMS abnormal termination in an RSR environment.	175	MTO and tracing operations	253
Isolated log transport processing in the logs at the active site	176	Offline dump formatting operations	255
Isolated log transport processing in the logs at the tracking site	177	Designing operating procedures	256
Catch-up processing	177	Operating procedure design using flowcharts and other graphic techniques	256
Online forward recovery at the active site	178	Operating procedure design using narrative	256
Remote takeover	180	IMS-to-IMS TCP/IP connection operations	257
How to coordinate recovery with DB2 for z/OS in a data sharing environment.	191	Viewing configuration and status information for IMS-to-IMS TCP/IP connections in IMS Connect	257
IMS recovery using Extended Recovery Facility	195	Stopping a connection to a remote IMS Connect instance for IMS-to-IMS TCP/IP communications	261
Planning for operating an XRF complex	196	Restarting a connection to a remote IMS Connect instance for IMS-to-IMS TCP/IP communications	262
XRF initialization	199	Stopping IMS Connect send client sockets on IMS-to-IMS TCP/IP connections	263
Actions performed during synchronization	202	Stopping IMS Connect communication with an IMSplex when MSC TCP/IP links are supported	264
Tracking in the XRF complex	205	Starting IMS Connect communication with an IMSplex when MSC TCP/IP links are supported	265
Processing during takeover.	214	Cleaning up MSC logical link resources in IMS Connect	265
Post-takeover	220	Stopping an MSC physical link in IMS Connect	266
Terminating subsystems	222	MSC operations	267
IMS DBCTL capabilities	222	MSC initialization	267
Recovery points	223	MSC termination	268
Creating recovery points	224	Changing logical link assignments	269
Creating recovery points and keeping the database quiesced	225	Restarting a logical link	271
Executing recovery-related functions	226	Switching TCP/IP and VTAM physical link types	271
Issuing DBRC commands	226	MSC TCP/IP link operations	273
Dumping the message queues	227	Commands that help control resources in an MSC environment.	275
Recovering the message queues	227	Displaying information about an MSC network	276
Archiving the OLDS	228	Logical link path control.	278
Making databases recoverable or nonrecoverable	228	Recovery considerations for multiple systems	279
Running recovery-related utilities.	228	Establishing maintenance procedures	281
		Setting up standard JCL	281
		Operator test procedures	281
Chapter 6. Developing operating procedures	229		
Operations personnel.	229	Chapter 7. Developing user procedures	285
Establishing operating procedure documents	232	Procedures for user terminal operators	285
Establishing operations interactions	232	User operator tasks	286
Plan for procedure requirements	233	Operating instructions for terminal operators	286
Planning the content of procedures	234	Potential use of IMS commands	288
How to set the level of operator control	236	Problem reporting for a remote terminal operator	291
Record-keeping procedures for the MTO	237	Connecting to IMS	292
Recording IMS control region activity	238	Communicating with IMS	293
Utility recording forms for IMS utilities	238	Transactions for communicating with IMS.	293
Application program recording forms	240	Operators-to-operator messages	296
Forms for recording problems or unusual events	240	IMS commands used to communicate with IMS	296
Forms for recording startup, shutdown, and system log activity	241	Administration support for errors	298
OM audit trail	243		
Planning availability of IMS service	243		
Checklist for defining the production cycle	244		
Devices available to the network	245		
Operations responsibilities for online control services	245		
MTO abilities and operator responsibilities	246		
Normal operator actions (DB/DC or DCCTL)	247		
Normal operator actions (DBCTL)	248		
Resources controlled by the MTO	249		

Chapter 8. Operations and IMS-supported devices 299

3270 Information Display System.	299
Connecting to IMS (3275 with dial feature-switched line).	299
Disconnecting from IMS (3275 with dial feature-switched Line)	300
Interacting with IMS	300
3270 terminal components that operate with IMS	300
Using programmed symbols for IBM 3270.	313
Remote 3270 errors (VTAM)	316
33/35 teletypewriter	317
Connecting to IMS using the 33/35 teletypewriter	317
Communicating with IMS using the 33/35 teletypewriter	317
Disconnecting from IMS.	318
System console	318
Local card reader	319
Local printer	320
Replace the magnetic tape	320
Disk data sets	320
Programmable remote systems	321
SLU-1 devices	321
Connecting to IMS using SLU-1 devices	322
Communicating with IMS using SLU-1 devices	322
Disconnecting SLU-1 devices from IMS.	324
SLU-2 devices	324
Using SLU-2 devices to connect and disconnect to IMS.	324
3290 information panel	325
NTO logical units	325
Connecting to IMS using NTO devices	325
How to communicate with IMS from an NTO device.	326
Disconnecting NTO devices from IMS	327
Special operational notes and restrictions	327
Master terminal	327

Chapter 9. Automated operations . . . 329

Advantages of automation	329
Deciding what to automate.	330
Tools for automated operations	330
IMS Automated Operator Interface (AOI)	330

REXX SPOC API	343
IMS time-controlled operations	344

Chapter 10. Type-1 automated operator (AO) application program (GU, GN, CMD, and GCMD calls) . . . 355

About the type-1 AO application program (GU, GN, CMD, and GCMD calls)	355
Supported application program environments	356
AO applications for shared-queues	356
Commands used with AO applications (CMD)	357
Format of commands.	357
Responses to commands.	357
Synchronization point processing.	358
Format identifiers for the /DISPLAY command	358
Getting messages from an AO exit routine or a terminal	360
Cross-reference of commands and command response messages.	360
Sample AO application (UETRANS).	363

Chapter 11. Type-2 automated operator (AO) application program (GMSG, ICMD, and RCMD calls) . . . 365

About the type-2 AO application program (GMSG, ICMD, and RCMD calls)	365
Issuing commands and retrieving command responses (ICMD and RCMD calls)	366
Supported application program environments	367
AO application security	367
Commands used with AO applications (ICMD)	367
Responses to commands.	368
Format identifiers for the /DISPLAY command	369
Restart and recovery considerations	370
Sample AO application (DFS AOPGM)	370

Notices 371

Trademarks	373
Privacy policy considerations	373

Bibliography. 375

Index 377

About this information

These topics provide guidance information for selecting tools and options for operating an IMS™ system, for recovering an IMS system, and for automating IMS operations and recovery tasks. The topics also describe how to develop procedures for the master terminal operator and the end user.

This information is available as part of the Information Management Software for z/OS® Solutions Information Center at pic.dhe.ibm.com/infocenter/dzichelp. A PDF version of this information is available in the information center.

Prerequisite knowledge

Before using this information, operators and system programmers need to have some foundation understanding of basic z/OS and IMS concepts, the IMS environment, and your installation's IMS system. Additionally, system programmers should understand administration of the IMS system and databases, z/OS control programs, and VSAM Access Service Methods.

You can learn more about z/OS by visiting the z/OS Basic Skills Information Center.

You can gain an understanding of basic IMS concepts by reading *An Introduction to IMS*, an IBM® Press publication. An excerpt from this publication is available in the Information Management Software for z/OS Solutions Information Center.

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

IMS function names used in this information

In this information, the term HALDB Online Reorganization refers to the integrated HALDB Online Reorganization function that is part of IMS Version 12, unless otherwise indicated.

How new and changed information is identified

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

Revision markers follow these general conventions:

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

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

New and changed information in the information center is denoted by blue carets (<< and >>) at the beginning and end of the new or changed information.


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).




- 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.



fragment-name:



- In IMS, a b symbol indicates one blank position.
- Keywords, and their minimum abbreviations if applicable, appear in uppercase. They must be spelled exactly as shown. Variables appear in all lowercase italic letters (for example, *column-name*). They represent user-supplied names or values.
- Separate keywords and parameters by at least one space if no intervening punctuation is shown in the diagram.
- Enter punctuation marks, parentheses, arithmetic operators, and other symbols, exactly as shown in the diagram.
- Footnotes are shown by a number in parentheses, for example (1).

Accessibility features for IMS Version 12

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

Accessibility features

The following list includes the major accessibility features in z/OS products, including IMS Version 12. 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.

Note: The Information Management Software for z/OS Solutions Information Center (which includes information for IMS Version 12) and its related publications are accessibility-enabled for the IBM Home Page Reader. You can operate all features by using the keyboard instead of the mouse.

Keyboard navigation

You can access IMS Version 12 ISPF panel functions by using a keyboard or keyboard shortcut keys.

For information about navigating the IMS Version 12 ISPF panels using TSO/E or ISPF, refer to the *z/OS TSO/E Primer*, the *z/OS TSO/E User's Guide*, and the *z/OS ISPF User's Guide Volume 1*. These guides describe how to navigate each interface, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

Related accessibility information

Online documentation for IMS Version 12 is available in the Information Management Software for z/OS Solutions Information Center.

IBM and accessibility

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

How to send your comments

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

- From any topic in the information center at pic.dhe.ibm.com/infocenter/dzichelp, click the **Feedback** link at the bottom of the topic and complete the Feedback form.
- Send your comments by e-mail to imspubs@us.ibm.com. Be sure to include the title, the part number of the title, the version of IMS, and, if applicable, the specific location of the text on which you are commenting (for example, a page number in the PDF or a heading in the information center).

Chapter 1. Controlling IMS

You can use the IMS Control Center and the TSO Single Point of Control (TSO SPOC) to issue commands to control IMS and IMS resources.

Issuing commands with the IMS Control Center

You can issue and view IMS type-1 and type-2 commands from the IMS Control Center. The Control Center command output is similar to the output displayed from the TSO SPOC, but it is formatted in a graphical interface based on the Windows system.

The IMS Control Center, included in the IBM DB2 Universal Database™ Control Center, uses the IMS Single Point of Control (SPOC) functions to enable IMS systems management from a workstation. Online wizards exist that help you build and issue commands.

Controlling IMS with the TSO SPOC application

You can use the TSO SPOC to issue operator commands in an IMSplex. The TSO SPOC application uses an ISPF panel interface and communicates with the IMS Operations Manager (OM). OM then communicates with all of the other address spaces in the IMSplex (for example, IMS) as required for operations.

There can be more than one TSO SPOC in an IMSplex. However, the TSO SPOC is optional in an IMSplex.

The TSO SPOC provides the following functions to an IMSplex:

- Presents a single system image for an IMSplex by allowing the user to issue commands to all IMS systems in the IMSplex from a single console.
- Displays consolidated command responses from multiple IMS address spaces.
- Sends a message to an IMS terminal connected to any IMS control region in the IMSplex by using the IMS /BROADCAST command.
- Allows users to create, query, update, and delete various IMS resources online.
- Displays the OM audit trail log which records command input, command responses, and selected system messages from across the IMSplex.
- Allows users to define input user exits to modify or reject command parameters, set return and reason codes, and send text messages to the TSO SPOC session.

There are several ways to issue commands in the IMS TSO SPOC application:

- By command line
- By retrieving a command
 - Using the ISPF RETRIEVE command
 - Using a command listed in the response area
 - Using the Command status panel
- By defining and using command shortcuts

You can use these methods in any combination at any time.

```

File   Action   Manage resources   SPOC   View   Options   Help
-----
PLEX1                                     IMS Single Point of Control
Command ==> _____

----- Plex . ____ Route . ____ Wait . ____
Response for:

CSLM000I (C) Copyright IBM Corp. 2000. All rights reserved.
F1=Help   f3=Exit   F4=Showlog   F6=Expand   F9=Retrieve   F12=Cancel

```

You can issue both IMS type-1 commands and type-2 commands by using the TSO SPOC interface. Enter the command next to the command prompt (Command ==> in the previous screen sample). Enter the IMSplex name in the Plex field. Enter the list of IMS systems to which to route the command, if applicable, in the Route field. If you specify an asterisk (*) in the Route field, the command is routed to all registered command processing clients in the IMSplex. If you specify a percent symbol (%) in the Route field, the command is routed to only one command processing client in the IMSplex that is registered for the command and that has MASTER capability. The Operations Manager chooses the command processing client. If no route list is specified, the default routing is to all registered command processing clients in the IMSplex.

After you type the command, press Enter. The command issued is shown in the Response for: field and the actual command response is shown below the Response for: field.

Tip: If you want to change the default PF key settings on SPOC ISPF panels, use the KEYS or KEYLIST ISPF command.

For more information about the TSO SPOC application, see the IMS TSO SPOC online tutorial. To see the IMS TSO SPOC online tutorial, select **Help > Tutorial** in the application.

Related concepts:

➡ A single point of control (SPOC) program in CSL (System Administration)

Related reference:

 CSLOMCMD: command request (System Programming APIs)

 ISP keylist settings

Starting and setting up the TSO SPOC

After you start the TSO SPOC through the IMS Application menu, you must add the IMS distribution libraries to the TSO user's environment.

To start the TSO SPOC, use the IMS Application menu.

Use one of the following methods to set up the TSO SPOC program:

- Edit the logon procedure to include the IMS distribution libraries or to issue ALLOCATE commands to make the data sets available to the TSO user.
The TSOLIB command of TSO can be used to make the load module data set available to the TSO user. TSOLIB is a command that establishes a STEPLIB-like data set without having to modify the logon procedure.
The data set allocated to ISPTABL is unique to the user. Because of the way ISPF uses tables, the same data set allocated to ISPTABL must also be allocated in the ISPTLIB concatenation and ahead of the IMS.SDFSTLIB data set. The data sets are shown in the following table.

Table 1. Data sets for SPOC setup

Usage	Data set
TSOLIB command or STEPLIB	IMS.SDFSRESL
FILE(ISPPLIB)	IMS.SDFSPLIB
FILE(ISPMLIB)	IMS.SDFSMLIB
FILE(ISPTLIB)	user.ISPTLIB IMS.SDFSTLIB
FILE(ISPTABL)	user.ISPTLIB
FILE(SYSPROC)	IMS.SDFSEXEC

- After the data sets are added to your TSO environment, you can call the TSO SPOC from other applications through a command interface.
- An example of calling TSO SPOC from a rexx program is:
- ```
/* rexx */Address ISPEXEC "SELECT CMD(DFSSPOC CMD(qry tran)
PLEX(plex1))"
```
- Additionally, you can call the TSO SPOC by entering:
- TSO DFSSPOC in any ISPF command line.
  - DFSSPOC in the ISPF option 6 command line followed by any of the optional parameters that appear in the DFSSPOC syntax diagram.
- Use the TSO ALTLIB command and ISPF's LIBDEF service. The DFSSPSRT exec provides an example of starting the DFSSPOC program using ALTLIB and LIBDEF.  
You can invoke DFSSPSRT from ISPF option 6 by entering:  
EXEC 'imslib.SDFSEXEC(DFSSPSRT)' 'HLQ(imslib)'


Related concepts:

 [IMS Application Menu \(System Administration\)](#)

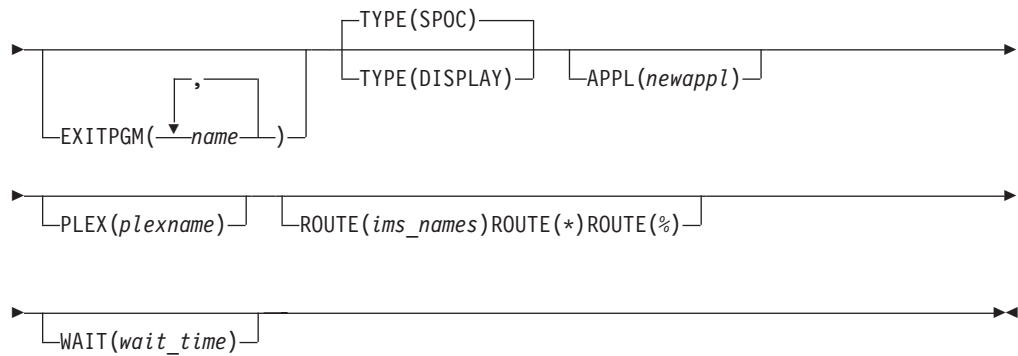
Syntax of the DFSSPOC command

The DFSSPOC command supports several parameters. External programs can issue IMS operator commands and display the command response. You can scroll through the data, but no other SPOC interactions are allowed.

The DFSSPOC command accepts the following parameters:







## DFSSPOC keywords

### DFSSPOC

Specifies the TSO SPOC command name.

### CMD(*command*)

Specifies that a command issue immediately. The response displays in the first SPOC panel. Only commands supported by the OM API can be used.

### OPT(*selection*)

Jumps to a specific Manage Resource menu without viewing intermediate menus where (*selection*) is a value (1, 2, 3, and so on) for a particular menu.

### EXITCMD

Specifies the user exit that TSO SPOC invokes before sending the command to OM. Can be either a single or a list of user exits. Command exits are called with a TSO command processor parameter list. All command exits are called before the program exits. On entry to the routine, register 1 points to the Command Processor Parameter List (CPPL), and is defined by macro IKJCPPL.

### EXITPGM

Specifies the user exit that TSO SPOC invokes before sending the command to OM. Can be either a single or a list of user exits. Program exits are called with a z/OS batch program parameter list. Command exits are called before the program exits. On entry to the routine, register 1 points to a standard parameter list. Register 1 points to a full word, which points to a half-word length followed by the parameter string.

### TYPE(DISPLAY)

Displays the command response only, with no other SPOC interactions.

### TYPE(SPOC)

Indicates that the normal TSO SPOC display be used. This is the default if TYPE is not specified.

### APPL(*newappl*)

Indicates a user specified application ID that is used by ISPF to fence off different applications. This is a one to four character value. The first character is alphabetic and the remaining characters are alphanumeric. If no value is specified, a default of 'CSLU' is used.

### PLEX(*plexname*)

Specifies the name of the IMSplex to which to issue the command. If no IMSplex name is specified, the user's default IMSplex name is used.

**ROUTE(*ims\_names*)**

Specifies which members of the IMSplex to route the command to. Names must be separated by commas.

**ROUTE(\*)**

Specifies that the command is routed to all registered command processing clients in the IMSplex.

**ROUTE(%)**

Specifies that the command is routed to only one command processing client in the IMSplex that is registered for the command and that has MASTER capability. The Operations Manager chooses the command processing client.

**WAIT(*wait\_time*)**

Specifies, in MM:SS format, how long OM waits for member responses before returning a response. The default wait time is 5 minutes.

**Related reference:**

 TSO SPOC user exit routines (Exit Routines)

**Starting the TSO SPOC for the first time**

When you start the TSO SPOC for the first time (after setup is complete), you must set the user preferences.

Perform the following tasks in the TSO SPOC application:

1. Select **Options > Preferences** to display the IMS Single Point of Control Preferences panel.
2. Set the default IMSplex value (this step is required).
3. Optionally, you can specify values for the other preferences or accept the default values. The Preferences panel provides field-level help for each field.
4. Press **Enter**.

You can now enter commands at the TSO SPOC command line.

**Type-1 and type-2 command responses from the TSO SPOC**

When you issue a type-1 command in the TSO SPOC, the command response is displayed in sequential format. The command response consists of messages prefixed by the member name. Information from each member is grouped together in the command response.

When you issue a type-2 command, the command response is displayed in tabular format, with the data displayed in columns. You can sort the list of messages in the command response by positioning the cursor on a column heading (for example, you can sort by member) and pressing Enter. You can also sort the messages by selecting **View > Sort** and selecting the column heading by which you want to sort.

**Restriction:** The OM API supports most type-1 commands.

**Related reference:**

 IMS command language overview (Commands)

**Displaying command status in the TSO SPOC**

You can use the TSO SPOC to display command status, reissue commands that were previously issued, or edit and delete commands.

In the command status panel of the TSO SPOC application, you can:

- Display the commands that you have previously issued in the same TSO SPOC session.
- Display the responses of commands that you have previously issued, if the responses are available.
- Reissue commands that you have previously issued and view the command responses.
- Delete commands.
- Edit commands.

To open the command status panel, select **SPOC > Command status** in the TSO SPOC application.

```
File Action Manage resources SPOC View Options Help

PLEX1 IMS Single Point of Control
Command ==> _____

----- Plex . _____ Route . _____ Wait . _____
Enter '/' to view command response, 'i' to reissue a command, 'd' to delete
a command, and 'e' to edit a command.

Act Status Command
--- ---
Complete DIS STATUS
Complete QRY TRAN NAME(SKS*) SHOW(ALL)
Complete QRY IMSPLEX SHOW(TYPE,STATUS,SUBTYPE)
 /NRE CHKPT 0 FMT ALL
 QRY IMSPLEX SHOW(ALL)
 QRY TRAN NAME(CDEBTRN3) SHOW(STATUS)
 START TRAN CDEBTRN3
 QRY TRAN NAME(CDEBTRN3) SHOW(ALL)
 DIS TRAN CDEBTRN3
 DIS TRAN CDEBTRN3 ALL
 DIS ACT
```

Figure 2. TSO SPOC Command Status panel

#### Related concepts:

“Reissuing commands in the TSO SPOC” on page 7

## Command shortcuts in the TSO SPOC

You can use the command shortcuts option of the TSO SPOC if you prefer to use shortened versions of commands or nicknames for commands. Define the shortcuts (the short commands or nicknames) in the TSO SPOC.

Use the ampersand (&) character as the first character of the shortcut if you want to use it as a nickname.

When you issue a short version of a command, the TSO SPOC appends the additional parameters to the short command. When you issue a nickname for a command, the TSO SPOC replaces the nickname with the full version of the command. The following example shows the SPOC Command Shortcuts panel with some user-defined command shortcuts.

|                            |                |                                |               |              |         |      |
|----------------------------|----------------|--------------------------------|---------------|--------------|---------|------|
| File                       | Action         | Manage resources               | SPOC          | View         | Options | Help |
| -----                      |                |                                |               |              |         |      |
| SPOC Command Shortcuts     |                |                                |               |              |         |      |
| Command ==> _____          |                |                                |               |              |         |      |
| -----                      |                |                                |               |              |         |      |
|                            |                | Plex . _____                   | Route . _____ | Wait . _____ |         |      |
| Act                        | Command        | Additional Parameters          |               |              |         |      |
| _____                      | _____          | _____                          |               |              |         |      |
| _____                      | &QRY1_____     | QRY IMSPLEX SHOW(STATUS)_____  |               |              |         |      |
| _____                      | QRY PGM_____   | SHOW(ALL)_____                 |               |              |         |      |
| _____                      | QRY TRAN_____  | SHOW(CLS PSB QCNT STATUS)_____ |               |              |         |      |
| _____                      | MON LINE ABC07 | PTERM ALL_____                 |               |              |         |      |
| -----                      |                |                                |               |              |         |      |
| ***** Bottom of data ***** |                |                                |               |              |         |      |

Figure 3. TSO SPOC Command Shortcuts panel

In this example, the first command shortcut is &QRY1 (a nickname). When you issue this command shortcut, the TSO SPOC replaces the nickname with the entire command, QRY IMSPLEX SHOW(STATUS). The second command shortcut is QRY PGM. When you issue this command shortcut, the TSO SPOC appends the additional parameters (in this example, SHOW(ALL)) to the short version of the command when it is issued. The full command QRY PGM SHOW(ALL) is issued.

**Related concepts:**

“Reissuing commands in the TSO SPOC”

**Entering command shortcuts in the TSO SPOC**

You can enter command shortcuts in the TSO SPOC instead of using the full command name. When you issue a short version of a command, the TSO SPOC appends the additional parameters to the short command.

To use command shortcuts in the TSO SPOC:

1. Specify the preference for command shortcuts in the IMS Single Point of Control Preferences panel of the TSO SPOC. Select **Options > Preferences** to display this panel.
2. Define your command shortcuts in the SPOC Command Shortcuts panel by selecting **SPOC > Command shortcuts**. To define a nickname, ensure that the first character of the command shortcut is an ampersand (&).

**Defining groups of IMS systems in the TSO SPOC**

You can use the **Route** field in the TSO SPOC application to issue commands to a specific IMS or IMS systems. However, the **Route** field is long enough for only two IMS systems. To work around this restriction, you can define groups of IMS systems in the group definitions panel of the TSO SPOC.

1. Select **Options > Set IMS groups** in the TSO SPOC to display the group definitions panel.
2. Specify the name of the group in the **Route** field to route commands to those specific IMS systems.

**Reissuing commands in the TSO SPOC**

When you issue a command, the command is saved in the TSO SPOC. You can reissue commands in several ways.

- After entering a command, position the cursor on the entered command and press **Enter**. The command is moved into the TSO SPOC command line, and you can edit the command or press **Enter** to issue the command again.
- Use the ISPF retrieve key to display commands that were previously entered. Because the ISPF retrieve key is an ISPF function, all commands entered from other applications are also retrieved.
- Use the ISPF RETP command to list previously entered commands. Because the ISPF retrieve key is an ISPF function, all commands entered from other applications are also retrieved. Select the number of the command shown in the list of commands and press Enter.
- Use the TSO SPOC command status panel. Enter “i” in the **Act** column to reissue the command. If you enter “i” next to several commands and press Enter, the commands are reissued in the order in which they are listed.
- Use the TSO SPOC command shortcuts panel. Enter “i” in the **Act** column to reissue the command. If you enter “i” next to several commands and press Enter, the commands are reissued in the order in which they are listed.

**Related concepts:**

“Displaying command status in the TSO SPOC” on page 5

“Command shortcuts in the TSO SPOC” on page 6

## Managing IMS resources by using TSO SPOC

The TSO SPOC application provides a set of resource management panels to help you manage IMS resources. You can use TSO SPOC to query IMS resources and, if dynamic resource definition (DRD) is enabled in your IMS systems, you can also use the TSO SPOC panels to create, delete, export, import, and update IMS resources.

You can access the panels from the Manage Resources menu in the TSO SPOC application or by selecting the **Manage Resources** option directly from the IMS Application Menu.

### Querying IMS resource information by using TSO SPOC

You can query the attributes and status of IMS resources as maintained in each IMS system or across the IMSplex. The information returned for IMS resources can include:

- The status of resources as set by commands and maintained locally by IMS systems
- The status of resources as set by commands and maintained globally across the IMSplex by Resource Manager (RM)
- The status of resources as set in the DBRC RECON data set
- The attributes of various IMS resources as set by system definition macros or by DRD commands
- The attributes status and attributes information about an IMSRSC repository

### Creating, deleting, exporting, importing, and updating IMS resources by using TSO SPOC

Before you can create, delete, export, import, or update IMS resources through TSO SPOC by using the dynamic resource definition commands, you should be aware of the following points:

- You must enable DRD in the target IMS systems.

- Changes made to IMS resources are not saved across an IMS cold start until they are exported to the resource definition data set (RDDS) or the repository. If a cold start or emergency restart with the COLDSYS parameter occurs before an export is performed, any changes to resource definitions are lost.
- After creating an application program or database by using the CREATE command, you must further define the application program or database, as well as its relationships to other resources, by running the appropriate generation utility, such as the Database Description Generation (DBDGEN) utility or the Program Specification Block Generation (PSBGEN) utility. You must perform an ACBLIB or ACBMBR online change to ensure that the created or modified DBD or PSB resources are available to the IMS system.

The IMS resources and resource descriptors that you can create, delete, export, import, and update in DRD-enabled systems include:

- Application programs
- Databases
- Fast Path routing codes
- Transactions

#### Related tasks:

 Enabling dynamic resource definition (System Definition)

## Viewing the OM audit trail log by using TSO SPOC

You can view the OM audit trail log by using the TSO SPOC application if the CSL Operations Manager (OM) is configured to produce an audit trail log. The OM audit trail log contains log records of commands, command responses, and system messages routed through OM.

OM audit trail logging is enabled by the AUDITLOG parameter in the CSL OM initialization parameters PROCLIB member (CSLOIxxx).

You can filter the records returned by TSO SPOC by specifying in the Preferences panel one or more IMSplex member names for which you want to see records. Alternately, you can specify the type of IMSplex member, such as IMS systems, automated operator programs, CQS, RM, and so forth.

When you are viewing the OM audit trail in TSO SPOC, a log record for command input, command output, or system message appears on a separate line. You can obtain additional information by:

- Clicking on a log entry for a system message, which produces a link that you can follow to a description of the message.
- Clicking on a log entry for a command response, which produces information about that command.

To search for a text string in the TSO SPOC Audit Trail panel, type **find** on the command line or click **View > Find** from the menu bar. A pop-up panel is displayed, prompting you to enter the text string. To find the same text string again, press PF5 or type **rfind** on the command line.

```

File Display View Options Help

PLEX1 IMSplex Audit Trail
Command ==>
----- Members. . Type . .
MbrName Time Message
IMS1 2006.086 14:04:21.148419 DFSxxxxI message text
IMS1 2006.086 14:08:18.898122 DFSxxxxW message text
IMS1 2006.086 14:10:08.752387 DFSxxxxE message text
OMUSER1 2006.086 14:15:08.752387 Response for: QRY IMSPLEX
IMS1 2006.086 15:08:21.148419 DFSxxxxI message text
IMS1 2006.086 15:09:18.898122 DFSxxxxW message text

F1=Help F7=Up F8=Down F12=Cancel

```

Figure 4. TSO SPOC IMSplex Audit Trail panel

When you are using the TSO SPOC ISPF interface, the log entries for system messages are color coded in the TSO SPOC display. Error messages are displayed in red. Warning messages are displayed in yellow. All other messages are displayed in white.

### Enabling the OM audit log display

You can enable the OM audit log display in order to display the current records that are being logged.

To enable the OM audit log display, follow these steps:

1. Optionally define your audit trail preferences in the TSO SPOC Preferences panel.
2. Select **Audit Trail** from the panel selection screen.
3. Enter the z/OS log data set name that contains the OM audit log. This name must match the z/OS log data set name specified in the AUDITLOG subparameter of the IMSPLEX parameter in the CSLOIxxx PROCLIB member.
4. Enter the earliest date and time for which you want to see records.
5. Enter the latest data and time for which you want to see records. Enter an asterisk in the end date and end time fields to see the current records being logged.

## Issuing Batch SPOC commands

You can use the Batch SPOC utility to submit IMS operator commands to an IMSplex. The Batch SPOC utility accepts any commands that are supported by the OM API.

The Batch SPOC utility uses:

- Program parameters to define the IMSplex environment, including IMSplex name, routing and the wait time.
- The SYSIN file as input for IMS operator commands.
- The SYSPRINT file to show SPOC-like formatted command response.

You can invoke the Batch SPOC utility using standard JCL statements. The following example shows a simple invocation, but you can call the utility using other valid JCL.



## Sample batch job with multiple commands

```
//SPOCJOB JOB ,
// MSGCLASS=H,NOTIFY=&SYSUID,USER=&SYSUID/*
//SPOC EXEC PGM=CSLUSPOC,
// PARM=(' IMSPLEX=PLEX1,ROUTE=IMS3,WAIT=30,F=WRAP')
//STEPLIB DD DISP=SHR,DSN=IMS.SDFSRESL
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
 QRY IMSPLEX SHOW(JOB,TYPE, +
 STATUS)

 QRY TRAN NAME(INV1*) SHOW(ALL) /* inventory appl */
/*EOF
```

The following program parameters define the IMSplex environment:

### IMSPLEX

Required parameter that specifies the 1- to 5-character suffix of the IMSplex name.

- F** Optional parameter that specifies the print format of the SPOC output. You can specify one of the following values:

#### WRAP

Wraps to the next line as needed. This is the default.

#### BYCOL

Lines of data are grouped together by the column.

#### BYRSC

Lines of data are grouped together by the resource.

### ROUTE

Optional parameter that specifies the SYSIDs of IMSplex members that are to execute the command. If ROUTE is not specified, all members of the IMSplex will execute the command. If more than one member is specified, enclose the list in parenthesis and separate the names with commas.

If ROUTE=\* is specified, the command is routed to all registered command processing clients in the IMSplex. If ROUTE=% is specified, the command is routed to only one command processing client in the IMSplex that is registered for the command and that has MASTER capability. The Operations Manager chooses the command processing client.

### WAIT

Optional parameter that specifies the wait time for individual commands. The wait value is in minutes and seconds (MMM:SS) or just seconds (SSSSS). OM will return a single response as soon as a response is received from all of the members of the IMSplex. If the interval expires, OM will return any responses from IMSplex members, plus an indication that some did not reply. The Batch SPOC utility will wait for each command to complete before issuing the next command. The default wait value is five minutes (5:00). The WAIT time applies to every command in the SYSIN file. The user can specify a wait time of zero seconds; in this case, the batch SPOC issues a command but does not wait for the response.

The SYSIN file is provided by the user and contains the commands that the user wants to execute. The commands are executed serially. When one command completes, the next command is executed until all records from the SYSIN file are processed. Continuation of the SYSIN control statements is specified by a plus sign (+) or a minus sign (-) as the last nonblank character of the line. A plus sign

removes the leading spaces from the next line; a minus sign keeps leading spaces. Comments can be included within the SYSIN file and are specified with the following format:

```
/* this is a comment */
```

The SYSPRINT file will have the formatted command response. If more than one command is issued, the responses appear in the same order as the commands appear in the SYSIN file. The default record length is 133. The command response is formatted in a style specified by the user. The user may specify DCB information in the JCL or in the data set allocation to allow longer records in the SYSPRINT file.

System Display and Search Facility (SDSF) can be used to view batch job output.

## Sample batch job output

The following example shows sample batch job output:

```
=====
Log for. . : QRY IMSPLEX SHOW(JOB,TYPE,STATUS)

IMSpIplex : PLEX1
Routing :
Start time. . . . : 2005.132 15:36:28.11
Stop time : 2005.132 15:36:29.17
Return code . . . : 00000000
Reason code . . . : 00000000
Command master. . : SYS3

IMSpIplex MbrName CC Member JobName Type Status
CSLPLEX1 OM10M 0 USRT002 USRT002 AOP ACTIVE
CSLPLEX1 OM10M 0 OM10M OM1 OM READY,ACTIVE
CSLPLEX1 OM10M 0 RM1RM RM1 RM READY,ACTIVE
CSLPLEX1 OM10M 0 SCI1SC SCI1 SCI READY,ACTIVE
CSLPLEX1 OM10M 0 IMS1 IMS1 IMS READY,ACTIVE
CSLPLEX1 OM10M 0 SYS1 SYS1 IMS READY,ACTIVE
=====
```

## Sample batch job output with no response

If no wait time, WAIT=0, is specified, the command response is not available and therefore will not be printed. The SYSPRINT file will only have short summary information for each command:

```
=====
Log for. . : QRY IMSPLEX SHOW(JOB,TYPE,STATUS)

IMSpIplex : PLEX1
Routing :
Start time. . . . : 2006.075 15:36:28.11
=====
```

---

## Modifying and controlling system resources

You establish the initial settings of IMS resources during IMS system definition.

### Modifying system resources online

IMS supports two separate functions that support modifying IMS resources online: dynamic resource definition (DRD) and the online change function. Authorized system operators and database administrators (DBAs) can change various system resources using IMS commands.

### *Dynamic resource definition*

With DRD, you can issue commands to create, update, and delete the resources in the MODBLKS data set dynamically, or use the enhanced Destination Creation exit routine (DFSINSX0), formerly called the Output Creation exit routine, to create transactions (and, if necessary, the programs that are associated with the transactions) instead of using the online change process. DRD enables you to avoid the online change process only for resources in the MODBLKS data set. To change other resources online (for example, resources in IMS.ACBLIB), you still need to use the online change process.

DRD is enabled by specifying MODBLKS=DYN in the DFSDFxx member or the DFSCGxxx member.

When dynamic resource definition (DRD) is enabled in your IMS system, you can issue the following type-2 commands to create, modify and delete runtime resource and descriptor definitions for application programs, databases, routing codes and transactions.

- CREATE
- UPDATE
- DELETE
- IMPORT

The CREATE, UPDATE, DELETE and, IMPORT commands in a DRD-enabled system are essentially the same as adding, modifying, or deleting the system definition macros for IMS runtime resource definitions.

The EXPORT command is used to export runtime resource and descriptor definitions from the online IMS system to a resource definition data set (RDDS) or IMSRSC repository. The resource and descriptor definitions can then be added to IMS dynamically through the IMPORT command.

Exporting runtime resource and descriptor definitions to a repository allows you to create resources for another IMS in the IMSplex.

Modifications to IMS resources that are made using the type-2 commands are recoverable across a warm or emergency restart. Runtime descriptor definitions are lost if an IMS system is cold started, unless you exported your resource and descriptor definitions to the RDDS or repository while IMS was running. If you want to modify resources, use the QUERY command to view the resource information before making the change.

You must ensure that:

- Commands are routed to and executed on the systems where you want the changes to be applied.
- Any changes made dynamically with the DRD commands are recovered across a cold start. Here are some examples of possible methods:
  - Using automatic export to and then automatic import from the RDDS data set.  
Instead of using automatic export, you can also issue the EXPORT command while IMS is running.
  - Issuing the EXPORT DEFN TARGET(REPO) command while IMS is running and then using automatic import from the repository.

| You can periodically issue the EXPORT DEFN TARGET(REPO)  
| OPTION(CHANGESONLY) command to write changes to the repository.

| The following IMS resources and resource descriptors can be defined dynamically:

- Application programs
- Databases
- Fast Path routing codes
- Transactions

You can enter the CREATE, UPDATE, DELETE, and MODIFY type-2 commands directly through:

- TSO SPOC
- IMS Control Center
- Batch SPOC
- REXX automation (through the REXX SPOC API)
- Your own programs that use the OM API

Similarly, you can use the Manage Resource panels to enter DRD commands by selecting an option from the IMS Application Menu. The Manage Resource panels are a hierarchy of ISPF panels that offer an alternate way to enter commands. From the Manage Resource panels, you can:

- Define and control IMS resources
- Select an action that you want to perform (for example, CREATE)
- Manage attributes of a resource

MORE: +

|          |                                    |      |                 |
|----------|------------------------------------|------|-----------------|
| * NAME   | Transaction name . . 4NEWTRAN      |      |                 |
| SET      |                                    |      |                 |
| AOCMD    | AOI command option . . . . .       | N    | CMD, N, Tran, Y |
| CLASS    | Class. . . . .                     | 1    | 1-999           |
| CMTMODE  | Commit mode. . . . .               | SNGL | Sngl, Mult      |
| CONV     | Conversational . . . . .           | N    | Y, N            |
| DCLWA    | Log write-ahead option . . . . .   | Y    | Y, N            |
| DIRROUTE | MSC direct routing option. . . . . | N    | Y, N            |
| EDITRTN  | Input edit routine . . . . .       |      |                 |
| EDITUC   | Edit to uppercase. . . . .         | Y    | Y, N            |
| EMHBSZ   | EMH buffer size. . . . .           |      | 12-30720        |

Figure 5. Example of a Manage Resource panel

You can use many IMS commands to perform similar control functions for different types of resources.

Online change function

There are two variations of online change:

- Local online change, which enables you to make online changes to IMS resources for one IMS.
- Global online change, which enables you to coordinate online changes to IMS resources across all IMS systems in an IMSplex.

Global online change is not supported in all IMS environments that support local online change. Environments that require the MODSTAT/MODSTAT2 data sets do

not support global online change. FDBR, XRF alternate and DBCTL-standby do not participate in global online process, however they keep track of the online change from the log records.

*Table 2. IMS environments that support online change*

| Environment   | Supports local online change? | Supports global online change? |
|---------------|-------------------------------|--------------------------------|
| DB/DC         | Yes                           | Yes                            |
| DBCTL         | Yes                           | Yes                            |
| DBCTL-standby | No                            | No                             |
| DCCTL         | Yes                           | Yes                            |
| FDBR          | No                            | No                             |
| RSR tracker   | Yes                           | No                             |
| XRF active    | Yes                           | Yes                            |
| XRF alternate | No                            | No                             |

**Related concepts:**

➡ Overview of the IMSRSC repository (System Definition)  
 “Creating, updating, deleting, and querying resource definitions dynamically” on page 25

➡ Overview of the local online change function (System Administration)

➡ Overview of dynamic resource definition (System Definition)

**Related tasks:**

➡ Exporting resource and descriptor definitions (System Definition)

➡ Importing resource and descriptor definitions (System Definition)

**Related reference:**

“List of commands with similar functions for multiple resources”

## List of commands with similar functions for multiple resources

These tables show the IMS commands and how they affect certain resources.

**Note:** In the following tables, a blank indicates that the command is not applicable for the resource.

**Subsections:**

- “Telecommunication line, physical terminal, or node” on page 16
- “Logical terminal” on page 16
- “MSC logical link (MSLINK) commands” on page 17
- “MSC physical link (MSPLINK) commands” on page 18
- “MSC logical link path (MSNAME) commands” on page 18
- “Transaction” on page 20
- “Transaction class” on page 21
- “Program” on page 22
- “Database” on page 22
- “Area” on page 22
- “Subsystem” on page 23

- “User” on page 23

## Telecommunication line, physical terminal, or node

The following table shows the IMS commands that affect the telecommunication line, physical terminal, or node resources. This table indicates whether each resource can perform the following functions after the command is issued:

- Receive input
- Send output
- Output message queuing

*Table 3. IMS commands that affect telecommunications line, physical terminal, or node resources.*

| IMS command | Receive input | Send output | Output message queuing |
|-------------|---------------|-------------|------------------------|
| /ASSIGN     | Y             | Y           | Y                      |
| /LOCK       | N             | N           | Y                      |
| /MONITOR    | Y             | N           | Y                      |
| /PSTOP      | N             | N           | Y                      |
| /PURGE      | N             | Y           | Y                      |
| /RSTART     | Y             | Y           | Y                      |
| /START      | Y             | Y           | Y                      |
| /STOP       | N             | N           | Y                      |
| /UNLOCK     | Y             |             | Y                      |

**Note:** /MONITOR, /PSTOP, /PURGE, and /RSTART refer to the telecommunication line or physical terminal, not to the node.

## Logical terminal

The following table shows the IMS commands that affect logical terminal resources. This table indicates whether these resources can perform the following functions after the command is issued:

- Receive input
- Send output
- Queuing from other terminals

*Table 4. IMS commands that affect logical terminal resources*

| IMS command | Receive input | Send output | Queuing from other terminals |
|-------------|---------------|-------------|------------------------------|
| /ASSIGN     | Y             | Y           | Y                            |
| /LOCK       | N             | N           | N                            |
| /PSTOP      | N             | N           | Y                            |
| /PURGE      | N             | Y           | N                            |
| /RSTART     |               |             |                              |
| /START      | Y             | Y           | Y                            |
| /STOP       | N             | N           | N                            |
| /UNLOCK     |               | Y           | Y                            |

## MSC logical link (MSLINK) commands

Several commands are used to control logical links to other systems (MSLINKs). Logical link definitions are defined with the MSLINK macro which, among other options, allows you to name the link, to associate the logical link with a logical link defined in a partner system, and to define the type of physical link the logical link can be used with.

The following table describes the effect of each of the commands that you can use to control logical link definitions.

*Table 5. Commands used to control MSC logical links definitions*

| MSC use and effect                                                                                                                                             | Type-1 command                                              | Type-2 command                                                  |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------|-----------------------------------------------------------------|
| Set or reset the ASR.                                                                                                                                          | /CHANGE LINK <i>link</i> ASR ON OFF                         | UPDATE MSLINK SET(ASR(ON OFF))                                  |
| Change the session resynchronization option.                                                                                                                   | /CHANGE LINK <i>link</i><br>FORCSESS SYNCSSESS <br>COLDSESS | UPDATE MSLINK SET<br>(SYNCOPT(FORCSESS SYNCSSESS <br>COLSDESS)) |
| Change the default modetable name.                                                                                                                             | /CHANGE LINK <i>link</i> MODE <i>name</i>                   | UPDATE MSLINK SET(MODETBL( <i>name</i> ))                       |
| Assign to a new physical link.                                                                                                                                 | /MSASSIGN LINK <i>link</i> MSPLINK <i>name</i>              | UPDATE MSLINK SET(MSPLINK( <i>name</i> ))                       |
| Stop the current link and sending and receiving of messages.                                                                                                   | /PSTOP LINK <i>link</i>                                     | UPDATE MSLINK STOP(COMM)                                        |
| Force the stoppage of a link.                                                                                                                                  | /PSTOP LINK <i>link</i> PURGE FORCE                         | UPDATE MSLINK STOP(COMM)<br>OPTION(FORCE)                       |
| Start a previously stopped MSLINK and start the queuing of or sending of messages to logical links on another system.                                          | /RSTART LINK <i>link</i>                                    | UPDATE MSLINK START(COMM)                                       |
| Start a previously stopped MSLINK and start the queuing of or sending of messages to logical links on another system, but use modetable for this session only. | /RSTART LINK <i>link</i> MODE <i>name</i>                   | UPDATE MSLINK START(COMM)<br>SET(MODETBL( <i>name</i> ))        |
| Start or stop an internal trace.                                                                                                                               | /TRACE SET ON OFF LINK <i>link</i>                          | UPDATE MSLINK START(TRACE) or<br>UPDATE MSLINK STOP(TRACE)      |
| Start or stop an XRF takeover trace.                                                                                                                           | /TRACE SET ON OFF LINK <i>link</i><br>TAKEOVER              | UPDATE MSLINK START(TKOTRC) or<br>UPDATE MSLINK STOP(TKOTRC)    |
| Start or stop the bandwidth mode.                                                                                                                              | Not applicable                                              | UPDATE MSLINK<br>SET(BANDWIDTH(ON OFF))                         |
| Change the send/receive buffer size.                                                                                                                           | Not applicable                                              | UPDATE MSLINK SET(BUFSIZE( <i>size</i> ))                       |
| Change the logical link name.                                                                                                                                  | Not applicable                                              | UPDATE MSLINK SET(MSLINK( <i>name</i> ))                        |
| Change the partner id.                                                                                                                                         | Not applicable                                              | UPDATE MSLINK SET(PARTNER( <i>id</i> ))                         |



## MSC physical link (MSPLINK) commands

Several commands are used to control physical links (MSPLINKs). The MSPLINK macro can define three different types of connections between two systems:

- Channel-to-channel (CTC)
- Memory-to-memory (MTM)
- VTAM®

The following table describes the effect of each of the commands that you can use with to control physical link paths to IMS systems.

*Table 6. Commands used to control MSC physical links*

| MSC use and effect                                                | Type-1 command  | Type-2 command                                |
|-------------------------------------------------------------------|-----------------|-----------------------------------------------|
| Disable logons to this physical link.                             | /PSTOP MSPLINK  | UPDATE MSPLINK<br>STOP(LOGON)                 |
| Enable logons to this physical link.                              | /RSTART MSPLINK | UPDATE MSPLINK<br>START(LOGON)                |
| Set or reset ASR for all assigned logical links.                  | Not applicable  | UPDATE MSPLINK<br>SET(ASR(ON OFF))            |
| Change the default modetable name for all assigned logical links. | Not applicable  | UPDATE MSPLINK<br>SET(MODETBL( <i>name</i> )) |
| Change the physical link name.                                    | Not applicable  | UPDATE MSPLINK<br>SET(MSPLINK( <i>name</i> )) |
| Change the VTAM node name.                                        | Not applicable  | UPDATE MSPLINK<br>SET(NODE( <i>name</i> ))    |

## MSC logical link path (MSNAME) commands

Several commands are used to control logical link paths (MSNAMEs). Logical link paths are defined by system identifier (SYSID) pairs that identify sending and destination systems. Logical link paths are named by the MSNAME macro.

The following table describes the effect of each of the commands that you can use to control logical link paths resources.

*Table 7. Commands used to control MSC logical link paths*

| MSC use and effect                                | Type-1 command   | Type-2 command                              |
|---------------------------------------------------|------------------|---------------------------------------------|
| Change the local system identification SIDs.      | Not applicable   | UPDATE MSNAME<br>SET(SIDL(id))              |
| Change the remote system identification SIDs.     | Not applicable   | UPDATE MSNAME<br>SET(SIDR(id))              |
| Assign a new logical link path to a logical link. | /MSASSIGN MSNAME | UPDATE MSNAME<br>SET(MSLINK( <i>name</i> )) |

Table 7. Commands used to control MSC logical link paths (continued)

| MSC use and effect                                                                                                                                                                                                                                                                                                                                                                                        | Type-1 command | Type-2 command                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|---------------------------------------|
| <p>Halt the queuing of primary requests for all remote terminals and programs represented by the MSNAME.</p> <p>Continued conversations and secondary requests are handled. Primary requests entered through an input terminal receive message DFS065. Requests from other systems that require use of the logical link path for a response are not accepted but remain queued in the sending system.</p> | /PURGE MSNAME  | UPDATE \MSNAME STOP(Q)<br>START(SEND) |
| Start a previously stopped MSNAME and start the queuing of or sending of messages to logical link paths.                                                                                                                                                                                                                                                                                                  | /START MSNAME  | UPDATE MSNAME<br>START(Q,SEND)        |

Table 7. Commands used to control MSC logical link paths (continued)

| MSC use and effect                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Type-1 command | Type-2 command                |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|-------------------------------|
| Stop the sending and receiving of primary request messages associated with the logical link path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | /STOP MSNAME   | UPDATE MSNAME<br>STOP(Q,SEND) |
| <p>When an MSNAME is stopped by an input system, primary requests for remote programs or terminals associated with the stopped logical link path are canceled, and message DFS065 is returned to the input terminal. Conversations in progress are allowed to continue.</p> <p>When an MSNAME is stopped by a destination system, messages received from other systems over a stopped logical link path cause a logical link path to be stopped in the sending system (input or intermediate), and MTOs of the sending and receiving systems to be notified by messages DFS2140 and DFS2142, respectively. The message remains queued in the sending system until the logical link is subsequently started in both systems.</p> |                |                               |

## Transaction

The following table shows the IMS commands that affect transaction resources. This table indicates whether these resources can perform the following functions after the command is issued:

- Message scheduling by transaction
- Message queuing by transaction

Table 8. IMS commands that affect transaction resources

| IMS command           | Message scheduling by transaction | Message queuing by transaction |
|-----------------------|-----------------------------------|--------------------------------|
| /ASSIGN               | Y                                 | Y                              |
| or<br>UPDATE TRAN SET |                                   |                                |

Table 8. IMS commands that affect transaction resources (continued)

| IMS command                                     | Message scheduling by transaction | Message queuing by transaction |
|-------------------------------------------------|-----------------------------------|--------------------------------|
| /LOCK<br>or<br>UPDATE TRAN SET(LOCK ON OFF)     | N                                 | Y                              |
| /MSASSIGN<br>or<br>UPDATE TRAN SET(MSNAME)      | Y                                 | Y                              |
| /PSTOP<br>or<br>UPDATE TRAN STOP(SCHD) START(Q) | N                                 | Y                              |
| /PURGE<br>or<br>UPDATE TRAN START(SCHD) STOP(Q) | Y                                 | N                              |
| /START<br>or<br>UPDATE TRAN START(Q,SCHD)       | Y                                 | Y                              |
| /STOP<br>or<br>UPDATE TRAN STOP(Q,SCHD)         | N                                 | N                              |
| /UNLOCK<br>or<br>UPDATE TRAN SET(LOCK ON OFF)   | Y                                 | Y                              |

## Transaction class

The following table shows the IMS commands that affect transaction class resources. This table indicates whether these resources can perform transaction scheduling by class after the command is issued.

Table 9. IMS commands that affect transaction class resources.

| IMS command           | Transaction scheduling by class |
|-----------------------|---------------------------------|
| /ASSIGN               | Y                               |
| /MSASSIGN             | Y                               |
| /START                | Y                               |
| /STOP                 | N                               |
| UPDATE TRAN           | Y or N <sup>1</sup>             |
| UPDATE TRAN SET       | Y                               |
| UPDATE TRAN START (Q) | Y                               |
| UPDATE TRAN STOP (Q)  | N                               |

<sup>1</sup> Whether the transaction class can perform transaction scheduling by class after the UPDATE command is issued depends on the parameters and keywords specified in this command.

## Program

The following table shows the IMS commands that affect program resources. This table indicates whether the program can be started after the command is issued.

*Table 10. IMS commands that affect program resources*

| IMS command               | Can the program be started? |
|---------------------------|-----------------------------|
| /ASSIGN                   | Y                           |
| /LOCK                     | N                           |
| /START                    | Y                           |
| /STOP                     | N                           |
| /UNLOCK                   | Y                           |
| UPDATE PGM SET(LOCK(ON))  | N                           |
| UPDATE PGM SET(LOCK(OFF)) | Y                           |
| UPDATE PGM START(SCHD)    | Y                           |
| UPDATE PGM STOP(SCHD)     | N                           |

## Database

The following table shows the IMS commands that affect database resources. This table indicates whether the database can be accessed after the command is issued.

*Table 11. IMS commands that affect database resources.*

| IMS command               | Can the database be accessed? |
|---------------------------|-------------------------------|
| /DBDUMP                   | N                             |
| /DBRECOVERY               | N                             |
| /LOCK                     | N                             |
| /START                    | Y                             |
| /STOP                     | N                             |
| /UNLOCK                   | Y                             |
| UPDATE                    | Y or N <sup>1</sup>           |
| UPDATE DB START(ACCESS)   | Y                             |
| UPDATE DB STOP(ACCESS)    | N                             |
| UPDATED DB STOP(UPDATES)  | Y                             |
| UPDATE DB STOP(SCHD)      | N                             |
| UPDATE DB SET(LOCK(ON))   | N                             |
| UPDATE DB SET (LOCK(OFF)) | N                             |

<sup>1</sup> Whether the database can be used after the UPDATE command is issued depends on the parameters and keywords specified in this command.

## Area

The following table shows the IMS commands that affect area resources. This table indicates whether the area can be accessed after the command is issued.

Table 12. IMS commands that affect area resources.

| IMS command               | Can the area be accessed? |
|---------------------------|---------------------------|
| /DBRECOVERY               | N                         |
| /START                    | Y                         |
| /STOP                     | N                         |
| UPDATE                    | Y or N <sup>1</sup>       |
| UPDATE AREA START(ACCESS) | Y                         |
| UPDATE AREA STOP(ACCESS)  | N                         |
| UPDATE AREA STOP(SCHD)    | N                         |

<sup>1</sup>Whether the database can be used after the UPDATE command is issued depends on the parameters and keywords specified in this command.

## Subsystem

The following table shows the IMS commands that affect subsystem resources. This table indicates whether the subsystem can be attached after the command is issued.

Table 13. IMS commands that affect subsystem resources

| IMS command | Can the subsystem be attached? |
|-------------|--------------------------------|
| /START      | Y                              |
| /STOP       | Y                              |
| /CHANGE     | N                              |

## User

The following table shows the IMS commands that affect user resources. This table indicates whether the user resource can perform the following functions after the command is issued:

- Receive input
- Send output
- Output message queuing

Table 14. IMS commands that affect user resources

| IMS command | Receive input | Send output | Output message queuing |
|-------------|---------------|-------------|------------------------|
| /ASSIGN     | Y             | Y           | Y                      |
| /RSTART     | Y             | Y           | Y                      |
| /START      | Y             | Y           | Y                      |
| /STOP       | N             | N           | Y                      |

#### **Related concepts:**

“Commands for IMS operations tasks” on page 47

“Modifying system resources online” on page 12

## **Modifying dependent regions**

To modify the assignment of classes to regions or to adjust the processing load among message regions, you can use the type-1 /ASSIGN command or the type-2 UPDATE command.

Use the /ASSIGN TRAN or UPDATE TRAN SET(CLASS(*new\_class\_number*)) command to modify the assignment of classes to regions.

## **Modifying telecommunication lines**

To discard response-mode output messages, you can use the /DEQUEUE command

Use the /DEQUEUE command to discard response-mode output messages before you enter an /RSTART LINE command.

## **How to modify terminals**

Terminals and links to terminals can be modified through various type-1 commands. You can use these commands and their variations to manage terminals.

Use the /DISPLAY PTERM and /DISPLAY NODE command to display the status of terminals and nodes. You can also use the /DISPLAY STATUS command, which shows all resources that have a status requiring operator intervention, including terminals and nodes.

Use the /ASSIGN LTERM command to modify the assignment of logical terminals to physical terminals or nodes. The new assignment remains in effect until the next cold start or until you issue another /ASSIGN command.

Use the /DEQUEUE command to discard full-function response-mode output so that the /RSTART command can reset terminal response mode.

You can reset static nodes that are hung in Fast Path input response mode by issuing the /STOP NODE and /START NODE commands in sequence.

Use the /COMPT command for VTAM terminals (nodes) to notify IMS that a terminal component is operable or inoperable.

IMS provides a VTAM I/O Timeout facility to detect VTAM hung nodes and determine what action, if any, should be taken. Use the /TRACE command to start and stop the VTAM I/O Timeout facility. Use the /IDLE command to deactivate a node and the /ACTIVATE command to activate a node. Use the /DISPLAY command to display all nodes that have I/O outstanding for a time period greater than that specified during system definition.

## **Modifying and controlling transactions**

You can use the /ASSIGN TRAN or UPDATE TRAN SET(CPRI(*new\_current\_priority*)) command to reassign the scheduling priorities



established for transactions during system definition. The new assignments remain in effect until the next cold start or until you issue another /ASSIGN or UPDATE command.

In a shared-queues environment, you can use the /ASSIGN TRAN or UPDATE TRAN SET(CLASS(*new\_class\_number*)) command to control which IMS subsystems can run certain types of transactions by assigning transactions to a particular class.

For example, you can define TRANA to class 4 on IMSA and to class 255 on IMSB and IMSC, so that only IMSA can run TRANA. If IMSA fails, you can reassign TRANA on either IMSB or IMSC to a class that these IMS subsystems can run.

**Attention:** Do not use the /STOP TRAN or UPDATE TRAN STOP(Q,SCHD) command to control which IMS subsystems can run certain types of transactions.

## Database control

You can use the type-1 /DBDUMP DB or type-2 UPDATE DB commands to control databases and access to databases.

Use the /DBDUMP DB or UPDATE DB STOP(UPDATES) command to stop online update access to a database so that you can produce an offline dump of the database.

Use the /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command to stop all online access to a database. Use it to recover a database offline.

Normally, IMS switches to using the next OLDS when you enter the /DBDUMP or /DBRECOVERY command. This switch does not occur if you specify the NOFEOV keyword on either command. When you enter the UPDATE DB STOP(ACCESS) command to stop online access to the database, IMS does not switch to the next OLDS. You can specify OPTION(FEOV) if you want IMS to switch to the next OLDS.

Specify the GLOBAL keyword on the /DBDUMP or /DBRECOVERY command to have the command apply to all subsystems sharing the database. The IRLM must be active if you use this keyword. The default is LOCAL, which specifies that the command applies only to the subsystem on which you enter the command.

## Creating, updating, deleting, and querying resource definitions dynamically

If dynamic resource definition (DRD) is enabled in your IMS system, you can create, update, delete, and query certain IMS resources online through several type-2 commands. You can also create runtime descriptor definitions to use as templates for creating additional resources.

You can use the following type-2 commands:

- CREATE
- UPDATE
- DELETE
- QUERY
- IMPORT

You can also create *runtime descriptor definitions*, which are a set of resource definitions that you can use as templates to create additional resources.

The following IMS resources and resource descriptors can be dynamically defined:

- Application programs
- Databases
- Fast Path routing codes
- Transactions

You can use Fast Path routing codes and transactions as soon as they are correctly defined by the CREATE command.

After creating an application program or database by using the CREATE command, you must further define the application program or database, as well as its relationships to other resources, by running the appropriate generation utility, such as the Database Description Generation (DBDGEN) utility or the Program Specification Block Generation (PSBGEN) utility.

You can also create and modify these resources by using the Managing Resources series of ISPF panels which are invoked through the IMS applications menu.

**Related concepts:**

 Dynamic resource definition (System Definition)

“Modifying system resources online” on page 12

## **Creating runtime resource definitions or runtime descriptor definitions dynamically**

The CREATE commands can be used to add resources dynamically without having to go through the online change process. You can use the CREATE command with the appropriate parameters to define a new runtime resource definition or runtime descriptor definition.

Runtime resource definitions and runtime descriptor definitions that are created by using the CREATE command exist until the next IMS cold start, unless they are deleted by using the DELETE command.

Runtime resource definitions and runtime descriptor definitions created dynamically by using the CREATE command are recoverable across an IMS warm start or an emergency restart. Resources and runtime descriptor definitions are lost if an IMS system is cold started, unless the resource and descriptor definitions were previously exported to a resource definition data set while IMS was running.

The CREATE command can only be issued through the Operations Manager API. The CREATE command can be used in all IMS environments, although some keywords are restricted. For example, you cannot use the CREATE DB command in an IMS DCCTL environment, and you cannot use the CREATE TRAN command in IMS DBCTL environment. The CREATE command is not valid in an XRF alternate IMS system, RSR tracker IMS system, or an FDBR region.

The CREATE command is only valid if DRD is enabled. To enable DRD, specify MODBLKS=DYN in either the DFSDFxxx or the DFSCGxxx PROCLIB member.

## **Creating databases and database descriptors dynamically**

Use the type-2 command CREATE DB to dynamically create a database. Use the type-2 command CREATE DBDESC to create a database descriptor.

You can reference existing databases or database descriptors to use as a model by specifying the LIKE() parameter.

The CREATE DB and CREATE DBDESC commands are valid only in DB/DC and DBCTL systems.

### **Creating application programs and application program descriptors dynamically**

Use the type-2 command CREATE PGM to dynamically create an application program. Use the type-2 command CREATE PGMDESC to create an application program descriptor.

You can reference existing application programs or application program descriptors to use as a model by specifying the LIKE() parameter.

The CREATE PGM and CREATE PGMDESC commands are valid in DB/DC, DBCTL, and DCCTL systems.

In a shared EMH queues environment, CREATE PGM creates a program even if there are messages on the shared EMH queues for the program. If the messages were placed on the shared EMH queues because the program was defined as Fast Path in another IMS system, but the program being created is defined as non Fast-Path, the IMS system in which the new program is being defined will be unable to access the messages on the shared EMH queues.

### **Creating Fast Path routing codes and Fast Path routing code descriptors dynamically**

Use the type-2 command CREATE RTC to dynamically create a Fast Path routing code. Use the type-2 command CREATE RTCDESC to create a Fast Path routing code descriptor.

You can reference existing Fast Path routing codes or Fast Path routing code descriptors to use as a model by specifying the LIKE() parameter.

The CREATE RTC and CREATE RTCDESC commands are valid only in DB/DC and DCCTL systems in which Fast Path is installed.

### **Creating transactions and transaction descriptors dynamically**

Use the type-2 command CREATE TRAN to dynamically create a transaction. Use the type-2 command CREATE TRANDESC to create a transaction descriptor.

You can reference an existing transaction or transaction descriptor to use as a model by specifying the LIKE() parameter.

The CREATE TRAN and CREATE TRANDESC commands are valid only in DB/DC and DCCTL systems.

In a shared message queues environment, the CREATE TRAN command creates a transaction, even if there are messages on the shared message queues for the transaction. If the messages were placed on the shared message queues using conversation, Fast Path, response, or serial attributes that differ from the attributes defined for the new transaction, the IMS system in which the new transaction is being created will be unable to access the messages on the shared message queues.

## Updating runtime resource definitions or runtime descriptor definitions dynamically

Use the UPDATE command with the appropriate parameters to dynamically update a runtime resource definition or resource descriptor. Changes to resource definitions or runtime descriptor definitions made by using the UPDATE command exist until the next IMS cold start, unless they are modified by another UPDATE command or the runtime resource definition or runtime descriptor definition is deleted by using the DELETE command.

Updates to runtime resource definitions and runtime descriptor definitions are recoverable across an IMS warm start or an emergency restart. Updates to resources and runtime descriptor definitions are lost if an IMS system is cold started, unless the resource and descriptor definitions were previously exported to the resource definition data set while IMS was running.

The UPDATE command can be used to update runtime descriptor definitions dynamically without having to go through the online change process. The UPDATE command can only be issued through the Operations Manager API. Each resource or descriptor is updated individually. DRD does not work like the online changes where either all resources are created, updated, and deleted, or none of the resources are created, updated, or deleted.

You cannot update the attributes of runtime resource definitions that are currently in use or that have work in progress. If a resource is in use, the UPDATE command fails. To ensure the success of an UPDATE command, stop the resource and let the work complete before you issue the UPDATE command. In an IMSplex environment, the UPDATE command might succeed on some IMS systems and fail on others.

To minimize the likelihood that an update of a resource definition will fail, perform the following steps before issuing the UPDATE command:

1. Stop the resource
2. Query the resource to check for work in progress
3. Dequeue transactions with the QUEUE command
4. Complete the work (if any)

If all of the attributes specified by an UPDATE command are already defined for a resource, IMS takes no action and returns a completion code of zero.

Except for a few resource-specific exceptions, the UPDATE command is not valid with the SET() parameter or with any of the resource descriptor parameters if DRD is not enabled in either the DFSDFxxx or the DFSCGxxx PROCLIB member.

You cannot modify resource definitions for resources that are delivered with the IMS product. For example, you cannot modify the Fast Path utility program DBF#FPU0. You can, however, use the UPDATE command to change the status of certain IMS-delivered resources. For example, you might start or stop DBF#FPU0.

You also cannot modify the runtime descriptor definitions that are delivered with the IMS product, other than to select them as your default resource descriptor for a given resource. DFSDSDB1, DFSDSPG1, DBFDSRT1, and DFSDSTR1 are examples of the descriptors delivered with the IMS product.

### Updating database attributes and database descriptors dynamically:

Use the type-2 command `UPDATE DB` to update certain attributes of a database definition dynamically. Use the type-2 command `UPDATE DBDESC` to update a database descriptor.

The **DB** and **DBDESC** parameters of the `UPDATE` command are valid only in DB/DC and DBCTL systems. You cannot specify the `UPDATE DB` command or the `UPDATE DBDESC` command in an XRF alternate IMS system, an RSR tracker IMS system, or in an FDBR region.

If you use the `UPDATE DB` command to change the status of a database that has a MODBLKS definition type, the definition type of the database changes to `UPDATE`, as displayed by the `QUERY DB` command.

If DRD is not enabled, you cannot use the `UPDATE` command to update runtime resource definition attribute values.

If a database is in use, you cannot update the status of the database definition.

Changing `RESIDENT(N)` to `RESIDENT(Y)` takes effect at the next restart.

### Updating application programs and application program descriptors dynamically:

Use the type-2 command `UPDATE PGM` to update certain attributes of the definition of an application program dynamically. Use the type-2 command `UPDATE PGMDESC` to update an application program descriptor.

The **PGM** and **PGMDESC** parameters of the `UPDATE` command are valid in DB/DC, DBCTL, and DCCTL systems. You cannot specify the `UPDATE PGM` command or the `UPDATE PGMDESC` command in an XRF alternate IMS system, an RSR tracker IMS system, or in an FDBR region.

If there is work in progress for the application program, you cannot update application program definitions. In a local queues environment, if there are queued messages for transactions associated with the program being updated, the update fails.

In a shared queues environment, the `UPDATE PGM` command updates a program even if there are messages queued for transactions that reference the program. If there are conflicts within a single IMS between the program attributes being changed and the attributes of the transactions and routing codes that reference the program, the update fails. However, if there are conflicts between the program attributes being changed and the attributes of messages queued for transactions that reference the program, this will not cause the update to fail. You must be careful when updating a program in a shared queues environment to set the attributes consistently between programs, transactions, and routing codes across all IMSs in the shared queues group. In a sysplex environment with multiple IMS systems, the update might succeed on some IMS systems and fail on others.

Changing `RESIDENT(N)` to `RESIDENT(Y)` takes effect at the next restart.

You must stop an application program before issuing an `UPDATE PGM` command if you are updating any of the following attributes of the application program definition:

- DOPT
- FP
- GPSB
- LANG
- PGMTYPE
- RESIDENT
- SCHDTYPE

If DRD is not enabled, the LOCK and TRANSTAT attributes are the only application program attributes that you can set by using the UPDATE PGM command.

For application programs delivered as part of the IMS product, you can issue the UPDATE PGM command with the following parameters:

- START
- STOP
- SET(LOCK(ON|OFF))
- SET(TRANSTAT(Y|N))

If you use the UPDATE PGM command to change any of the following attributes of an application program that has a MODBLKS definition type, the definition type of the application program changes to UPDATE, as displayed by the QUERY PGM command:

- DOPT
- FP
- GPSB
- LANG
- PGMTYPE
- SCHDTYPE
- TRANSTAT

You cannot update any of the following attributes of an application program if DRD is not enabled:

- DOPT
- FP
- GPSB
- LANG
- PGMTYPE
- RESIDENT
- SCHDTYPE

#### **Updating Fast Path routing codes and Fast Path routing code descriptors dynamically:**

Use the type-2 command UPDATE RTC to dynamically update the definition of a Fast Path routing code. Use the type-2 command UPDATE RTCDESC to update a Fast Path routing code descriptor.

The **RTC** and **RTCDESC** parameters of the CREATE command are valid only in DB/DC and DCCTL systems in which Fast Path is installed.

If the Fast Path routing code is in use, the UPDATE RTC command fails if you are modifying either the INQ attribute or the PGM attribute of the definition of a Fast Path routing code.

You cannot update the following attributes of the definition of a Fast Path routing code if DRD is not enabled:

- INQ
- PGM

If you use the UPDATE RTC command to change either the INQ attribute or the PGM attribute in the definition of a Fast Path routing code that has a MODBLKS definition type, the definition type of the Fast Path routing code changes to UPDATE, as displayed by the QUERY RTC command.

### **Updating transactions and transaction descriptors dynamically:**

Use the type-2 command UPDATE TRAN to dynamically update the definition of a transaction. Use the type-2 command UPDATE TRANDESC to update a transaction descriptor.

The **TRAN** and **TRANDESC** parameters of the UPDATE command are valid only in DB/DC and DCCTL systems.

If DRD is not enabled, you cannot use the UPDATE command to update definitional attributes. You can, however, use the UPDATE command when DRD is not enabled to control the locking, tracing, queuing, and scheduling of transactions.

In a local queues environment, you cannot update the definitional attributes of a transaction if there are queued messages for the transaction being updated. In a shared queues environment, the UPDATE TRAN command updates the definitional attributes of a transaction, even if there are messages queued for the transaction on the shared queue.

You must be careful when updating the definitional attributes of a transaction to set attributes consistently across all IMSs in a shared queues group. Updating a transaction does not remove queued messages that are associated with the transaction from the queue. Also, queued messages are not automatically updated to be consistent with the changes made to associated transactions. Inconsistencies between queued messages and updated transactions can cause problems. In a sysplex environment with multiple IMS systems, an update might succeed on some IMS systems and fail on others.

If a transaction is in use or has work in progress, the UPDATE TRAN command fails if you are modifying any of the following attributes of the transaction definition:

- AOCMD
- CMTMODE
- CONV
- DCLWA
- DIRROUTE
- EDITRTN
- EDITUC
- EMHBSZ

- EXPRTIME
- FP
- INQ
- MSGTYPE
- MSNAME
- PGM
- PLCTTIME
- RECOVER
- REMOTE
- RESP
- SERIAL
- SIDL
- SIDR
- SPASZ
- SPATRUNC
- TRANSTAT
- WFI

If you use the UPDATE TRAN command to change any of the following attributes in the definition of a transaction that has a MODBLKS definition type, the definition type of the transaction changes to UPDATE, as displayed by the QUERY TRAN command:

- AOCMD
- CMTMODE
- CONV
- DCLWA
- DIRROUTE
- EDITRTN
- EDITUC
- EMHBSZ
- FP
- INQ
- MSGTYPE
- MSNAME
- PGM
- RECOVER
- RESP
- SERIAL
- SIDL
- SIDR
- SPASZ
- SPATRUNC
- WFI



## Deleting runtime resource definitions or runtime descriptor definitions dynamically

Use the DELETE command to dynamically delete runtime resource definitions such as databases, application programs, Fast Path routing codes, or transactions. Also use the DELETE command to dynamically delete any runtime resource definition descriptors.

If you delete resource definitions or runtime descriptor definitions from all IMSs in an IMSplex, and the IMSRSC repository is used, you must also issue a DELETE DEFN command to delete the resource definitions or runtime descriptor definitions from the repository.

**Recommendation:** Delete the resource definitions or runtime descriptor definitions from all IMSs first, and then delete them from the repository.

Deleted runtime resource definitions and runtime descriptor definitions reappear if an IMS system is cold started, unless all resource and descriptor definitions are exported to a resource definition data set before IMS terminates and are imported from the resource definition data set when IMS cold starts. Deleted resources and descriptors remain deleted across either a warm start or an emergency restart. The DELETE command is recoverable.

In a shared message queues environment, DELETE TRAN dynamically deletes the transaction even if there are messages on the shared MSG queues for the transaction. If the transaction is deleted from all IMS systems, no IMS can access the messages on the shared MSG queues. Similarly, in a shared EMH queues environment, DELETE PGM deletes an application program even if there are messages on the shared EMH queues for the program. If the program is deleted from all IMS systems, no IMS system can access the messages on the shared EMH queues.

The DELETE commands are similar to local online change (/MODIFY) or global online change (INITIATE OLC) for IMS resources in the MODBLKS data set, except that the DELETE command works without requiring any offline procedures such as a MODBLKS system definition or the execution of the online change copy utility. Moreover, the DELETE command deletes resources and runtime descriptor definitions individually, while the /MODIFY and INITIATE OLC commands either delete a list of resources as a group or deletes no resources at all if an error occurs.

You cannot delete resources that are currently in use or that have work in progress. If a resource is in use, the delete fails. In a sysplex environment, the delete might succeed on some IMS systems and fail on others.

To minimize the likelihood that a delete of a resource will fail, perform the following steps before issuing the DELETE command:

1. Stop the resource
2. Query the resource to check for work in progress
3. Dequeue transactions with the QUEUE command
4. Complete the work (if any)

Examples of resources in use include a database being accessed by an application program, a /STOP or UPDATE command in progress for the resource, a scheduled application program, an active routing code, a transaction with messages queued.

You cannot issue the DELETE command against resources or runtime descriptor definitions that are identified with a definition type of IMS and delivered with the IMS product. For example, you cannot delete the IMS-defined application program DBF#FPU0 or the IMS-defined runtime descriptor definitions DFSDSDB1, DFSDSPG1, DBFDSRT1, and DFSDSTR1.

In a DB/DC environment, all of the parameters of the DELETE command are valid.

In a DBCTL environment, only the following forms of the DELETE command are valid:

- DELETE DB
- DELETE DBDESC
- DELETE DEFN TYPE(DB | DBDESC | PGM | PGMDESC)
- DELETE PGM
- DELETE PGMDESC

In a DCCTL environment, only the following forms of the DELETE command are valid:

- DELETE DEFN TYPE(PGM | PGMDESC | RTC | RTCDESC | TRAN | TRANDESC)
- DELETE PGM
- DELETE PGMDESC
- DELETE RTC
- DELETE RTCDESC
- DELETE TRAN
- DELETE TRANDESC

The commands DELETE RTC and DELETE RTCDESC are only valid in DB/DC or DCCTL environments in which Fast Path is installed.

### **Querying runtime resource definitions or runtime descriptor definitions dynamically**

Use the QUERY command with the appropriate parameter to query database and program resources. All of the QUERY commands are type-2 commands that you can issue from the OM API.

These commands can be issued through the TSO SPOC or the Manage Resources options in the IMS Applications menu. These commands can also be issued to an IMSplex using the batch SPOC utility.

The QUERY commands return information based on the keyword specified. For example, issue the QUERY DB command to view the current database resource or database descriptor settings.

Several type-2 QUERY commands support dynamic resource definition (DRD):

- QUERY DB
- QUERY DBDESC
- QUERY OLC
- QUERY PGM
- QUERY PGMDESC
- QUERY RTC

- QUERY RTCDESC
- QUERY TRAN
- QUERY TRANDESC

## Modifying ETO user IDs and assignments of ISC users

You can use the /ASSIGN, /DISPLAY USER DEADQ, /DEQUEUE, and the /DISPLAY QCNT MSGAGE commands to modify various aspects of assignments of ISC users.

For dynamic users, use the /ASSIGN command to change the assignment of an LTERM to another user (also called a subpool for ISC). The new assignment remains in effect until the next cold start or until you issue another /ASSIGN command. For static terminals that are not ISC, the LTERM can be reassigned to another terminal or node.

Use the /DISPLAY USER DEADQ command to list all message queues that are eligible for dead letter status. Use the /ASSIGN command to assign a dead letter queue to another user. Use the /DEQUEUE command to discard a dead letter queue.

In a shared-queues environment, use the /DISPLAY QCNT MSGAGE command to determine which messages, if any, are eligible for dead letter status.

## Modifying Multiple Systems Coupling resources

After modifying MSC resources, use the /MSVERIFY command to ensure that the assignment produced a valid configuration. The /MSVERIFY command verifies the consistency of MSC system identifications (SYSIDs) and logical link paths (MSNAMEs) across two systems.

All changes made by the online commands remain in effect until the next cold start or until you issue another online command.

### Related concepts:

“MSC operations” on page 267

## Modifying security options

Commands to modify security options for RACF® are described.

Use the /MODIFY PREPARE RACF, /DISPLAY MODIFY, and /MODIFY COMMIT commands to reinitialize RACF information if you are not using a RACF data space. If you are using a RACF data space, use the RACF SETROPTS RACLIST command rather than the IMS /MODIFY command.

Use the /SECURE APPC command to control the RACF security level for input from LU 6.2 devices. Use the /DISPLAY APPC command to show the security level that is currently in effect. When IMS starts, the default is full security.

Use the /SECURE OTMA command to control the RACF security level for input from specified OTMA clients or for the entire OTMA z/OS cross-system coupling facility (XCF) group. Use the /DISPLAY OTMA command to show the security level that is currently in effect. When IMS starts, the default is full security.

Use the /SECURE OTMA REFRESH command to refresh security information cached in online memory after making a RACF change to a specific user ID.

Cached OTMA security information can be refreshed globally by the OTMA TMEMBER, or by user ID, without impacting other cached user IDs.

## Displaying and terminating conversations

Use the /DISPLAY CONV command to show the status of all conversations, held or active. You can terminate a conversation if necessary with the /EXIT command, but you should only do this after warning the end user.

## Modifying and controlling subsystems

Use the /CHANGE command to delete an invalid network identifier (NID). If you need to disconnect from a specific subsystem, use the /STOP command. If the /STOP command does not work, use the z/OS MODIFY command.

## Controlling OTMA input messages

Use the OTMA descriptor, the /START TMEMBER INPUT command, or a client-bid protocol message from an OTMA member to set the maximum transaction instance block (TIB) or maximum active input message count for an OTMA member.

The OTMA descriptor can be used to establish the initial maximum active input message count for an OTMA member. The client-bid protocol message can under-ride the value that is set by the OTMA descriptor, but not override it. The /START TMEMBER INPUT command can always be used to override the value that is specified by the OTMA descriptor or client-bid protocol message. If the value is not specified by any of these methods, the system default of 5000 TIBs is used.

OTMA monitors the growth of the active input messages from members. A warning message DFS1988W is sent to the console to indicate that a certain percentage of the TIB limit for the member has been reached. The message is sent when the member reaches 80% of the TIB limit and then for each additional 5% thereafter. When the maximum limit is reached, an error message DFS1989E is sent to the console. Any subsequent OTMA input messages are rejected with a new OTMA sense code X'30'.

Use the /DISPLAY TMEMBER or /DIS OTMA command to display the active input message count for an OTMA member.

## Recovery during the IMSRSC repository data set update process

The Repository Server (RS) uses a duplex pair of data sets for each IMSRSC repository, so that the RS can always recover data sets to the last completed and verified write activity.

The duplex pair of data sets is the primary repository index data set (RID) and repository member data set (RMD) (COPY1) and the secondary RID and RMD (COPY2). The RS always requires both the primary and secondary data set pairs to be available; the repository is otherwise stopped.

The repository can have an optional third defined data set pair, the spare RID and RMD (SPARE), which is used during the SPARE recovery process.

The RS writes to the repository in two phases:

1. In the first phase, the COPY1 RID and RMD are updated. At the crossover point, the change is considered completed. A successful return code is sent to the requesting client.
2. In the second phase, the same updates are written to the COPY2 RID and RMD. The request is completed.

The following sequence shows how the RS performs recovery of data sets if an I/O error occurs during the first phase of the update process (write to the COPY1 data set):

1. The request in progress fails with an error.
2. The repository becomes unavailable. The client is notified through any client connection exits with the Repository unavailable status.
3. If the RS verifies that a spare data set pair is unavailable, the repository is stopped. The client is notified that the repository is stopped. Administrator intervention is required to delete the data set in error and define a new data set and restart the repository. During repository start processing, data is copied from the valid data set to the newly added data set.
4. If the RS verifies that a valid spare data set pair is available and the other conditions for recovery are met, then spare recovery processing is started:
  - a. Any client connection exits are driven with the Repository recovery started status.
  - b. When recovery processing is completed, data is copied from the COPY2 data set to the SPARE data set pair and the status of the SPARE data set pair is changed to COPY1. The repository becomes available, and any client connection exits are driven with the Repository recovery ended successfully status.
  - c. If an error occurs during recovery processing, any client connection exits are driven with the Repository recovery error status.

If the RS fails during the second phase of the update process (write to the COPY2 data set), the RS performs the same tasks, except that data is copied from the COPY1 data set to the SPARE data set pair and the status of the SPARE data set pair is changed to COPY2. As the first phase of the update process is completed, the request in progress is considered completed. A successful return code is returned to the caller.

If the automated SPARE recovery process fails, the repository might be in a CLOSE state.

If a repository remains in CLOSED state at the end of SPARE recovery processing, perform the following steps to make the repository available:

1. Re-allocate the discarded repository data sets.
2. Issue the LIST FRPBATCH command to check if the repository is in a STOP state. If it is not, issue the STOP FRPBATCH command to change the state to STOP.
3. Issue the DSCHANGE FRPBATCH command to change the re-allocated discarded data sets back to a SPARE state.
4. Issue the START FRPBATCH command to start the repository.  
This resumes the data set recovery process.

For automated recovery to take place, either the primary RID and RMD or the secondary RID and RMD must be valid and a SPARE repository data set pair must be available.

If a SPARE repository data set pair is not defined when the failure occurs, the repository is stopped. You must define new data sets for the primary or secondary data sets and start the repository. Perform the following procedure:

1. Allocate new repository data sets to take the place of the failed primary or secondary data sets and to create a SPARE repository data set pair.
  2. Issue the DSCCHANGE FRPBATCH command to discard the repository data sets in error.
  3. Issue the UPDATE FRPBATCH command to replace the failed data sets with the newly allocated data sets for the existing repository.
  4. Issue the START FRPBATCH command to start the repository.
- This resumes the data set recovery process.

You can also use the F reposervername,ADMIN command equivalents to the FRPBATCH commands.

The following topics describe the recovery procedures after the failure of both the primary and secondary RIDs and RMDs.

**Related concepts:**

- ➞ Overview of the IMSRSC repository (System Definition)
- ➞ IMSRSC repository and RS catalog repository data sets (System Definition)
- ➞ Recovery of changed runtime resource and descriptor definitions (System Definition)
- ➞ Deleting runtime resource and descriptor definitions (System Definition)
- ➞ Updating IMSRSC repository specifications in the RS catalog repository (System Administration)

**Related reference:**

- ➞ F reposervername,ADMIN (Commands)

**Related information:**

- ➞ Commands for FRPBATCH (System Programming APIs)

## **Recovering IMSRSC repository data sets when all IMS systems and an RM system are active in an IMSplex**

You can recover IMSRSC repository data sets when all IMS systems and an RM system are active in an IMSplex.

Perform the following procedure:

1. Define new primary and secondary repository data sets for the repository that failed, by using the FRPBATCH UPDATE command.
  2. Start the repository by using the FRPBATCH START command or the F reposervername,ADMIN START command.
  3. List the repository by using the FRPBATCH LIST command or the F reposervername,ADMIN DISPLAY command.
  4. Write the resource definitions to the repository from each active IMS by using the IMS EXPORT command.
    - If your IMS systems are cloned, route the following command to one of the active IMS systems in the IMSplex and have it write its resource definitions to the repository for all of the IMS systems in the IMSplex.
- ```
EXPORT DEFN TARGET(REPO) TYPE(ALL) NAME(*) SET(IMSID
(IMSids_of_all_IMSs_sharing_repository))
```

- If your IMS systems are not cloned, route the following command to each of the IMS systems in the IMSplex so that each IMS writes its own runtime definitions to the repository.

```
EXPORT DEFN TARGET(REPO) TYPE(ALL) NAME(*)
```

This command writes the runtime resource definitions from each IMS to the repository.

IMS can continue processing while the repository data sets are being recovered.

Related concepts:

➞ Overview of the IMSRSC repository (System Definition)

Related reference:

➞ EXPORT command (Commands)

➞ F reposervername,ADMIN (Commands)

➞ LIST command for FRPBATCH (System Programming APIs)

➞ START command for FRPBATCH (System Programming APIs)

➞ UPDATE command for FRPBATCH (System Programming APIs)

Recovering IMSRSC repository data sets when one or more IMS systems are down but an RM system is active

You can recover IMSRSC repository data sets when one or more IMS systems are down but IMS logs are available and an RM system is active.

Perform the following procedure:

1. Define new primary and secondary repository data sets for a repository that failed, by using the FRPBATCH UPDATE command.
2. Start the repository by using the FRPBATCH START command or the F *reposervername*,ADMIN START command.
3. List the repository information by using the FRPBATCH LIST command or the F *reposervername*,ADMIN DISPLAY command.
4. Write the resource definitions to the repository by performing one of the following steps:
 - If your IMS systems are cloned, route the following command to one of the active IMS systems in the IMSplex and have it write its resource definitions to the repository for all of the IMS systems in the IMSplex.

```
EXPORT DEFN TARGET(REPO) TYPE(ALL) NAME(*) SET(IMSID  
(IMSids_of_all_IMSs_sharing_repository))
```
 - If your IMS systems are not cloned, take one of the following steps:
 - Issue the EXPORT DEFN TARGET(REPO) TYPE(ALL) NAME(*) command at each active IMS system to route the command to each active IMS system in the IMSplex. This command writes the runtime resource definitions from each active IMS to the repository.
 - Generate a resource definition data set (RDDS) for each IMS system that is down. The RDDS can be generated from the IMS logs by using the DFSURCL0 utility. Then, run the CSLRUP10 utility to write the resource definitions from each generated RDDS to the repository.

Related concepts:

➡ Overview of the IMSRSC repository (System Definition)

Related reference:

➡ EXPORT command (Commands)

➡ Repository to RDDS utility (CSLURP20) (System Utilities)

➡ Create RDDS from Log Records utility (DFSURCL0) (System Utilities)

➡ F reposervername,ADMIN (Commands)

➡ LIST command for FRPBATCH (System Programming APIs)

➡ START command for FRPBATCH (System Programming APIs)

➡ UPDATE command for FRPBATCH (System Programming APIs)

Recovering IMSRSC repository data sets when no IMS systems or RM systems are active

You can recover IMSRSC repository data sets when all IMS systems and RM systems are down and no IMS log is available but backup copies of the repository data sets are available.

Perform the following procedure:

1. Define the primary and secondary data sets from the backup copies.
2. Start the repository by using the FRPBATCH START command or the F *reposervername*,ADMIN START command. The repository is started from the backup copies of data.
3. List the repository information by using the FRPBATCH LIST command or the F *reposervername*,ADMIN DISPLAY command.
4. Cold start IMS systems from the repository.
5. To bring the repository and IMS systems to the same state as the time of failure, reissue the commands that you used to create, update, or delete the resource definitions between the time the backup was generated and the IMS cold start. The type-1 or type-2 commands that were issued from OM API can be obtained from the OM audit log or command responses sent to OM API.

Related concepts:

➡ Overview of the IMSRSC repository (System Definition)

Related tasks:

“Cold starting an IMS system that uses the IMSRSC repository” on page 76

Related reference:

➡ F reposervername,ADMIN (Commands)

➡ LIST command for FRPBATCH (System Programming APIs)

➡ START command for FRPBATCH (System Programming APIs)

➡ UPDATE command for FRPBATCH (System Programming APIs)

Controlling log data set characteristics

You might need to tune and modify log data set characteristics after monitoring or after changing your requirements for system availability, integrity, or operator handling. You can change characteristics for the OLDS, WADS, and RECON data sets through the IMS commands.

Changing online log data set characteristics

Because you can restart IMS (warm start or emergency restart) with all input on SLDS, you can reallocate the OLDSs between a shutdown (or failure) and a subsequent restart. To restart IMS using SLDSs as input, you must delete the PRIOLDS and SECOLDS records from the RECON data sets.

Changing block size for the OLDS

You can change the space, location, or block size of your OLDS or reallocate an OLDS on the same volume and with the same space.

1. Shut down IMS.
2. Archive all OLDSs.
3. Delete PRIOLDS and SECOLDS records from the RECON data sets by using the DELETE.LOG command of the Data Base Recovery Control utility program (DSPURX00).
4. Either scratch all OLDSs and reallocate them with the new BLKSIZE, or change the BLKSZ= parameter in the DFSVSMxx PROCLIB member to the new block size.
5. Verify WADS space allocation.
6. Restart IMS (from SLDS).

Changing the OLDS block size can affect the space required for the IMS write-ahead data set (WADS) . For example, if you have your WADS sized to hold a hundred 22K buffers, and you increase the OLDS block size to 24K, you might need to increase the size of the WADS if you want it to be able to hold a hundred of the larger 24K buffers. If you reallocate the WADS, make sure that you restart IMS by using the /NRE FORMAT WA or /NRE FORMAT ALL command.

Related tasks:

“Changing write-ahead data set characteristics” on page 43

Related reference:

 Execution data sets (Installation)

Changing the OLDS mode from single to dual

You must change your OLDSDEF specification in the DFSVSMxx member in IMS.PROCLIB to change the OLDS mode from single to dual. IMS initialization requires that at least three pairs of OLDSs be available. You must also reconsider data set placement.

When changing from single to dual OLDS, each data set in a pair of OLDSs must have the same space allocation (number of blocks).

Changing the mode from single to dual or from dual to single requires changes in the following operating procedures:

- Skeletal JCL for archive (ARCHJCL member)
- Skeletal JCL for log recovery (LOGCLJCL member)
- Batch JCL for Log Recovery utility
- Batch backout JCL for online transactions and BMPs

If you warm start the IMS subsystem after changing from single to dual OLDSs, the /DISPLAY OLDS command does not show the secondary OLDSs as IN USE until they have been archived once. The command does, however, show dual OLDS logging.

1. Shut down IMS.

2. Archive all OLDSs.
3. Allocate dual OLDSs.
4. Delete the OLDS records from the RECON data sets, using the DELETE.LOG command. The primary OLDS records will be deleted.
5. Change OLDSDEF specification to dual.
6. Change IMS startup procedure (OLDS DD statements), if required.
7. Compile DFSMDA macros, if required.
8. Modify operating procedures.
9. Restart IMS.

Changing the OLDS mode from dual to single

You must change your OLDSDEF specification in the DFSVSMxx member in IMS.PROCLIB to change the OLDS mode from dual to single. IMS initialization requires that at least three pairs of OLDSs be available. You must also reconsider data set placement.

Changing the mode from single to dual or from dual to single requires changes in the following operating procedures:

- Skeletal JCL for archive (ARCHJCL member)
- Skeletal JCL for log recovery (LOGCLJCL member)
- Batch JCL for Log Recovery utility
- Batch backout JCL for online transactions and BMPs

If you warm start the IMS subsystem after changing from single to dual OLDSs, the /DISPLAY OLDS command does not show the secondary OLDSs as IN USE until they have been archived once. The command does, however, show dual OLDS logging.

1. Shut down IMS.
2. Archive all OLDSs.
3. Allocate dual OLDSs.
4. Delete the OLDS records from the RECON data sets, using the DELETE.LOG command. The primary OLDS records will be deleted.
5. Change OLDSDEF specification to dual.
6. Change IMS startup procedure (OLDS DD statements), if required.
7. Compile DFSMDA macros, if required.
8. Modify operating procedures.
9. Restart IMS.

Changing the number of OLDS buffers

You must change BUFNO by changing the OLDSDEF specification for BUFNO in the DFSVSMxx member in IMS.PROCLIB to change the number of OLDS buffers.

When you modify BUFNO, you should consider also modifying the region size for the VSAM common storage area (CSA). The amount of storage fixed for OLDS buffers is [BUFNO x BUFFERSIZE].

WADS space is also affected by BUFNO.

To change the number of OLDS buffers:

1. Shut down IMS.
2. Change the OLDSDEF specification for BUFNO.

3. Verify CSA size.
4. Verify WADS space allocation.
5. Restart IMS.

Related tasks:

“Changing write-ahead data set characteristics”

Related reference:

 Execution data sets (Installation)

Changing space, location, or allocation of an OLDS without shutting down IMS

You can change the space, location, or allocation of an OLDS by change the block size for the OLDS. If your IMS is operating in a non-XRF environment, you can modify space, location, or allocation without shutting down IMS.

1. /STOP OLDS nn
2. Archive all OLDSs.
3. Delete PRIOLDS and SECOLDS records in the RECON data sets, using the DELETE.LOG command.
4. Scratch and reallocate OLDSs.
5. /START OLDS nn

Related concepts:

 Formatting newly initialized (reinitialized) volumes for an OLDS (System Definition)

Changing write-ahead data set characteristics

The write-ahead data set (WADS) contains log records that reflect completed operations and are not yet written to an online log data set. You can change the write-ahead data set (WADS) mode, add or remove a WADS, or change the space, location, or allocation of the WADS.

Related tasks:

“Changing block size for the OLDS” on page 41

“Changing the number of OLDS buffers” on page 42

Changing the WADS mode from single to dual

You can change the write-ahead data set (WADS) mode from single to dual. The WADS can be dynamically allocated and deallocated.

To reflect the new mode for WADSs, you must update the following parameters:

- Skeletal JCL for the Log Recovery utility (LOGCLJCL member).
 - IMS startup procedure (WADS= execution parameter).
 - All recovery procedures implemented to recover WADS errors or to close unclosed OLDSs using WADS.
1. Shut down IMS.
 2. Allocate new WADS.
 3. Define a DFSMDA member.
 4. Add DD statement in IMS JCL, if necessary.
 5. Code WADS=D in IMS JCL.
 6. Modify operating procedures.
 7. Restart IMS with FORMAT WADS keywords.

Changing the WADS mode from dual to single

You can change the write-ahead data set (WADS) mode from dual to single. The WADS can be dynamically allocated and deallocated.

To reflect the new mode for WADSs, you must update the following parameters:

- Skeletal JCL for the Log Recovery utility (LOGCLJCL member).
 - IMS startup procedure (WADS= execution parameter).
 - All recovery procedures implemented to recover WADS errors or to close unclosed OLDSs using WADS.
1. Shut down IMS.
 2. Code WADS=S in IMS JCL.
 3. Delete DFSMDA member.
 4. Remove DD statement in IMS JCL, if necessary.
 5. Modify operating procedures.
 6. Restart IMS.

Adding a spare WADS

The write-ahead data set (WADS) contains log records that reflect completed operations and are not yet written to an online log data set.

1. Allocate a spare WADS.
2. Update WADSDEF specification in the DFSVSMxx member of IMS.PROCLIB.
3. Define DFSMDA member.
4. Add DD statement in IMS JCL, if necessary.
5. Modify operating procedures.
6. /START WADS*n* (or wait until IMS restart).

Removing a spare WADS

You might need to remove a spare write-ahead data set (WADS). WADS contains log records that reflect completed operations and are not yet written to an online log data set.

1. Issue the /STOP WADS*n* command (and wait for dynamic deallocation).
2. Scratch the spare WADS.
3. Update WADSDEF statement in the DFSVSMxx member of IMS.PROCLIB.
4. Remove DD statement in IMS JCL.
5. Modify operating procedures.

Changing space, location, or allocation of a WADS

All write-ahead data sets (WADS) must have the same space allocation (number of tracks) and be on the same type of device.

1. Shut down IMS.
2. Scratch and reallocate WADSs.
3. Restart IMS with FORMAT WADS keywords.

Changes in the WADS block size after IMS initialization are accomplished by changing the OLDS block size. If you are changing the OLDS block size from being a multiple of 2048 (running in non-z/Architecture mode) to being a multiple of 4096 (running in z/Architecture® mode), make sure that you run the /NRE FORMAT WA or /NRE FORMAT ALL command after changing the OLDS block size.

Changing system log data set characteristics

Converting from single to dual SLDSs requires modification in the skeletal JCL (ARCHJCL member) and in all your operational procedures using SLDSs. In the operational procedures, consider using the secondary SLDS when you experience errors in the primary one. No modification is required for online processing because IMS dynamically allocates SLDSs.

Changing the BLKSIZE requires modification in the skeletal JCL (ARCHJCL member). All SLDSs required for online processing must have the same BLKSIZE.

Changing RECON data set characteristics

You can change many aspects of the RECON data set by adding or removing spares, replacing the active data set, or changing it from single to dual mode.

Adding a spare RECON data set

IMS normally works with two active RECON data sets. If one RECON data set becomes unavailable, a spare RECON data set will be activated, if a spare is available.

Recommendation: For both online and batch, use dynamic allocation for RECON data sets, and run with at least three RECON data sets.

Note: The spare data set must be in VSAM CREATE mode.

1. Delete the cluster.
2. Compile DFSMDA macro for the spare data set or add DD statement in IMS JCL and batch JCL if you do not use dynamic allocation.

Removing a spare RECON data set

There may be certain instances where you need to remove a RECON data set (for instance testing in a nonproduction system or maybe a damaged DASD).

Recommendation: For both online and batch, use dynamic allocation for RECON data sets, and run with at least three RECON data sets.

1. Delete the cluster.
2. Delete DFSMDA member in IMS.SDFSRESL or remove DD statement from IMS JCL and batch JCL if you do not use dynamic allocation.

Replacing an active RECON data set

There might be circumstances for which you need to replace an active RECON data set, for instance on a test system.

Recommendations: Stop all IMS subsystems and batch jobs accessing the RECON data sets.

For both online and batch, use dynamic allocation for RECON data sets, and run with at least three RECON data sets.

1. Define a spare data set with new space or allocate a spare data set at a new location.
2. Issue the CHANGE.RECON REPLACE (RECON n) command.
3. Define a new spare data set.
4. Continue normal processing.

Changing the RECON data set mode from single to dual

IMS normally works with two active RECON data sets. If one RECON data set becomes unavailable, a spare will be activated if a spare is available.

Recommendations: Stop all IMS subsystems and batch jobs accessing the RECON data sets.

For both online and batch, use dynamic allocation for RECON data sets, and run with at least three RECON data sets.

1. Define a spare RECON data set.
2. Issue the `CHANGE.RECON DUAL` command.
3. Define a new spare data set.
4. Continue normal processing.

Connecting and disconnecting subsystems

IMS can only connect to another subsystem if that subsystem is identified in the subsystem member in `IMS.PROCLIB`. You must specify the subsystem member name in the `SSM EXEC` parameter before an IMS subsystem can access databases in an external subsystem (another program executing in a z/OS address space).

When specified to and accessed by IMS, the subsystem member name cannot be changed without stopping IMS.

Connections between an IMS subsystem and another subsystem can occur with or without operator intervention, and without regard to which subsystem is available first.

- Automatic connection: When the `SSM EXEC` parameter is specified, IMS automatically establishes the connection when it processes the parameter.
- Operator-controlled connection: When the `SSM EXEC` parameter is not specified, the connection is established when the `/START SUBSYS SSM` command is entered. The `/START SUBSYS SSM` command specifies the subsystem member in `IMS.PROCLIB` that IMS uses in connecting to subsystems.

Disconnecting IMS from another subsystem is initiated with the `/STOP SUBSYS` command, and normally completes without operator intervention. IMS will quiesce all dependent region external subsystem activity prior to stopping the subsystem.

If an external subsystem fails, IMS might not be able to disconnect properly from the failed subsystem. In this case, IMS might appear to hang in the disconnect process and the external subsystem might not be able to reconnect to IMS until the threads that were open at the time of the failure are properly closed.

To prevent this problem, issue the `/STOP SUBSYS` command before the failing external subsystem is shut down. After the external subsystem is brought back up, reconnecting to IMS is automatic.

If it is not possible to issue the `/STOP SUBSYS` command before the external subsystem is shut down, issue the `/DISPLAY ACTIVE` or `/DISPLAY ALL` command. These commands enable you to display the external subsystem threads that are still connected to IMS. Each thread has to be canceled individually. After all threads are cancelled and the external subsystem is brought back up, reconnecting to IMS is automatic.

Related concepts:

 Specifying the SSM= EXEC parameter (System Definition)

Commands for IMS operations tasks

You can use various commands to complete IMS operation tasks.

The following table describes commands for IMS operation tasks.

Table 15. Commands for IMS operations tasks

Operations task	Command
Activating VTAM nodes	/ACTIVATE
Activating MSC links	/ACTIVATE
Allocating a conversation with an LU name and TP name	/ALLOCATE
Altering assignments for IMS resources	/MSASSIGN
Altering relationships between IMS resources	/ASSIGN
Assigning a new logical link	/MSASSIGN MSNAME <i>name</i> LINK <i>link</i> UPDATE MSNAME SET(MSLINK(<i>name</i>))
Assigning a new physical link	/MSASSIGN LINK <i>link</i> MSPLINK <i>name</i> UPDATE MSLINK SET(MSPLINK(<i>name</i>))
Assigning a user structure to a nonexistent node	/OPNDST NODE x USER y
Beginning HALDB Online Reorganization for one or more partitions	INITIATE OLREORG /INITIATE OLREORG
Canceling an input message	/CANCEL
Canceling an output message	/DEQUEUE QUEUE
Changing applications	UPDATE PGM
Changing database buffers	UPDATE POOL
Changing databases	UPDATE DB
Changing default mode table name	UPDATE MSPLINK SET(MODETBL(<i>name</i>)) /CHANGE LINK <i>link</i> MODE <i>name</i> UPDATE MSLINK SET(MODETBL(<i>name</i>))
Changing HALDB Online Reorganization impact on overall system performance, for one or more partitions	UPDATE OLREORG SET(RATE(<i>rate</i>)) /UPDATE OLREORG SET(RATE(<i>rate</i>))
Changing IMS Fast DB Recovery surveillance	/CHANGE
Changing IMS resources	/CHANGE UPDATE
Changing Fast Path routing code definitions	UPDATE RTC
Changing local system identification (SIDs)	UPDATE MSNAME SET(SIDL(<i>id</i>))
Changing logical link name	UPDATE MSLINK SET(MSLINK(<i>name</i>))
Changing partner ID	UPDATE MSLINK SET(PARTNER(<i>id</i>))
Changing physical link name	UPDATE MSPLINK SET(MSPLINK(<i>name</i>))

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Changing session resynchronization option	/CHANGE LINK <i>link</i> FORCSESS SYNCSESS COLDSESS UPDATE MSLINK SET(SYNCOPT(FORCESS SYNCSESS COLDSESS))
Changing transactions	UPDATE TRAN
Changing the RECON data set	/RMCHANGE
Changing the ready state for a terminal component	/COMPT /RCOMPT
Changing remote system identification SIDs	UPDATE MSNAME SET(SIDR(<i>id</i>))
Changing send or receive buffer size	UPDATE MSLINK SET(BUFSIZE(<i>size</i>))
Changing VTAM node name	UPDATE MSPLINK SET(NODE(<i>name</i>))
Checkpoint, taking	/CHECKPOINT
Class, assigning	/ASSIGN UPDATE TRAN SET(CLASS)
Cold start of components of IMS	/ERESTART
Cold start of IMS	/NRESTART
Connecting to a restarted IRLM	F job,RECONNECT
coupling facility structures, listing status	/CQQUERY
CQS checkpoint, initiating	/CQCHKPT
CQS structure checkpoint, requesting	/CQSET
Creating IMS resources online (application programs, databases, Fast Path routing codes, or transactions)	CREATE
Creating records in the RECON data set	/RMINIT
DBRC, commands for	/RMxxxxxx
Deallocating a user for an ISC node	/QUIESCE
Defining IMS resources online (application programs, databases, Fast Path routing codes, or transactions)	CREATE
Delivering asynchronous output from an LU 6.2 device	/ALLOCATE
Deleting affinities from VTAM	/CHECKPOINT
Deleting IMS resources	/CHANGE /CHECKPOINT DELETE
Deleting information in the RECON data set	/RMDELETE
Destination of messages, setting	/SET
Directing output to a component of a PTERM	/ASSIGN
Disabling logons to this physical link	/PSTOP MSPLINK <i>name</i> UPDATE MSPLINK STOP(LOGON)
Disconnecting a VTAM terminal	/CLSDST /RCLSDST
Displaying an LTERM	/DISPLAY LTERM <i>ltermname</i> QUERY LTERM NAME(<i>ltermname</i>)

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Displaying a node	/DISPLAY NODE <i>nodename</i> QUERY NODE NAME(<i>nodename</i>)
Displaying a user	/DISPLAY USER <i>username</i> QUERY USER NAME(<i>username</i>)
Displaying a userid	/DISPLAY USER <i>useridname</i> QUERY USERID NAME(<i>useridname</i>)
Displaying assigned MSLINKS	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(MSLINK) QUERY MSNAME SHOW(MSLINK)
Displaying assigned MSPLINKs	/DISPLAY ASMT LINK <i>link</i> MODE QUERY MSLINK SHOW(MSPLINK) QUERY MSNAME SHOW(MSPLINK)
Displaying assigned MSNAMEs	/DISPLAY ASMT LINK <i>link</i> QUERY MSLINK SHOW(MSNAME) QUERY MSPLINK SHOW(MSNAME)
Displaying bandwidth mode	/DISPLAY LINK OPTION BUFSIZE QUERY MSLINK SHOW(BANDWIDTH)
Displaying buffer size	/DISPLAY LINK OPTION BUFSIZE QUERY MSLINK SHOW(BUFSIZE)
Displaying CTC address	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(ADDR)
Displaying database buffer pool usage statistics	QUERY POOL /DISPLAY POOL DBAS
Displaying global queue counts	/DISPLAY LINK <i>link</i> QCNT QUERY MSNAME SHOW(QCNT)
Displaying HALDB Online Reorganization status	QUERY DB STATUS(OLR) /DISPLAY DB OLR
Displaying IMS Fast DB Recovery status	F fdbproc,STATUS
Displaying IMS resources and status	/DISPLAY QUERY
Displaying IMSplex member status or attribute information	QUERY MEMBER
Displaying IRLM status	F irlmproc,STATUS ,ALLD,ALLI ,MAINT ,STOR ,TRACE
Displaying links with specified status	QUERY MSNAME STATUS(<i>status</i>)
Displaying MSNAMEs (logical link paths)	QUERY MSNAME QCNT(<i>condition,count</i>) QUERY MSLINK STATUS(<i>status</i>)
Displaying local queue counts	/DISPLAY LINK <i>link</i> QUERY MSNAME SHOW(QCNT)
Displaying local and remote system IDs	/DISPLAY ASMT MSNAME <i>msname</i> QUERY MSNAME SHOW(SYSID)
Displaying the LTERM name of the MTO	/RDISPLAY QUERY LTERM STATUS(MTO,SMTO)
Displaying maximum session count	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(MSLINK)
Displaying members of the IMSplex	QUERY IMSPLEX

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Displaying mode table name	/DISPLAY LINK <i>link</i> MODE [for MSC] /DISPLAY NODE <i>nodename</i> MODE [for non-MS] QUERY MSLINK SHOW(MODETBL) QUERY NODE SHOW(MODETBL)
Displaying MSPLINKs with specified status	QUERY MSPLINK SHOW(MSNAME)
Displaying MSPLINKs of a specified type	QUERY MSPLINK TYPE(<i>type</i>)
Displaying OLCSTAT data set information	QUERY OLC LIBRARY (OLCSTAT)
Displaying partner ID	/DISPLAY LINK <i>link</i> MODE QUERY MSLINK SHOW(PARTNER)
Displaying physical link type	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(TYPE)
Displaying send and receive counts	/DISPLAY LINK <i>link</i> /DISPLAY NODE <i>nodename</i> QUERY MSLINK SHOW(COUNT) QUERY NODE SHOW(COUNT)
Displaying status of assigned MSLINKs	QUERY MSPLINK SHOW(MSLINK)
Displaying status of a database or area	/DISPLAY DB /DISPLAY AREA QUERY DB QUERY AREA
Displaying status and rate information about HALDB Online Reorganizations that are in progress	QUERY OLREORG
Displaying status of MSLINKs	/DISPLAY LINK <i>link</i> /DISPLAY STATUS LINK QUERY MSLINK SHOW(STATUS)
Displaying status of MSNAMEs	/DISPLAY MSNAME <i>msname</i> /DISPLAY STATUS MSNAME QUERY MSNAME SHOW(STATUS)
Displaying status of MSPLINKs	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(STATUS)
Displaying VTAM node name	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(NODE)
Dumping MSDBs	/DBDUMP
Enabling logons to this physical link	/RSTART MSPLINK <i>name</i> UPDATE MSPLINK START(LOGON)
Ending special modes	/END
Exclusive mode for a terminal, setting	/EXCLUSIVE
External subsystem, entering commands for	/SSR
Fast DB Recovery surveillance, changing	/CHANGE
Forcing stop link	/PSTOP LINK <i>link</i> PURGE FORCE UPDATE MSLINK STOP(COMM) OPTION (FORCE)
Formatting a terminal screen	/FORMAT

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Gathering diagnostic data while IMS is active	/DIAGNOSE
Generating JCL for utilities	/RMGENJCL
IMS monitor, starting or stopping	/TRACE
Initiating IMS Fast DB Recovery of tracked databases	F fdbproc,RECOVER
Initiating a session with a terminal	/OPNDST
IRLM, abending	F irlmproc,ABEND
IRLM, changing MAXCSA (for IRLM Version 2.1)	F irlmproc,SET,CSA=nnn,TRACE=nnn
IRLM, displaying status	F irlmproc,STATUS ,ALLD,ALLI ,MAINT ,STOR ,TRACE
IRLM, reconnecting to	F proc, RECONNECT
IRLM, releasing locks	F irlmproc,PURGE,imsname
IRLM, setting timeout value	F irlmproc,START,TMOUT=n F irlmproc,STOP,TMOUT
IRLM, starting	S irlmproc
IRLM, stopping	P irlmproc
IRLM, tracing	S irlmproc,TRACE=YES TRACE CT
ISC node shutdown	/QUIESCE
ISC session cold start	/ASSIGN
Language Environment® (LE) attributes and parameters, changing	UPDATE LE DELETE LE
Limiting count for a transaction, setting	/ASSIGN UPDATE TRAN SET(LCT)
Listing all of the databases for which HALDB Online Reorganizations are in progress	QUERY DB STATUS(OLR)
Listing information in the RECON data set	/RMLIST
Logging off from a terminal	/RCLSDST
Logging on to IMS from a VTAM terminal	/OPNDST
LTERM, enqueueing messages to or dequeuing messages from	QUEUE LTERM
LTERM to node, assigning	/ASSIGN
LTERM to PTERM, assigning	/ASSIGN
LTERM to user or ISC half session, assigning	/ASSIGN
Message control error exit routine, activating	/DEQUEUE
Message Format Service (MFS) format, using on a terminal	/FORMAT
MFS test mode for terminals, setting	/TEST
Modifying applications	/MODIFY INITIATE UPDATE PGM

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Modifying databases	/MODIFY INITIATE UPDATE DB
Modifying IMS resources	INITIATE /MODIFY UPDATE
Modifying IMS resources globally for an IMSplex	INITIATE OLC
Modifying Fast Path routing code definitions	/MODIFY INITIATE UPDATE RTC
Modifying transactions	UPDATE TRAN
Monitoring and displaying the status of the specified databases or partitions (including those HALDB Online Reorganizations that are in progress)	/DISPLAY DB OLR
MTO, displaying LTERM name, line, and node	/RDISPLAY QUERY LTERM STATUS(<i>mta,smta</i>) SHOW(<i>node</i>)
Notifying DBRC about additional information, which gets recorded in the RECON data set	/RMNOTIFY
Online change, performing	INITIATE /MODIFY
Output errors, testing for	/LOOPTEST
Output segments for application program, assigning	/ASSIGN
Output to secondary MTO, sending	/SMCOPY
Preventing updates to a database	/DBDUMP /DBRECOVERY UPDATE AREA STOP(ACCESS) UPDATE DATAGRP STOP(ACCESS) UPDATE DB STOP(UPDATES) UPDATE DB STOP(ACCESS)
Priority for a transaction, assigning	/ASSIGN UPDATE TRAN SET(NPRI)
Querying a logical terminal (LTERM)	/DISPLAY LTERM <i>ltermname</i> QUERY LTERM NAME(<i>ltermname</i>)
Querying a node	/DISPLAY NODE <i>nodename</i> QUERY NODE NAME(<i>nodename</i>)
Querying a user	/DISPLAY USER <i>username</i> QUERY USER NAME(<i>username</i>)
Querying a user ID	/DISPLAY USER <i>useridname</i> QUERY USERID NAME(<i>useridname</i>)
Querying assigned Multiple Systems Coupling (MSC) logical links (MSLINKs)	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(MSLINK) QUERY MSNAME SHOW(MSLINK)
Querying assigned Multiple Systems Coupling (MSC) physical links (MSPLINKs)	/DISPLAY ASMT LINK <i>link</i> MODE QUERY MSLINK SHOW(MSPLINK) QUERY MSNAME SHOW(MSPLINK)

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Querying assigned MSNAMEs	/DISPLAY ASMT LINK <i>link</i> QUERY MSLINK SHOW(MSNAME) QUERY MSPLINK SHOW(MSNAME)
Querying bandwidth mode	/DISPLAY LINK OPTION BUFSIZE QUERY MSLINK SHOW(BANDWIDTH)
Querying buffer size	/DISPLAY LINK OPTION BUFSIZE QUERY MSLINK SHOW(BUFSIZE)
Querying CTC address	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(ADDR)
Querying global queue counts	/DISPLAY LINK <i>link</i> QCNT QUERY MSNAME SHOW(QCNT)
Querying HALDB Online Reorganization status	QUERY DB STATUS(OLR) /DISPLAY DB OLR
Querying IMS Fast DB Recovery status	F fdbproc,STATUS
Querying IMS resources and status	/DISPLAY QUERY
Querying IMSplex member status or attribute information	QUERY MEMBER
Querying IRLM status	F irlmproc,STATUS ,ALLD ,ALLI ,MAINT ,STOR ,TRACE
Querying links with specified status	QUERY MSNAME STATUS(<i>status</i>)
Querying MSNAMEs (logical link paths)	QUERY MSNAME QCNT(<i>condition,count</i>) QUERY MSLINK STATUS(<i>status</i>)
Querying local queue counts	/DISPLAY LINK <i>link</i> QUERY MSNAME SHOW(QCNT)
Querying local and remote system IDs	/DISPLAY ASMT MSNAME <i>msname</i> QUERY MSNAME SHOW(SYSID)
Querying the LTERM name of the MTO	/RDISPLAY QUERY LTERM STATUS(MTO,SMTO)
Querying maximum session count	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(MSLINK)
Querying members of the IMSplex	QUERY IMSPLEX
Querying mode table name	/DISPLAY LINK <i>link</i> MODE [for MSC] /DISPLAY NODE <i>nodename</i> MODE [for non-MSC] QUERY MSLINK SHOW(MODETBL) QUERY NODE SHOW(MODETBL)
Querying MSPLINKs with specified status	QUERY MSPLINK SHOW(MSNAME)
Querying MSPLINKs of a specified type	QUERY MSPLINK TYPE(<i>type</i>)

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Querying OLCSTAT data set information	QUERY OLC LIBRARY (OLCSTAT)
Querying partner ID	/DISPLAY LINK <i>link</i> MODE QUERY MSLINK SHOW(PARTNER)
Querying physical link type	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(TYPE)
Querying send and receive counts	/DISPLAY LINK <i>link</i> /DISPLAY NODE <i>nodename</i> QUERY MSLINK SHOW(COUNT) QUERY NODE SHOW(COUNT)
Querying status of assigned MSLINKs	QUERY MSPLINK SHOW(MSLINK)
Querying status of a database or area	/DISPLAY DB /DISPLAY AREA QUERY DB QUERY AREA
Querying status and rate information about HALDB Online Reorganizations that are in progress	QUERY OLREORG
Querying status of MSLINKs	/DISPLAY LINK <i>link</i> /DISPLAY STATUS LINK QUERY MSLINK SHOW(STATUS)
Querying status of MSNAMEs	/DISPLAY MSNAME <i>msname</i> /DISPLAY STATUS MSNAME QUERY MSNAME SHOW(STATUS)
Querying status of MSPLINKs	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(STATUS)
Querying VTAM node name	/DISPLAY ASMT MSPLINK <i>name</i> QUERY MSPLINK SHOW(NODE)
Queuing or dequeuing messages	/DEQUEUE QUEUE
RACF, controlling	/SECURE
Reconnecting to a restarted IRLM	F job,RECONNECT
Reconnecting to coupling facility structures	F job,RECONNSTR
Recovering databases by using the Database Recovery Facility	/RECOVER
Releasing IRLM locks	F irlmproc,PURGE,imsname
Remote takeover, initiating	/RTAKEOVER
Removing IMS from a VTAM generic resources group	/CHECKPOINT /STOP
Removing an IMS from the IMSplex	/CHE FREEZE or DUMPQ or PURGE LEAVEPLEX
Resetting values set during IMS system definition	/ASSIGN UPDATE

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Restarting a component of IMS	/OPNDST /RELEASE /RSTART /START /UNLOCK UPDATE TRAN START(Q,SCHD,TRACE) UPDATE AREA START(ACCESS) UPDATE DATAGRP START(ACCESS) UPDATE DB START(ACCESS) UPDATE DB START(ACCESS) AREA(*) UPDATE DB SET(LOCK(OFF))
Restarting IMS after failure	/ERESTART
Restarting XRF alternate subsystem	/ERESTART
Resuming a conversation	/RELEASE
Resuming HALDB Online Reorganization for one or more partitions	INITIATE OLREORG /INITIATE OLREORG
Scheduling a suspended transaction	/DEQUEUE UPDATE TRAN START(SUSPEND)
Secondary MTO, sending output to	/SMCOPY
Security, controlling	/SECURE
Sending messages to terminals or other IMS subsystems	/BROADCAST LTERM QUEUE LTERM
Setting the RATE for a HALDB Online Reorganization	INITIATE OLREORG SET(RATE(rate)) /INITIATE OLREORG SET(RATE(rate)) UPDATE OLREORG SET(RATE(rate)) /UPDATE OLREORG SET(RATE(rate))
Setting or resetting ASR	/CHANGE LINK <i>link</i> ASR ON OFF UPDATE MSLINK SET(ASR(ON OFF)) UPDATE MSPLINK SET(ASR(ON OFF))
Setting timeout value for IRLM	F irlmproc,START,TMOUT=n F irlmproc,STOP,TMOUT
Sharing startup JCL between IMS dependent regions	/START
Shutting down Common Queue Server (CQS)	P cqsproc
Shutting down the Common Service Layer (CSL)	F sciproc, SHUTDOWN CSLLCL/CSLPLEX
Shutting down IMS	/CHECKPOINT
Shutting down an ISC node	/QUIESCE
Shutting down Operations Manager (OM)	P omjob
Shutting down Resource Manager (RM)	P rmjob
Shutting down Structured Call Interface (SCI)	P scijob
Signing on to IMS from a non-VTAM terminal	/IAM
Signing on or off IMS terminals	/SIGN
Specifying that a HALDB master is capable of being reorganized online	INIT.DB OLRCAP CHANGE.DB OLRCAP

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Specifying whether to delete the inactive data sets after the copying phase of a HALDB online reorganization completes	UPDATE OLREORG OPTION(DEL NODEL) /UPDATE OLREORG OPTION(DEL NODEL)
Specifying whether to release ownership of OLR at normal or abnormal shutdown of IMS	UPDATE OLREORG OPTION(REL NOREL) /UPDATE OLREORG OPTION(REL NOREL)
Starting bandwidth mode	UPDATE MSLINK SET(BANDWIDTH(ON OFF))
Starting Common Queue Server (CQS)	S cqsproc
Starting HALDB Online Reorganization for one or more partitions	INITIATE OLREORG /INITIATE OLREORG
Starting IMS	/NRESTART /ERESTART
Starting IMS Fast DB Recovery	S fdbbrproc
Starting IMS resources	/START UPDATE TRAN START(Q,SCHD,TRACE) UPDATE AREA START(ACCESS) UPDATE DATAGRP START(ACCESS) UPDATE DB START(ACCESS) UPDATE DB START(ACCESS) AREA(*)
Starting an internal trace	/TRACE SET ON OFF LINK <i>link</i> UPDATE MSLINK START(TRACE)
Starting IRLM	S irlmproc
Starting a line, link, node, PTERM, or user	/RSTART UPDATE MSLINK(<i>linkname</i>) START(COMM)
Starting msname queuing or sending	/START MSNAME <i>name</i> UPDATE MSNAME START(Q,SEND))
Starting an XRF takeover trace	/TRACE SET ON OFF LINK <i>link</i> TAKEOVER UPDATE MSLINK START(TKOTRC)
Statistics, recording to IMS log	/CHECKPOINT
Stopping bandwidth mode	UPDATE MSLINK SET(BANDWIDTH(ON OFF))
Stopping a line, a logical link, or a logical terminal	/PSTOP UPDATE MSLINK STOP(COMM)
Stopping database access	/LOCK /STOP UPDATE AREA STOP(SCHD) UPDATE DATAGRP STOP(SCHD) UPDATE DB SET(LOCK(ON)) UPDATE DB STOP(SCHD)
Stopping HALDB Online Reorganization for one or more partitions	TERMINATE OLREORG /TERMINATE OLREORG
Stopping IMS resources	/STOP UPDATE AREA STOP(SCHD) UPDATE DATAGRP STOP(SCHD) UPDATE DB STOP(SCHD) UPDATE TRAN STOP(Q,SCHD)

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Stopping input to a line, link, or terminal	/PURGE UPDATE MSNAME STOP(Q),START(SEND)
Stopping input messages with a transaction code	/PURGE UPDATE TRAN START(SCHD)STOP(Q)
Stopping IMS Fast DB Recovery	F fdbproc,DUMP F fdbproc,STOP
Stopping IMS Fast DB Recovery tracking	F fdbproc,TERM
Stopping input for msname	/PURGE MSNAME <i>name</i> UPDATE MSNAME STOP(Q) START(SEND)
Stopping an internal trace	/TRACE SET ON OFF LINK <i>link</i> UPDATE MSLINK STOP(TRACE)
Stopping IRLM	P irlmproc
Stopping an LTERM, node, program, PTERM, or transaction	/LOCK SET(LOCK(ON))
Stopping a transaction	/STOP UPDATE TRAN STOP(Q,SCHD)
Stopping msname queuing or sending	/STOP MSNAME <i>name</i> UPDATE MSNAME STOP(Q,SEND)
Stopping output to a line, link, or terminal	/PSTOP UPDATE MSLINK STOP(COMM)
Stopping output for msname	UPDATE MSNAME START(Q) STOP(SEND)
Stopping output to a programmable remote station	/MONITOR
Stopping scheduling of a message with a transaction code	/PSTOP UPDATE TRAN STOP(SCHD) START(Q)
Stopping an XRF takeover trace	/TRACE SET ON OFF LINK <i>link</i> TAKEOVER UPDATE MSLINK STOP(TKOTRC)
Subsystem, external, entering commands for	/SSR
Suspending a conversation	/HOLD
Switching data sets for XRF	/SWITCH
Terminating a conversation	/EXIT
Terminating global online change on all IMS systems in an IMSplex	TERMINATE OLC
Terminating IMS without a dump	F job,STOP F job,STOPxxx F job,FORCExxx
Terminating IMS with a dump	F job,DUMP F job,DUMPxxx
Terminating IRLM abnormally	F irlmproc,ABEND
Terminating I/O for terminals	/IDLE
Terminating preset mode	/RESET
Terminating special modes	/END
Test mode for terminals, setting	/TEST
Testing for output errors	/LOOPTEST

Table 15. Commands for IMS operations tasks (continued)

Operations task	Command
Tracing IMS resources	/TRACE UPDATE TRAN START(TRACE) UPDATE MSLINK START(TRACE)
Tracing IRLM	S irlmproc,TRACE=YES TRACE CT
Transactions, enqueueing messages to or dequeuing messages from	QUEUE TRAN
Transport Manager Subsystem, defining	DEFINE
Transport Manager Subsystem, displaying	DISPLAY
Transport Manager Subsystem, setting	SET
Transport Manager Subsystem, starting	START
Transport Manager Subsystem, stopping	STOP
Undoing /ASSIGN	/ASSIGN /NRE CHECKPOINT 0 UPDATE TRAN
Undoing /EXCLUSIVE	/END /START
Undoing /HOLD	/RELEASE /EXIT
Undoing /IAM	/IAM
Undoing /IDLE	/ACTIVATE
Undoing /LOCK	/UNLOCK UPDATE DB SET(LOCK(OFF))
Undoing /LOOPTEST	/END
Undoing /MSASSIGN	/MSASSIGN /NRE CHECKPOINT 0 UPDATE MSLINK NAME(linkname) SET(MSLINK(plinkname) UPDATE MSNAME NAME(msname) SET(MSLINK(linkname)
Undoing /SET	/RESET
Undoing /START, /RSTART, /PSTART, /PURGE, or /MONITOR	/STOP UPDATE
Undoing /START, /STOP, /COMPT, /RCOMPT, /PSTART, /PURGE, or /MONITOR	/RSTART UPDATE
Undoing /START, /MONITOR, /RSTART, /STOP, or /PSTART	/PURGE UPDATE
Undoing /START, /MONITOR, /RSTART, /STOP, or /PURGE	/PSTOP UPDATE
Undoing /START, /RSTART, /STOP, /PSTOP, or /PURGE	/MONITOR UPDATE
Undoing /TEST	/END /IAM /START
Unloading an area from a data space	/VUNLOAD

Table 15. Commands for IMS operations tasks (continued)


Operations task	Command
Verifying local transactions and LTERMs with remote transactions and LTERMs	/MSVERIFY
Verifying /MSASSIGN	/MSVERIFY
Warm start of IMS	/NRESTART
Writing a message to the IMS log	/LOG
XRF, switching data sets	/SWITCH

Important: In an IMSplex, type-2 commands are issued to the Operations Manager, which then routes the commands to the subsystem IMSplex members. When commands are issued globally in an IMSplex, they can behave differently from commands that are issued to single IMS systems.

Related concepts:

 Restrictions for HALDB Online Reorganization (Database Administration)

Related reference:

 Equivalent IMS type-1 and type-2 commands (Commands)

“List of commands with similar functions for multiple resources” on page 15

Chapter 2. Starting or restarting IMS

Starting or restarting an IMS subsystem means initializing the IMS control region. You must first restore the operating environment and the master terminal operator must insure that critical applications, such as any dependent regions, are initialized.

To restore the operating environment, the MTO must ensure the following are also initialized:

- All IRLMs
- All CQSs
- Any dependent regions (including utilities)
- All CSL address spaces (ODBM, OM, RM, and SCI), if operating as an IMSplex using CSL
- Any connections to VTAM
- All communication lines and terminals

The term MTO used in this information refers to the IMS master terminal operator for the DB/DC and DCCTL environments, and to the operator for the DBCTL environment. All commands shown assume that the command-recognition character is the forward slash (/).

Starting an IMSplex

Starting an IMSplex involves starting required CSL components, such as SCI and OM, starting the IMS control regions, and starting any other optional IMSplex components, such as CSL RM, CSL ODBM, IMS Connect, and so on.

The following briefly summarizes starting an IMSplex:

1. Start all SCIs.
2. Start all OMs.
3. If using a resource structure, start CQS.
4. If using RM, start all RMs.
5. Start all IMS control regions with the following parameters:
 - CSLG= parameter specified on the control region execution or the DFSPBxxx member of the IMS PROCLIB data set.
 - Parameters specified in the DFSDCxxx member of the IMS PROCLIB data set if you do not want to use the defaults:
 - RCVYSTSN
 - RCVYFP
 - SRMDEF
 - OCMD and CSLT parameters specified in the DFSVSMxx member of the IMS PROCLIB data set.
 - Parameters specified in the DFSCGxxx member of the IMS PROCLIB data set:
 - CMDSEC
 - IMSPLEX

- NORSCCC
 - OLC
 - OLCSTAT
 - Appropriate parameters specified in the BPE and CQS PROCLIB member data sets.
6. Start any other optional IMSplex components, such as ODBM and IMS Connect.
 7. If you use a SPOC, you can log on to TSO, allocate the appropriate files to the TSO user ID, and issue the DFSSPOC command from ISPF option 6.

Starting the CSL

If you are running IMS as an IMSplex with CSL, you need to start the CSL manager address spaces (ODBM, OM, RM, and SCI) before starting the IMS control region. You can start the CSL manager address spaces using the z/OS START command from a z/OS system console or by using JCL.

The Common Service Layer (CSL) is comprised of multiple address spaces. The CSL is considered “started” after each of the CSL manager address spaces you require is started.

Initialization parameters must be defined for each address space before CSL can be started. To customize the parameters of a specific CSL manager address space, use the following members of the PROCLIB data set:

- CSLDIxxx, which initializes the Open Database Manager
- CSLOIxxx, which initializes the Operations Manager
- CSLRIxxx, which initializes the Resource Manager
- CSLSIxxx, which initializes the Structured Call Interface

The parameters that you define using the PROCLIB data set vary on the address space that you are customizing, but can include the IMSplex name, the address space name, the use of Automatic Restart Manager (ARM), and more.

For ODBM, IMS also provides a CSLDCxxx member of the IMS.PROCLIB data set, which contains parameters that define the IMS data store connections to ODBM.

ODBM, OM, RM, and SCI are initially started by running their respective startup procedures. To run an ODBM, OM, RM, or SCI procedure, you can either issue a z/OS START command or code and run JCL. ODBM, OM, RM, and SCI can be restarted using the z/OS Automatic Restart Manager (ARM). You can specify the parameter ARMRST either in the startup procedures or in the execution parameters of ODBM, OM, RM, and SCI. The default is ARMRST=Y, so if you do not specify the parameter, all of the CSL managers will try to register with ARM.

Before you can start any of the CSL components, they must be defined to the IMS system.

The IMS PROCLIB data set members DFSCGxxx and DFSDFxxx define how the IMS control region interacts with the CSL.

Related concepts:

➡ CSL definition and tailoring (System Definition)

Related reference:

➡ DFSDFxxx member of the IMS PROCLIB data set (System Definition)

➡ DFSCGxxx member of the IMS PROCLIB data set (System Definition)

Starting the CSL SCI

The system operator can start the CSL SCI in one of three ways: as a z/OS started procedure, as JCL, or as a started procedure in an IMS control region.

You can start SCI in the following ways:

- The system operator can use a z/OS started procedure. CSLSC000, a sample started procedure that is shipped with IMS, can be modified for your installation, or you can identify your own started procedure. To start an SCI address space with a started procedure, issue the z/OS START command as follows: S scijobname.scijobname is the job name of the SCI address space to be started.
- The system operator can submit JCL that starts the started procedure.
- The IMS control region can use a started procedure in an RMENV=N environment by using the SCIPROC parameter in the DFSCGxxx or DFSDFxxx PROCLIB member data sets. If RMENV=N is specified and there is no SCI address space active in the IMSplex, IMS will start the specified procedure.

Restarting the CSL SCI

After the Structured Call Interface (SCI) is started, if it is abnormally terminated, it can be restarted by using the z/OS Automatic Restart Manager (ARM). SCI must complete initialization for ARM to restart the address space if an abend occurs. Use of ARM to restart SCI is the default.

Starting the CSL ODBM

You can start ODBM either before or after starting IMS, IMS Connect, or z/OS Resource Recovery Services (RRS).

If ODBM and an IMS system are not members of the same IMSplex, ODBM automatically connects to the IMS system only if the IMS system is started before ODBM.

To enable ODBM to connect to the IMS system automatically regardless of the order in which ODBM and IMS are started, use either the DFSCGxxx member of the IMS.PROCLIB data set or the Common Service Layer section of the DFSDFxxx PROCLIB member to define the IMS system as a member of the same IMSplex as ODBM.

The options for starting ODBM include:

- The system operator can use a z/OS started procedure. For example, CSLODBM, a sample started procedure that is shipped with IMS, can be modified for your installation, or you can identify your own started procedure. To start an ODBM address space with a started procedure, issue the z/OS START command as follows:
S odbmjobname

odbmjobname is the job name of the ODBM address space to be started.

- The system operator can submit JCL that starts the started procedure.

Restarting the CSL ODBM

In the event that ODBM terminates abnormally, it can be restarted by using the z/OS Automatic Restart Manager (ARM).

ODBM must complete initialization for ARM to restart the address space if an abend occurs. Use of ARM to restart ODBM is the default.

Starting the CSL OM

The system operator can start the CSL OM three ways: using a z/OS started procedure, or by submitting JCL, or by using a started procedure in an IMS control region.

You can start OM in the following ways:

- The system operator can use a z/OS started procedure. For example, CSLOM000, a sample started procedure that is shipped with IMS, can be modified for your installation, or you can identify your own started procedure. To start an OM address space with a started procedure, issue the z/OS START command as follows: S omjobname omjobname is the job name of the OM address space to be started.
- The system operator can submit JCL that starts the started procedure.
- The IMS control region can use a started procedure in an RMENV=N environment by using the OMPROC parameter in the DFSCGxxx or DFSDFxxx member of the IMS PROCLIB data sets. If RMENV=N is specified and there is no OM address space active in the IMSplex, IMS will start the specified procedure.

Restarting the CSL OM

If, after OM is started, if it is abnormally terminated, it can be restarted using the z/OS Automatic Restart Manager (ARM). OM must complete initialization for ARM to restart the address space if an abend occurs. Use of ARM to restart OM is the default.

Starting the CSL RM

The system operator can use a z/OS started procedure. CSLRM000, a sample started procedure that is shipped with IMS, can be modified for your installation, or you can identify your own started procedure.

You can start RM in the following ways:

- To start an RM address space with a started procedure, issue the z/OS START command as follows:

```
S rmjobname
```

rmjobname is the job name of the RM address space to be started.

- The system operator can submit JCL that starts the started procedure.

You can also add a resource structure to your CFRM policy.

Related tasks:

 Defining the CFRM policy (System Administration)

Restarting the CSL RM

If the resource manager (RM) terminates abnormally, it can be restarted using the z/OS Automatic Restart Manager (ARM). RM must complete initialization for ARM to restart the address space if an abend occurs. Use of ARM to restart RM is the default.

Starting and stopping the IMSRSC repository

An IMSRSC repository can have a status of closed, opening, open, stopping, or stopped. Start or stop a repository to make it available or unavailable for client connections.

CLOSED

The repository is defined to the Repository Server (RS) and it is started, but its associated data sets are closed. The RS accepts client connections to the repository, but because the repository has not yet been opened and its associated data sets have not been allocated by the RS, you cannot write or read from the repository in this state.

If a repository remains in CLOSED state at the end of SPARE recovery processing, perform the following steps to make the repository available:

1. Re-allocate the discarded repository data sets.
2. Issue the LIST FRPBATCH command to check if the repository is in a STOP state. If it is not, issue the STOP FRPBATCH command to change the state to STOP.
3. Issue the DSCHANGE FRPBATCH command to change the re-allocated discarded data sets back to a SPARE state.
4. Issue the START FRPBATCH command to start the repository.

This resumes the data set recovery process.

OPENING

The repository is in the process of being opened.

OPEN

Started and opened. The repository is available for client connections and read and update requests.

Starting a repository makes it possible for connections to be made to it. The following functions are performed when a repository is started:

- The repository becomes available for client connections.
- If the repository is set with the ADD or UPDATE FRPBATCH commands to AUTOOPEN=Y, or the repository was started with the OPEN=YES option in the START FRPBATCH command, the RS opens the repository, readying it for client function requests.
- If the repository is not set to open either explicitly or automatically, the repository and RS catalog repository data sets are allocated when a client first connects to the repository.

STOPPING

The repository is in the process of being stopped.

Stopping a repository has the following results:


- The repository can no longer accept new client connections.
- If the repository data sets are allocated, the RS tries to close and deallocate them. The RS waits until all active client request threads are completed before closing and deallocating the data sets.


STOPPED

Stopped and closed. The RS does not accept client connections to the repository. The repository data sets are closed and deallocated.

The start and stop state of a repository is persistent. For example, a repository that is stopped remains stopped after the RS is restarted, regardless of the setting for the AUTOOPEN option. The AUTOOPEN option only determines when data sets are allocated and opened if the repository is started.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

 IMSRSC repository and RS catalog repository data sets (System Definition)

Related information:

 Commands for FRPBATCH (System Programming APIs)

Opening the IMSRSC repository

You can control when an IMSRSC repository is opened with the AUTOOPEN option for the ADD or UPDATE FRPBATCH commands, or with the OPEN option for the START FRPBATCH command.

- If the repository is set to AUTOOPEN=Y, its data sets are allocated and opened when it starts. If the repository was in the started state when the Repository Server (RS) last shut down, the data sets are allocated and opened when the RS is restarted.
- If the repository is set to AUTOOPEN=N and it is started with OPEN=YES, it is opened immediately.
- If the repository is set with AUTOOPEN= N and it is not started with OPEN=YES, it is opened at the time of the first client connection.

The RS performs basic validation of IMSRSC repository data sets at open time. A repository cannot be opened unless there is both a valid COPY1 and COPY2 of the IMSRSC repository data set.

If there is one valid copy and a valid SPARE IMSRSC repository data set, automated SPARE recovery occurs. In the SPARE recovery process, the SPARE IMSRSC repository data set is reassigned to take the place of the missing COPY1 or COPY2, and its data sets are populated from the valid copy.

At the conclusion of data set validation:




- The available IMSRSC repository data sets are allocated and open.
- Empty data sets are identified.
- All data sets have valid characteristics for their identified function as either a repository member data set (RMD) or repository index data set (RID).
- Populated data sets (with resource definitions) are considered to be valid repository data sets.
- The repository's set of data sets is known to be consistent.

If any data set fails to meet these basic requirements, an error message is issued and the repository open process terminates.

During repository open processing, either with the AUTOOPEN option or during the first access, all of the repository data sets are opened using VSAM local shared resources (LSR). As the data sets are opened using VSAM LSR, if the repository data sets are empty, informational messages are issued that indicate that the data sets are empty. An informational messages received for the spare data set pair requires no user action. The spare data set is opened to ensure that it is empty, so that there are no errors during SPARE recovery process.

The RS closes the IMSRSC repository data sets when the repository is stopped, during the SPARE recovery process and during RS shutdown.

Related concepts:

-  Overview of the IMSRSC repository (System Definition)
-  IMSRSC repository and RS catalog repository data sets (System Definition)
-  Recovery in an IMSplex (System Administration)

Related information:

-  Commands for FRPBATCH (System Programming APIs)

Starting the IMSRSC repository

To start the IMSRSC repository, you can either use a START command issued from the z/OS MODIFY (F) interface or issue START FRPBATCH command.

To start the repository, do one of the following:

- Issue the following command from the z/OS MODIFY (F) interface:
`F reposervername,ADMIN START repositoryname`
- Issue the START FRPBATCH command.

Related concepts:

-  Overview of the IMSRSC repository (System Definition)

Related reference:

-  F reposervername,ADMIN (Commands)
-  START command for FRPBATCH (System Programming APIs)


Stopping the IMSRSC repository

To stop the IMSRSC repository, you can either use a STOP command issued from the z/OS MODIFY (F) interface or issue STOP FRPBATCH command.

To stop the repository, do one of the following:

- Issue the following command from the z/OS MODIFY (F) interface:
`F reposervername,ADMIN STOP repositoryname`
- Issue the STOP FRPBATCH command.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

Related reference:


 F reposervername,ADMIN (Commands)

 STOP command for FRPBATCH (System Programming APIs)

Starting and stopping the Repository Server

To start the Repository Server (RS), submit a customized RS startup procedure. To stop the RS, use an RS command issued from the z/OS MODIFY (F) interface.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

Related reference:

 F reposervername,ADMIN (Commands)

Starting the Repository Server

To start the Repository Server (RS), submit a customized RS startup procedure.

Perform the following procedure:

1. Make sure that you have performed system definition for the IMSRSC repository.
2. Customize the following startup procedure:

```
//FRP12A  PROC RGN=0M,SOUT=A,
//*
//IEFPROC EXEC PGM=BPEINI00,REGION=0M,
//  PARM='BPEINIT=FRPINI00,BPECFG=BPECONFIG,FRPCFG=FRPCONFG'
//*
//STEPLIB DD  DSN=RESLIB_dataset_name,DISP=SHR
//*
//PROCLIB DD  DSN=PROCLIB_dataset_name,DISP=SHR
//*
//FRPPRINT DD SYSOUT=&SOUT
//SYSPRINT DD SYSOUT=&SOUT
//SYSUDUMP DD SYSOUT=&SOUT
//*
```

Make sure that the values specified in this JCL for the BPECFG= and FRPCFG= parameters match the values specified in the FRP and BPE configuration members of the IMS PROCLIB data set.

3. Submit the RS startup procedure.

If the RS is started successfully as the active RS, then the FRP2002I message is issued.

If another RS is already started in the same z/OS cross-system coupling facility group, a FRP2001I message is issued, indicating that the subordinate RS is started.

During startup of both the active and subordinate RS, the FRP2006I message is issued.

When the RS has started successfully and is ready to accept client connections, the FRP2025I message is issued.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

Related tasks:


“Stopping the Repository Server”

 Defining the IMSRSC repository (System Definition)

“Cold starting an IMS system that uses the IMSRSC repository” on page 76

Related reference:

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

 BPE configuration parameter member of the IMS PROCLIB data set (System Definition)

 FRP messages (Repository Server) (Messages and Codes)

Related information:

 FRP2002I (Messages and Codes)

 FRP2001I (Messages and Codes)

 FRP2006I (Messages and Codes)

 FRP2025I (Messages and Codes)

Starting subordinate Repository Servers

To start a Repository Server (RS) as a subordinate, you only need to define it with the same z/OS cross-system coupling facility group as an RS that is already running.

However, all parameters in the FRPCFG member for both the master RS and its subordinates should be the same, except for the RSNAMES= parameter. RSNAMES= should be unique for each RS, including the master RS and its subordinates. RSNAMES= is used to register to the Structured Call Interface (SCI) if IMSPLEX() is specified. SCI requires unique names for the clients.

Make sure that the subordinate RSs use the same RS catalog repository data sets as the master RS.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

 Recovery in an IMSplex (System Administration)

Stopping the Repository Server

To stop the Repository Server (RS), you can use an RS command issued from the z/OS MODIFY (F) interface.


To stop the RS, issue the following command from the z/OS MODIFY (F) interface:
F reposervername,SHUTDOWN

To stop all the RS address spaces, issue the following command from the z/OS MODIFY (F) interface. The command can be directed to a master RS or subordinate RS:

F reposervername,SHUTDOWN ALL

You can also stop the RS by issuing the z/OS STOP (P) command.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

Related tasks:

“Starting the Repository Server” on page 68

Related reference:

 F reposervername,SHUTDOWN (Commands)

How to start the IMS control region

You usually start the IMS control region with a z/OS START command from a z/OS system console.

If, depending on the size of your installation, you have an alternate system console for the IMS MTO, be sure the z/OS system operator relays all IMS messages (that is, those beginning with BPE, CQS, DFS, DSP, DXR, and ELX) to the IMS MTO. If the z/OS system console is physically separate from the IMS MTO, you should have a local telephone or paging line to enable communication between the z/OS operator and the IMS MTO. The MTO should also have easy access to the operators responsible for DASD and tape units.

When the IMS control region is started, IMS issues two DFS1929I messages. Message DFS1929I lists the system parameters and their values defined for the IMS system. This information will help you understand the configuration of a particular system. Understanding the configuration will be helpful in diagnosing problems. The system parameters are also listed in the job log.

IMS issues message DFS1929I once during initialization and again when restart completes. The first DFS1929I message reflects the system parameter values that are defined in the DFSPBxxx PROCLIB member. The second DFS1929I message reflects the system parameter values that are actually in effect for the current execution of IMS.

For online IMS systems, the control region automatically starts the DBRC address space. If you are running an IMS DB/DC or DBCTL environment, you can specify LS0=S as an EXEC parameter for the control region to start the DL/I address space. The following rules apply:

- IMS initialization waits until there is a separate address space available for DBRC.
- You can start DBRC manually (if, for example, a JCL error prevented DBRC from starting).
- When you restart IMS, you can change the specification of LS0=; IMS does not use any prior specifications.

After the control region has started, the MTO enters an IMS command to start the system. Select one of three ways of starting IMS:

1. Cold start
2. Warm start (also referred to as normal restart)
3. Emergency restart

Your operating procedures should include guidelines for the MTO in selecting the right type of startup.

You can alternatively use automatic restart instead of having the MTO enter an IMS command to restart the system. Specify AUTO=Y in the control region JCL to request automatic restart. IMS selects a suitable checkpoint for restart and notifies the MTO. You cannot use automatic restart if you:

- Want to perform a cold start
- Need to specify any options on the restart command (such as formatting system data sets, changing security options, or restarting only one portion of a DB/DC system)
- Need to override DBRC after a failure in which DBRC is unable to mark the subsystem record in the RECON data set as abnormally terminated, such as in the case of a power, CPC, z/OS, or DBRC failure

Setting the z/OS TOD clock

The z/OS time-of-day (TOD) clock setting is critical to the integrity of the IMS log. The TOD clock is also critical to the proper functioning of IMS restart, data sharing, XRF tracking and takeover, and RSR tracking and takeover.

Recommendations: Use one of the following procedures:

- Use an External Time Reference (ETR) device. Be sure to set it to use Universal Coordinated Time (also known as Greenwich Mean Time) with the appropriate offset for your local time zone.
- Set the z/OS TOD clock during IPL using Universal Coordinated Time, then set the local time.

Attention: During IPL, do not set the TOD clock to a time and date **earlier** than the immediate prior shutdown or failure. Setting the TOD clock back has severe IMS database integrity and recovery implications. If you do set the TOD clock back, you must reallocate the OLDSs using a different block size, reinitialize the DBRC RECON data sets, take image copies of all DBDSs, and cold start IMS.

Considerations for setting local time

Set your system's time zone offset (the difference between local time and Universal Coordinated Time (UTC)) to an integer multiple of 15 minutes. This is consistent with world time zone definitions and ensures that the local time reported by IMS matches the local time reported by the operating system.

IMS will function properly even if the time zone offset is not a multiple of 15 minutes. However, due to the way IMS stores time values internally, not following this recommendation might result in IMS reporting a local time that is inconsistent with the local time reported by the operating system.

More about IMS time stamps and setting the time zone offset

IMS stores all date/time stamps internally as 12 bytes in the following packed-decimal format:

YYYYDDDF HHMMSSth mijAQQ\$

The date and time fields (YYYYDDDF and HHMMSSthmij) are in Universal Coordinated Time (UTC). The A field consists of 4 bits of flags used internally by IMS. The QQ\$ field contains the signed packed-decimal local time zone offset for when the time stamp was created. To get the local time, IMS adds the time zone offset—stored in quarter hours (fifteen-minute intervals)—to the UTC stamp. For example, the time stamp:

2000353F 06420588 4242032D

represents the UTC date and time of:

06:42:05.884242 on day 353 of 2000 (December 18th)

The QQ\$ field of this time stamp is 32D, which represents a local time zone offset of minus 32 quarter hours (or minus 8 hours west of Greenwich). Thus, the local time represented by this time stamp is:

22:42:05.884242 on day 352 of 2000 (December 17th)

This format allows IMS internal time stamps to encode both the UTC time (to provide an always-incrementing time that is not subject to daylight savings time changes), and local time (to determine the “wall-clock” time for when the time stamp was generated). If your time zone offset is set to a multiple of fifteen minutes, IMS can accurately use the offset to convert from the UTC stored in the 12-byte time stamp into the local time.

In this example, the local time will be consistent with the time reported by the operating system. However, if your time zone offset is not set to a multiple of fifteen minutes, the local time derived from this conversion will not be consistent with the time reported by the operating system.

In the following example, the time zone offset is set to +01:55 (one hour and fifty-five minutes east of Greenwich). A time stamp is generated at 15:00:00.000000 UTC time (16:55:00.000000 local time, using the current time zone offset setting). IMS will construct an internal time stamp of:

2001240F 15000000 0000008C

The encoded time zone offset is 8C, or +8 quarter hours (2 hours). Because IMS cannot store a time zone offset with a precision greater than fifteen minutes, it rounds to the nearest fifteen minute boundary.

When this time stamp is converted back into local time, IMS reports a time of 17:00:00.000000, not 16:55:00.000000. This does not cause problems for IMS, which uses the more accurate UTC form of the time stamp internally. However, you will notice this rounding effect when you compare IMS-generated local times to the current operating system time. For example, if you issue an IMS checkpoint command on the system described above, the DFS058I message will indicate two different times, as shown below:

16:55:00 DFS058I 17:00:00 CHECKPOINT COMMAND IN PROGRESS

The first time value (16:55) is returned by the operating system; this value is generated by using the z/OS time zone offset of +01:55. The second time value (17:00) is returned by IMS; this value is generated from an internal IMS time stamp that uses a time zone offset of +08 quarter hours (2 hours).

This behavior may also be encountered by application programs that compare the local time stamp in the I/O PCB with the value returned from the z/OS TIME macro. The I/O PCB local time represents the time when the message was received by IMS. The data in this time field is derived from an internal IMS 12-byte time stamp stored with the message.

In the following example, a message is received by IMS at 12:42 UTC on the system described above. The local time is 14:37, based on the z/OS time zone


offset of +01:55. However, the time presented to the application by IMS in the I/O PCB local time field is 14:42, due to the rounding of the time zone offset.

Now, suppose the message is scheduled at 12:43 UTC, one minute after it arrived. If the application program issues a z/OS TIME macro request, it will receive back a local time of 14:38. If the application were to compare this time to the I/O PCB time, it would appear as if the message was being processed before it was enqueued (enqueue time = 14:42; process time = 14:38). This is only a problem for an application program if it compares the I/O PCB time to a time generated from an operating system time service.

Related tasks:

“Setting or changing local time”

Related reference:

 External Time Reference (ETR)

Setting or changing local time

You can reset local time whenever necessary, for example when changing from standard time to daylight savings (or summer) time.

1. To set or change the local time, edit the CLOCKxx member (the TIMEZONE keyword) of the SYS1.PARMLIB data set to set the current offset for local time.
2. If you have an ETR device, use it to set and change the time zone offset. If you do not have an ETR device, complete one of the following steps to set or change the local time:

Option	Description
Use the SET CLOCK command to set the time zone offset.	Using the SET CLOCK command does not allow you to set a precise value for your system's time zone offset. Consequently, using this method might cause IMS to report a local time that is inconsistent with the local time reported by the operating system.
Re-IPL your system with the updated CLOCKxx member.	Your time zone offset is set to the precise value specified in the member.

Related concepts:

“Setting the z/OS TOD clock” on page 71

Cold starting IMS in a non-shared queues environment

You perform a cold start when starting IMS for the first time, or after changing the nucleus in a system definition. IMS never uses information from a previous shutdown as input to a cold start.

In a nonshared-queues environment, a cold start assumes empty message queues, so IMS discards any messages that exist. You must close the last OLDS used by the online system before you attempt a cold start. You should also archive the last OLDS before a cold start, so the OLDS is available if a database recovery is required.

Cold starting IMS in a shared-queues environment

You perform a cold start when starting IMS for the first time, or after changing the nucleus in a system definition. IMS never uses information from a previous shutdown as input to a cold start.

In a shared-queues environment, IMS does not discard messages on shared queues. If you want to discard messages on shared queues, you must scratch the shared-queue structures on the coupling facility, and scratch and reallocate the CQS structure recovery data sets.

Cold starting IMS with dynamic resource definition

You perform a cold start when starting IMS for the first time, or after changing the nucleus in a system definition. IMS never uses information from a previous shutdown as input to a cold start.

During a cold start, resource and descriptor definitions are automatically imported from the IMSRSC repository, resource definition data set (RDDS), or the MODBLKS data set, or IMS comes up with no resources.

During a cold start, the partition status or access state will be copied from the HALDB master. If you export the HALDB master status to the RDDS or the repository, the status of the master and its partitions will be obtained from the RDDS or the repository.

Automatic import of resource and descriptor definitions is defined through the AUTOIMPORT parameter in the DYNAMIC RESOURCES section of the DFSDFxxx member of the IMS PROCLIB data set

If AUTOIMPORT=AUTO is specified, these conditions determine the source of the imported resource and descriptor definitions:


If all of the following conditions are true, automatic import from the IMSRSC repository is enabled:

- IMS is enabled with DRD.
- The REPOSITORY section of the DFSDFxxx member of the IMS PROCLIB data set is defined with TYPE=IMSRSC.
- IMS is enabled with RM services (RMENV=N is not specified in the DFSCGxxx member of the IMS PROCLIB data set or in the CSL section of the DFSDFxxx member of the IMS PROCLIB data set)
- The CSLRIxxx member of the IMS PROCLIB data set is defined for the repository with TYPE=IMSRSC.
- The repository contains stored resource definitions for the IMS system.
- RM is started with the repository enabled.
- The RS address space is started and available.

If all of the following conditions are true, automatic import from an RDDS is enabled:

- IMS is enabled with DRD.
- The REPOSITORY section of the DFSDFxxx member of the IMS PROCLIB data set is not present.
- Two or more system RDDSs are defined in the RDDSDSN() parameter of the DFSDFxxx member of the IMS PROCLIB data set.
- All of the defined RDDSs can be allocated and read.
- At least one of the RDDSs contains valid resource and descriptor definitions.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

Related reference:

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

Cold starting an IMS system that uses an RDDS

IMS systems that use dynamic resource definition to manage application programs, databases, Fast Path routing codes, and transactions can save resource definitions to a resource definition data set. During a cold start, IMS can use the definitions from the MODBLKS data set or the resource definition data set (RDDS) to create runtime resource definitions.

If your IMS system is not configured to import the resource definitions automatically at cold start, you can issue the IMPORT command after the restart process is complete. The IMPORT command, as well as all of the other type-2 dynamic resource definition commands, requires that the SCI and OM CSL components be started.

You can successfully start an IMS for the first time with DRD enabled and no resources defined yet. After the restart process completes, you can define resources dynamically by using the appropriate CREATE commands.

Before you shut down a running IMS system that uses dynamic resource definition for a cold start, you must export the resource and descriptor definitions that are currently defined in your online system to the RDDS. If you do not, any changes made to the resource definitions in the online system will be lost at the next cold start or emergency restart with the COLDSYS parameter.


During a cold start, the partition status or access state will be copied from the HALDB master. If you export the HALDB master status to the RDDS or the repository, the status of the master and its partitions will be obtained from the RDDS or the repository.

If an error occurs while trying to access an RDDS during the automatic import process, automatic import terminates abnormally. IMS either continues with the cold start process without any runtime resource definitions defined, or IMS stops the cold start process. If an error occurs during automatic import processing because of an invalid resource or descriptor definition, one of the following situations occurs:


- Automatic import processing terminates abnormally, and IMS cold start terminates abnormally with a U3397abend.
- Automatic import processing continues, and the resource in error is marked and given a not-initiated status (NOTINIT).

The action IMS takes depends on the values specified for the RDDSERR and the IMPORTERR parameters in the DFSDFxxx member of the IMS PROCLIB data set.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

Related tasks:

 Backing out DRD command-related changes made to your online system (System Definition)

Cold starting an IMS system that uses the IMSRSC repository

If IMS is using the IMSRSC repository, and AUTOIMPORT=AUTO or REPO is defined in the DYNAMIC RESOURCES section of the DFSDFxxx member of the IMS PROCLIB data set, the stored resource definitions are read from the repository during IMS cold start if the repository contains the resource definitions for the IMS that is cold starting.

During a cold start, the partition status or access state will be copied from the HALDB master. If you export the HALDB master status to the RDDS or the repository, the status of the master and its partitions will be obtained from the RDDS or the repository.

The stored resource definitions from the repository are used to create the runtime resource definitions in the IMS system that is cold starting.

An IMS resource list in the repository is used during an IMS cold start to identify all the resource and descriptor definitions that are to be imported during the cold start.

The DFS3395I message is issued, indicating the stored resource definitions are automatically imported from the repository. The DFS3396I message with a return code is issued when the automatic import from the repository is completed.

An RM error trying to read the stored resource definitions from repository results in the DFS4401E message with the RM return and reason code.

If REPOERR=ABORT is specified in the DFSDFxxx member, IMS terminates with a 3368 abend. The DFS3397E message with a return code of 16 is issued before the abend.


If REPOERR=NOIMPORT is specified in the DFSDFxxx member, IMS continues processing with no resources imported.

If there are no resource definitions in the repository for the IMS, the stored resource definitions from the resource definition data set (RDDS) or MODBLKS data set are read if AUTOIMPORT=AUTO is specified. If the RDDS or MODBLKS data set is not defined or is empty, or AUTOIMPORT=REPO is specified, then IMS will come up with no resources.

If there are no resources for the IMS in the repository, IMS issues the DFS4405W message, indicating that the repository is empty.

At the end of cold start processing, you must issue the EXPORT DEFN TARGET(REPO) command to export the runtime resource definitions imported from the RDDS or MODBLKS data set to the repository.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

 Resource lists for the IMSRSC repository (System Definition)

Related tasks:

“Recovering IMSRSC repository data sets when no IMS systems or RM systems are active” on page 40

Related reference:

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

 EXPORT command (Commands)

Related information:

 DFS4404I (Messages and Codes)

 DFS3395I (Messages and Codes)

 DFS3396I (Messages and Codes)

 DFS4401E (Messages and Codes)

 DFS3397E (Messages and Codes)

 DFS4405W (Messages and Codes)

 DFS3374W (Messages and Codes)

Warm starting IMS (or normal restart)

A warm start (or normal restart) is the most common way of initializing IMS, and is also the recommended way of restarting IMS after a controlled shutdown.

Use the **/NRESTART** command to perform a warm start or a normal restart. You can use the **FORMAT** parameter to format various system data sets (something you might want to do if you have reallocated any of the data sets). If you are running an IMS DB/DC or DCCTL environment, and you request reformatting of any of the message queue data sets, you must specify **BLDQ** on the **/NRESTART** command if IMS was previously shut down with either the **DUMPQ** or **PURGE** keyword specified.

If the MTO needs to dequeue transactions after a warm start, use the **/DEQUEUE SUSPEND** command.

If you want to reload MSDBs during restart, use the **MSDBLOAD** keyword on the **/NRESTART** command.

Related concepts:

 Resource lists for the IMSRSC repository (System Definition)

Emergency starting IMS

After IMS terminates without a controlled shutdown, you must perform an emergency restart. The one exception is if an emergency restart fails; in this case, you must perform a cold start with the **/ERESTART COLDSYS** command.

For an emergency restart, IMS restarts from the point of failure. Before restarting IMS, the MTO must know whether system data sets must be reallocated and

reformatted during restart. The following table shows symptoms (a message or abend number) that indicate whether a system data set needs reallocation and reformatting.

Table 16. Symptoms indicating system data set reallocation is needed. Key: N/A=Not Applicable

System data set	Write error	Read error	Data set full	Open error
WADS	DFS0414I	DFS739I	N/A	DFS3256I
Message queue	U0759	U0759	U0758	DFS986A
Restart data set (RDS)	DFS3127I	U0970	N/A	U0970
MSDB checkpoint or MSDB dump	DFS2718I, DFS2722I	N/A	N/A	DFS2681I, DFS2713I, DFS2714I

The following situations cause the abends shown in the previous table.

- U0758** An IMS internally generated /CHECKPOINT DUMPQ command failed after messages DFS206I, DFS207I, or DFS208I because one of the message queue data sets is full.
- U0759** An error occurred while attempting to read or write the message queue data sets.
- U0970** The restart data set (RDS) or checkpoint table was not initialized, or restart was unable to open the input log.

If the WADS has not previously been formatted, you should format it by specifying the FORMAT WADS keywords on the /ERESTART command. If you want to reformat any of the message queue data sets, specify the BLDQ keyword. If no checkpoint specifying either DUMPQ or SNAPQ has been taken since the last cold start, specify CHECKPOINT 0 on the /ERESTART command.

During Extended Specified Task Abnormal Exit (ESTAE) routine processing, IMS issues message DFS0617I if the RDS buffers have been purged and the RDS has been successfully closed. However, if you do not receive this message, you do not necessarily have to reformat or reallocate the RDS.

Requirement: You should reformat the RDS if you receive message DFS3127I or abend U0970, either of which indicate a read or write problem with the RDS. However, if the error has nothing to do with the RDS, you should not reformat the RDS. If errors occur on both MSDB checkpoint data sets:

1. Issue the /ERESTART command.
2. Issue the /DBDUMP DB or UPDATE DB STOP(UPDATES) command for the MSDBs.
3. Shut down the system (in a controlled way).
4. Run the MSDB Dump Recovery utility.
5. Restart the system using an /NRESTART command, including the MSDBLOAD keyword so that IMS can reload the MSDBs.

Emergency restarts and dynamic resource definition

IMS systems that use dynamic resource definition automatically recover the resources and runtime descriptor definitions during emergency restart by

processing the checkpoint records and X'22' log records during an emergency restart, unless the COLDSYS parameter is specified for the /ERESTART command.

If the COLDSYS parameter is specified, the resource definitions are read from the resource definition data set (RDDS) or the IMSRSC repository.

Related concepts:

➡ Overview of the IMSRSC repository (System Definition)
“Emergency restart” on page 97

➡ Resource lists for the IMSRSC repository (System Definition)

Related reference:

➡ /ERESTART command (Commands)

SLDS input to warm start and emergency restart

SLDS input might be required to perform a warm start or emergency restart. When the SLDS is required for restart, IMS dynamically allocates it. IMS also deallocates any data sets previously allocated with the dynamic allocation macro (DFSMDA) with IMSLOGR or IMLOGR2 as the DD name for the TYPE=SLDS statement.

DBRC provides the data set name and volume information for the dynamically allocated SLDSs. However, DBRC does not provide the unit type information, so a DFSMDA member with the name of IMSLOGR must provide the unit type.

In some cases, you may want to force IMS to read from the SLDS instead of the OLDS during restart—for example, to prevent IMS from reading an OLDS that you know is invalid or has been corrupted. To force IMS to read from the SLDS, you need to delete the OLDS entry from the PRIOLDS and SECOLDS records in the RECON data set. To delete an OLDS entry from the RECON data set, use the DELETE.LOG OLDS(*ddname*) command, where *ddname* is the name of the primary OLDS.

Attention: When image copy, change accumulation, and log data sets are listed in the RECON status record as being under the control of a catalog management system (CA | IC | LOG DATA SETS CATALOGED=YES), the restart of IMS will fail if the SLDS data sets are not actually cataloged. The volume serial number (VOLSER) is not passed back to IMS for data set allocation when the data sets are not cataloged. Likewise, if the data sets are SMS-managed and not cataloged (CA | IC | LOG DATA SETS CATALOGED=NO), the restart of IMS might fail if SMS migrated the SLDS data sets to a different volume serial number. In this case, the volume serial number will be passed back to IMS but it does not match the volume serial number where the data set resides. Therefore, it is important to set the CATDS keyword on the CHANGE.RECON or INIT.RECON command correctly.

Specifying security options at startup

Depending on your system definition, the MTO might be able to specify whether IMS activates various security options during startup. Each restart command options affects security differently, and can be used with the /NRESTART command or the /ERESTART COLDSYS command.

The following table shows the restart command options that affect security.

Table 17. Security options on the /NRESTART or /ERE COLDSYS command.

System security macro name	Macro options	JCL EXEC parameters	/NRE or /ERE command input	System security level after restart
SECLVL	=NOSIGN	SGN=N	USER	Signon active. Single signon per user for static users. ¹
			TRANAUTH	Signon and transaction authorization active.
			CMDAUTH	Signon, static, and ETO command authorization active.
			CMDAUTHE	Signon and ETO command authorization active.
			MULTSIGN	Signon active. Multiple signon per user (static users). ¹
	=SIGNON	SGN=Y	NOUSER	Signon not active.
	=TRANAUTH	TRN=Y	NOUSER	Signon and transaction authorization not active.
			NOTRANAUTH	Transaction authorization not active. Signon active.
	=NOTRAN	TRN=N	TRANAUTH	Signon and transaction authorization active.
	=FORCTAN	TRN=F		Signon and transaction authorization enforced.
TYPE	=RACFCMD	RCF=C		
			NOUSER	Signon and command authorization for ETO not active.
			NOCMDAUTHE	Command authorization for ETO not active.
			NOCMDAUTH	Command authorization for ETO/static users not active.
		RCF=S	NOUSER	Signon for ETO and static users not active.
			NOCMDAUTHE	Command authorization for ETO/static users not active.
			NOCMDAUTH	Command authorization for ETO/static users not active.
			CMDAUTHE	Command authorization for ETO active.
			CMDAUTH	Command authorization for static users not active.
	=NORACFCM	RCF=N		
			CMDAUTHE	Signon and command authorization for ETO active.
			CMDAUTH	Signon and command authorization for ETO and static users active.

Note:

1. This also applies to dynamic users signing on to ETO terminals where the user structure name is different from the user ID.

Starting the IRLM

You use the internal resource lock manager (IRLM) to perform block-level or sysplex data sharing. If you are running an IMS DB/DC or DBCTL environment, start the IRLM from the z/OS system console using the z/OS START command.

You must start IRLM subsystems before starting the IMS batch or online subsystem. The IMS online subsystem identifies itself to the IRLM subsystem during IMS initialization or restart.

Related concepts:

 Defining an IRLM for sysplex data sharing (System Administration)

Starting the CQS

In an IMS shared-queues environment, IMS starts the CQS subsystem during IMS startup or restart if the CQS subsystem is not already running. You can use the z/OS START command to start the CQS subsystem if you want to ensure CQS is running before you start IMS.

After CQS is running, the IMS online subsystem identifies itself to the CQS subsystem during IMS initialization or restart.

If you register CQS with the z/OS Automatic Restart Manager, you do not have to restart CQS after a CQS failure. Because CQS and IMS are separate subsystems, you do not have to restart CQS after an IMS failure.

You can start CQS in one of three ways:

- As a z/OS task, using the z/OS START command
- As a z/OS batch job
- As a client task — Some clients (such as IMS) automatically start CQS, when appropriate


Example: If IMS is the client, define the CQS name in the DFSSQxxx or the DFSDFxxx member in the IMS PROCLIB data set and specify the CQS name on the SHAREDQ parameter of the IMS procedure. IMS does not start CQS if the CQS is needed only to manage a resource structure.

A CQS that supports only a resource structure must be started manually because IMS does not start this CQS.

Related tasks:

 [Preparing to start the CQS address space \(System Definition\)](#)

Related reference:

 [DCC procedure \(System Definition\)](#)

 [DFSSQxxx member of the IMS PROCLIB data set \(System Definition\)](#)

 [Defining Sysplex Policies](#)

Starting dependent regions

IMS dependent regions are z/OS regions that are separate from the IMS control region, and that are logically dependent upon the IMS control region.

The following types of programs execute in dependent regions:

- IMS batch message programs (BMPs) and message processing programs (MPPs)
- IMS Fast Path (IFP) message-driven application programs
- Fast Path DEDB online utility programs
- Java™ batch processing (JBP) and Java message processing (JMP) applications
- Coordinator controller (CCTL) application programs (in the IMS DBCTL environment)

You can start IMS dependent regions from the z/OS system console or the IMS master terminal. From the z/OS system console, use the z/OS START command to start a reader procedure (usually IMSRDR) to read the JCL for the region. From the IMS master terminal, use the IMS /START REGION command to start a region. If you do not specify a procedure name on the START or /START REGION command, IMS uses the IMSMSG procedure.

You should not start IMS message regions and Fast Path regions with automatic priority group (APG) initiators because IMS cannot control the dispatching priority

of an APG-initiated job or uniquely identify such jobs. You might need to start and stop z/OS initiators for IMS message regions and Fast Path regions that have APG initiators.

IMS can dynamically allocate up to 255 regions. The total number of regions includes the number of regions allocated during IMS system definition plus the number of dynamically allocated regions or CCTL threads.

Message processing regions

One of the jobs in the IMS.JOBS data set contains the message region JCL. Parameters on the JCL EXEC statement define the transaction classes that are eligible for processing in the region. You can define more than one set of JCL if necessary.

The default job name is IMSMSG. After you start a message region, transactions associated with the transaction classes supported by the region can be scheduled. You can change these classes using the /ASSIGN TRAN or UPDATE TRAN SET(CLASS(*new_class_number*)) command after you have started the region.

Batch message processing regions

You can start BMP regions from the z/OS console or the IMS master terminal. Generally, you use the z/OS console and a z/OS reader procedure to start a BMP, unless the JCL for the job (using procedure IMSBATCH) is in the IMS.JOBS data set.

A JCL EXEC parameter indicates that the job is a BMP.

Fast Path message-driven regions

You can start IMS Fast Path message-driven regions from the z/OS console or the IMS master terminal. Generally, you use the z/OS console and a z/OS reader procedure to start a Fast Path message-driven region, unless the JCL for the job (using procedure IMSFP) is in the IMS.JOBS data set.

A JCL EXEC parameter indicates that the job is a Fast Path message-driven application program.

Java dependent regions

You can start two types of Java message-driven dependent regions, either a JMP or a JBP region.

The JMP region is similar to an MPP region and is started in the same manner as an MPP region. The default job name is IMSJMP.

The JBP region is similar to a BMP region and is started in the same manner as a BMP region. The default job name is IMSJBP.

DEDB online utility regions

You can start all DEDB online utilities using procedure FPUTIL. Specify the utility in the TYPE operand in the control data set.

CCTL regions

The CCTL region is not part of an IMS subsystem. See your CCTL documentation for information on how to start the CCTL.

Restarting CQS

After CQS completes the structure initialization, it continues with the restart. CQS can do either a cold restart or a warm restart. Restarting CQS affects only the units of work that this CQS manages. Restarting does not back out or restore any units of work owned by another CQS.

When you have completed restarting all structure pairs for a particular CQS, the CQS ready message is issued (CQS0020I).

Since changes to resource structures are not checkpointed or logged, restarting CQS does not affect units of work for resource structures.

The frequency of system checkpoints affects restart. CQS must read more log records when checkpoints are infrequent than when checkpoints occur more often. Also, the amount of logging that one CQS performs can affect another CQS during restart. All CQSs write to the same log, so a CQS restarting must read all log records written by all CQSs.

CQS takes an initial system checkpoint at the end of a restart.

CQS warm start

During a warm start, CQS reads the log records from the last system checkpoint, restores the environment for committed data objects, and backs out uncommitted data objects to prepare CQS to regain synchronization with the client and resume processing. Normally, a CQS warm start is automatic and you do not need to take any action.

When CQS warm starts, it reads the checkpoint data set to find the log token representing the last system checkpoint. When CQS finds this log token, it initiates a warm start. If CQS fails to find this log token in the checkpoint data set, it reads the log token from the structure. If CQS finds the log token, CQS issues WTOR CQS0031A to allow you to confirm the use of this token. At this point, you can do one of the following:

- Confirm the log token
- Cold start CQS
- Cancel CQS
- Specify a new log token

If you specify a new log token and CQS fails to find this log token, CQS issues WTOR CQS0032A. At this point, you can take one of the following actions:

- Cold start CQS
- Cancel CQS
- Specify a new log token

Sometimes, CQS purges log records that are required for restart. CQS purges log records in the following situations:

- During a structure checkpoint
- When the log becomes full and no more data sets are available for logging

Important: CQS might not have any log records if it is only managing resource structures.

Recommendation: If a CQS does not accept a log token during CQS restart, cold start the CQS. In cases where multiple CQSs are running, it is possible that log records for a CQS that previously failed and was not restarted are purged while another CQS performs a structure checkpoint.

Related concepts:

 Using CQS structure checkpoint (System Administration)

CQS cold start

When CQS restarts after a structure cold start, CQS cold start processing is automatic. You do not need to take any action. CQS takes a system checkpoint and then CQS restart is complete.

When CQS cold starts after a structure warm start or a structure recovery, CQS reads the structure to find unresolved work. CQS backs out requests to move data but completes requests to delete data. CQS performs a system checkpoint and restart is complete.

No log records are read or processed when CQS is cold started.

To cold start CQS, you must:

1. Scratch the CQS system checkpoint data set for the structure.
2. Reply COLD to the CQS0031A WTOR.

CQS registration with the z/OS Automatic Restart Manager

CQS, if requested, can register with the z/OS Automatic Restart Manager (ARM). The ARM is a z/OS recovery function that can improve the availability of started tasks. When a task fails or the system on which it is running fails, the ARM can restart the task without operator intervention.

Recommendation: Register with the z/OS ARM regardless of the types of structures that CQS is using.

To enable the ARM, you can specify ARMRST=Y in one of two ways:

- In the CQSIPxxx member of the IMS PROCLIB data set
- As an execution parameter

An abend table exists in module CQSARM10. The table lists the CQS ends abnormally for which the ARM does not restart CQS after the abend occurs. You can modify this table.

IBM provides policy defaults for automatic restart management. You can use these defaults or define an ARM policy to specify how CQS should be restarted. The policy can specify different actions to be taken when the system fails and when CQS fails. When ARM is enabled, CQS registers to ARM with an ARM element name of CQS + *cqsssn* + CQS. Use this ARM element name in the ARM policy to define the ARM policy for CQS.

cqsssn is the CQS name that can be defined either as a CQS execute parameter, or with the **SSN=** parameter in the CQSIPxxx member of the IMS PROCLIB data set. For example, if **SSN=CQSA**, then the ARM element name is CQSCQSACQS.

Related reference:

 [Using z/OS Automatic Restart Management](#)

Restarting CQS after CQS resource cleanup failures

If CQS abends and you receive message CQS0102E with module CQSRSM00, log records might be missing from the CQS log. Message CQS0102E indicates that a failure occurred during CQS resource cleanup. This failure might prevent CQS log records in internal buffers from being externalized to the CQS log.

If this situation occurs, perform one of the following actions:

- If the terminating CQS is the only CQS running for its set of structures, restart the CQS immediately.
- If other CQSs are running, either immediately restart the terminated CQS or initiate a structure checkpoint on one of the surviving CQSs.

Successfully restarting the failed CQS or taking a structure checkpoint is necessary to preserve the state of the data on the shared queues in the event that a structure rebuild is needed.

Restarting CQS queue structures

Restarting queue structures involves restarting a structure pair. The structure restart function ensures that the data in the structure is correct. Structure restart deals only with the status of data in a specific coupling facility structure, not units of work specific to a given CQS.

Before CQS can restart, CQS must recover each queue structure defined to CQS, if needed. CQS connects to both the primary and the overflow queue structures for each structure pair defined. CQS determines whether the structures need to be warm or cold started and performs the necessary recovery functions. If more than one structure pair is defined to CQS, one structure pair can be warm started and another can be cold started.

CQS structure allocation

A structure is allocated the first time a CQS connects to it and remains allocated until you explicitly delete it using the z/OS SETXCF command.

When a CQS connects to a structure, the structure might be empty; that is, it might contain no data, might contain only CQS control information, or it might contain client data. The structure can be empty if:

- This is the first time a CQS is accessing the structure.
- You scratched the structure to perform a structure cold start.
- A structure failure occurred and the structure must be recovered.

CQS structure warm start

CQS warm starts both a primary structure and its overflow structure if either the primary structure contains data, or one SRDS contains valid structure checkpoint data and the CQS log contains valid data.

During a structure warm start, CQS determines the status of the structure and initiates a structure recovery if necessary. If a structure recovery is needed, CQS

allocates the structure and repopulates it from either the CQS log and the SRDS, which contains valid client data from a previous checkpoint, or from the CQS log by itself.

After a structure warm start has completed, CQS determines whether a future recovery is possible based on the status of the structure recovery data sets and the log stream. If the primary structure contains client data, but neither the SRDS nor the log can be used for future recovery, CQS issues a CQS0009W message.

Recommendation: If CQS issues a CQS0009W message, initiate a structure checkpoint as soon as possible. If a structure checkpoint does not complete successfully and the structure fails, CQS cannot recover the structure.

After a structure warm start, CQS can be cold started or warm started. If the log records needed for the CQS restart have been deleted, you might have to cold start the CQS.

CQS structure cold start

Common Queue Server (CQS) cold starts both a primary structure and its overflow structure if the primary structure is empty and both of the structure recovery data sets are empty.

During a structure cold start, CQS deletes:

- The overflow structure
- All the log records in the log stream for the structure

To cold start a structure, you must:

1. Shut down any CQS systems that are connected to the structure. For example, issue the following command, where *cqs_name* is the CQS subsystem name.

```
P cqs_name
```

The response will be shown as follows:

```
CQS0021I CQS SHUTDOWN COMPLETE
CQS0101I CQS CLEANUP SUCCESSFUL
```

2. Delete the primary and overflow structures on the coupling facility. For example, issue the following command, where *structure_name* is the structure name.

```
SETXCF FORCE,STRUCTURE,STRNAME=structure_name
```

The response will be shown as follows:

```
IXC353I THE SETXCF FORCE REQUEST FOR STRUCTURE
structure_name WAS COMPLETED:
STRUCTURE WAS DELETED
```

3. Scratch both structure recovery data sets (SRDS 1 and 2). For example, use the Access Method Services program IDCAMS to scratch the data sets.
4. Reallocate both structure recovery data sets (SRDS 1 and 2). For example, use the Access Method Services program IDCAMS to reallocate the data sets.
5. Restart a CQS. CQS initialization will reallocate the structure. For example, use the CQS startup procedure.
6. Verify that the structure is empty. For example,
 - Issue the following z/OS command, where *structure_name* is the structure name.

```
D XCF,STR,STRNAME=structure_name
```

The response for the z/OS command will be shown as follows:

```

...
STORAGE INCREMENT SIZE: 256 K
ENTRIES: IN-USE:          3 TOTAL:      4238,    0% FULL
ELEMENTS: IN-USE:          3 TOTAL:      4238,    0% FULL
EMCS:     IN-USE:          0 TOTAL:      6634,    0% FULL
...

```

When structure cold start completes, CQS automatically performs cold start restart processing.

Restarting resource structures

You can cold start a resource structure by shutting down any CQS systems that connect to the resource structure, deleting the resource structure on the coupling facility, and restarting a CQS to reallocate the resource structure.

To cold start a resource structure, you must:

1. Shut down any CQS systems that connect to the resource structure. For example, issue the following command, where *cqs_name* is the CQS subsystem name.

```
P cqs_name
```

The response will be shown as follows:

```

CQS0021I CQS SHUTDOWN COMPLETE
CQS0101I CQS CLEANUP SUCCESSFUL

```

2. Delete the resource structure on the coupling facility. For example, issue the following command, where *structure_name* is the resource structure name.

```
SETXCF FORCE,STRUCTURE,STRNAME=structure_name
```

The response will be shown as follows:

```

IXC353I THE SETXCF FORCE REQUEST FOR STRUCTURE
structure_name WAS COMPLETED:
STRUCTURE WAS DELETED

```

3. Restart a CQS. CQS initialization will reallocate the resource structure. For example, use the CQS startup procedure.
4. Verify that the resource structure is empty. For example, issue the following command, where *structure_name* is the resource structure name.

```
D XCF,STR,STRNAME=structure_name
```

The response will be shown as follows:

```

...
STORAGE INCREMENT SIZE: 256 K
ENTRIES: IN-USE:          3 TOTAL:      4238,    0% FULL
ELEMENTS: IN-USE:          3 TOTAL:      4238,    0% FULL
EMCS:     IN-USE:          0 TOTAL:      6634,    0% FULL
...

```

When structure cold start completes, CQS automatically performs cold start restart processing.

CQS structure recovery for restarting

A structure might need to be recovered if the structure is empty or if it contains only Common Queue Server (CQS) control information. Data from the last structure checkpoint and the z/OS log stream are used to recover a structure.

The structure is first repopulated with data objects from the structure recovery data sets. CQS then reads the log, starting at the time of the structure checkpoint, and updates the structure with changes that occurred after the structure checkpoint.

If the primary structure is empty and neither SRDS contains valid structure checkpoint data, CQS determines whether it can use just the CQS log for recovery. If the first log record in the log stream is the “Beginning of Log” record, the log stream contains all the log records required for recovery, and CQS can use the log records to complete the structure recovery.

If CQS finds that a previous structure rebuild did not complete successfully, it initiates another rebuild.

If the primary structure contains only CQS control information and the first CQS that connected to the structure (the one that allocated the structure) cannot determine whether a rebuild is needed, CQS initiates a rebuild if either SRDS is valid or if all the log records are available.

If neither SRDS is valid and the log records have been deleted by a previous structure checkpoint, CQS cannot rebuild the structure. When this happens and a rebuild is necessary, CQS issues write-to-operator-with-reply (WTOR) CQS0034A message asking what you want to do. You can cold start the structure or cancel this CQS.

Related concepts:

 CQS structure recovery (System Administration)

Starting Transaction Manager

For remote terminal operators to enter transactions, you must initialize the IMS Transaction Manager (for the IMS DB/DC and DCCTL environments).

The procedures are different for the different access methods and protocols:

- Virtual Telecommunications Access Method (VTAM)
- Communications with devices
- Advanced Program-to-Program Communications (APPC/MVS)
- Intersystem Communication (ISC)

Connecting to the VTAM network

You use the /START DC command to establish communications between IMS and VTAM. The /START DC command tells VTAM to pass queued logon requests to IMS for all terminals defined to VTAM as belonging to IMS.

Before establishing a session with IMS for any terminal, verify that the following conditions exist:

- VTAM and the Network Control Program (NCP) are active.
- Controllers and logical units are active. If they are not automatically activated by VTAM, the VTAM network operator can use the VTAM VARY command to activate them.
- Controllers are powered on, appropriately configured, initialized, and activated for VTAM and NCP.

Recommendation: Start VTAM before starting the IMS control region.

You also need to use the /START DC command in either of the following situations:

- If VTAM is stopped and restarted while the IMS control region is running

- If you terminate communications between IMS and VTAM using the /STOP DC command

The /START DC command activates the following processes:

- Initiating IMS Transaction Manager processing
- Opening the VTAM access method control block (ACB)
- Enabling the IMS VTAM logon exit routine
- Opening the primary and secondary VTAM master terminals if these terminals were not opened during IMS initialization (due to the specification of VACBOPN=DELAY in the DFSDCxxx IMS.PROCLIB member)

If VTAM is active when IMS starts, the IMS VTAM ACB is opened; otherwise, the ACB is opened by the /START DC command. Any logon requests received by VTAM before the IMS /START DC command, but after the IMS VTAM ACB has been opened, are enqueued by VTAM until the /START DC command completes. If the queuing of logon requests by VTAM causes operational problems, the specification of VACOPN=DELAY in the DFSDCxxx PROCLIB member can be used to delay the opening of the VTAM ACB. When the master terminals are generated to use VTAM, the master terminal operator will be unable to enter either the /NRE or /ERE commands (including /ERE BACKUP) or the /START DC command. These commands can be issued using automatic restart (AUTO=Y) or automated operations.

Before you can establish a session between IMS and a logical unit or node, the logical unit or node must be active, connected, enabled, and available to VTAM and started in IMS. You can activate logical units or nodes when VTAM starts using VTAM start options, or a VTAM operator can activate them using the VTAM VARY command.

Normally after you start IMS, all VTAM terminals have a started status. If you stop a VTAM terminal with an IMS /STOP command, the terminal remains stopped after an IMS warm start or emergency restart. You must use the /START NODE command before issuing the /OPNDST (open destination) command to reactivate the terminal.

You must establish sessions between terminals and IMS before anyone can use them to transmit data. You can initiate sessions in several ways:

- VTAM can automatically log them on (based on the VTAM terminal definition).
- The end user can log on.
- You can issue an IMS /OPNDST command (if they are not automatically logged on).
- You can issue a VTAM VARY command.

Connecting to devices

When you start IMS, IMS automatically starts communications between IMS and devices such as spools, readers, printers, punches, tapes, and disks. After you start IMS, the telecommunication lines have a stopped status.

You can start the telecommunication lines, and devices if necessary, using the /START or /RSTART commands. Normally, you use the /START command after an IMS cold start and the /RSTART command after a warm start. The /RSTART command starts the devices so that they have the same mode they had when they were stopped at the previous system shutdown.

If you stop a device using the /STOP command, the device remains stopped after an IMS warm start or emergency restart.

Connecting to APPC/MVS

Before an LU 6.2 conversation can be established with IMS, the following must be active: APPC/MVS, VTAM, and NCP.

Recommendation: Start APPC/MVS and VTAM before starting IMS.

When you specify APPC=Y in the IMS procedure, IMS automatically starts communication between IMS and APPC/MVS. If you specify APPC=N in the IMS procedure, you can use the /START APPC command to establish communications.

You also need to use the /START APPC command in either of the following situations:

- If APPC/MVS is stopped and restarted while IMS is running
- If you terminate communication between IMS and APPC/MVS by using a /STOP APPC command

Connecting ISC sessions from CICS to IMS

You can connect CICS® ISC sessions to IMS in two distinct ways. You can initiate the session explicitly or the CICS operator can initiate the session through a command.

You can connect CICS ISC sessions to IMS in one of the following ways:

- You can initiate the session explicitly by specifying AUTOCONNECT YES on the DEFINE CONNECTION command in the CICS system definition (CSD) utility, or by specifying **CONNECT=**AUTO in the DFHTCT **TYPE=**TERMINAL program. When CICS is started, it attempts to establish all sessions for which AUTOCONNECT is specified.

In order for session to be initiated, IMS must be active when CICS is started.

- The CICS operator can initiate a session by entering the following command:
CEMT SET TERMINAL(*termid*) [ACQUIRED | COLDACQ | RELEASED]
where *termid* is the four-character session name (defined on the DEFINE SESSIONS command in the CSD utility), or the TRMIDNT in the DFHTCT TYPE=TERMINAL program. The following describes the effects of the keywords for the CEMT SET TERMINAL command:

ACQUIRED

Specifies normal resynchronization with the IMS half-session. CICS is in session with the logical unit represented by the terminal.

COLDACQ

Specifies that no resynchronization is done.

COLDACQ is a special form of ACQUIRED, in which no resynchronization is required. If the previous session ended abnormally, using COLDACQ overrides CICS integrity control, which can lead to integrity problems. Also, check the CSMT log for an activity keypoint (similar to an IMS checkpoint) after the restart of a session following a CICS failure. If there is no activity keypoint, reissue the CEMT SET TERMINAL(*termid*) COLDACQ command after the next emergency restart.

RELEASED

Specifies that CICS is not in session with the logical unit represented by the terminal.

Specifying RELEASED for a terminal that has an active session causes the session to be terminated. Running transactions are allowed to finish unless you also specify PURGE or FORCEPURGE.

The display area of the CEMT [INQUIRE | SET] TERMINAL command shows the status of the CICS half-session. If the session initiates successfully, the status changes from REL (released) to ACQ (acquired). If the session does not initiate successfully, CICS writes an error message to the transient data destination CSMT.

- An application program can initiate a session implicitly by using the ALLOCATE command. A program can use the ALLOCATE command only for the SEND/RECEIVE interface.

Recommendation: Use the ALLOCATE SYSID form of the command because it allows CICS to select an available session.

If a session is not immediately available, control is returned to the application program if either of the following occur:

- The program issues a HANDLE CONDITION command.
- The program specifies NOQUEUE on the ALLOCATE command.

Otherwise, the command is queued until a session becomes available.

Each of the following conditions causes a session to be “not immediately available”.

- All sessions to the specified system (or the specified session) are in use.
 - The only available sessions are not bound.
 - The only available sessions are contention losers.
- CICS automatic task initiation (ATI) can initiate a session.

To determine the primary and secondary RU sizes, CICS uses the values for the RECEIVESIZE and SENDSIZE parameters on the DEFINE CONNECTION | SESSION command in the CSD utility or the values for the RUSIZE and BUFFER parameters on the DFHTCT TYPE=TERMINAL | SYSTEM program.

Related concepts:

“Terminating an ISC session from CICS” on page 120

Related reference:

 CICS Transaction Server for z/OS

Starting IMS systems and global command status

Global command status is the status of a shared IMS resource as set by online commands and as seen by all IMSplex members that have the resource defined.

If you are operating in an IMSplex environment that specifies PLEXPARM(GSTSTRAN(Y), GSTSDB(Y) or GSTSAREA(Y)), you must understand how the status of IMS resources as set by online commands is maintained at both the global and the local level, as well as how resource status is recognized by IMSplex members as individual IMS systems are started and stopped across the IMSplex.

For example, if the Resource Manager (RM) maintains global command status information for a database and the database is stopped by a command, all members of the IMSplex, whether they were running at the time the command was issued or not, see that the database is stopped. That same database can have a local status in an IMS system that is different than the global command status maintained in RM. Stopping a database globally while allowing access locally can be useful for preventing updates to the database from other IMSplex members while a local batch program processes the database.

RM maintains global command status for IMS resources by resource type. The maintenance of global command status is enabled for a resource type by specifying the PLEXPARM= parameter in the DFSDFxxx or DFSCGxxx PROCLIB member during IMS initialization.

You can assign global status to a resource dynamically by using one of the following commands:

- UPDATE IMS SET(PLEXPARM(GSTSDB(Y))) identifies that global database status is to be maintained on database commands.
- UPDATE IMS SET(PLEXPARM(GSTSAREA(Y))) identifies that global area status is to be maintained on area commands.
- UPDATE IMS SET(PLEXPARM(GSTSTRAN(Y))) identifies that global transaction status is to be maintained on transaction commands.

If global status is not explicitly specified for a resource type, RM does not maintain global status.

When RM maintains global status for a resource type, any commands issued against specific instances of the resource change the global status of the resource instance in RM. You can limit the affect of commands on the global status of a resources by using the SCOPE() parameter with the appropriate type-2 command:

- The SCOPE(ACTIVE) parameter limits the scope of commands to the active IMS systems that the command is routed to, and no global information is changed.
- The SCOPE(ALL) parameter specifies the scope of commands to the active IMS systems that the command is routed to and updates the RM status if applicable.

If a command is issued against a resource with global status before an IMS system is started, the starting IMS system can read the global status of resources in RM and apply them locally to the IMS system if the resources are valid.

The first IMS system to start in an IMSplex sets the PLEXPARM values that it knows about. IMS systems that start in the IMSplex subsequent to the first, use the PLEXPARM status definitions already in RM.

When an IMS system restarts and global status exists for the resource type in RM, the type of start determines which status takes precedence:

- For cold starts, the global status of databases, partitions, areas, and transactions takes precedence.
- For warm starts, the restarting IMS system first applies the local status saved in the log records and then applies the global status stored in RM only if it changed while the IMS system was down.

Before attempting to restart IMS again, complete one of the following actions:

1. Cancel IMS or shutdown IMS after restart completes. After the error condition reported in the DFS3308 message is fixed, restart IMS. IMS will be synchronized with RM status, and no additional action is required.
2. Allow the IMS restart to complete, then take the following actions:
 - Issue the appropriate UPDATE IMS SET(PLEXPARM) command to enable global status at this IMS.
 - Issue the appropriate QRY TRAN|DB|AREA SHOW(GLOBAL) commands to see the resources with global status.
 - Issue appropriate command with the SCOPE(ACTIVE) parameter to ensure that the IMS is synchronized with RM status.

RM stores global status information in its resource structure. If the resource structure should fail, RM repopulates the resource structure by having each active IMS system in the IMSplex update the resource structure with its local resources that have global status. Any global status for resources owned by an IMS system that is not active is not restored to the resource structure.

Related concepts:

“IMS restart and global resource status in an IMSplex that uses RM” on page 99

Restarting a component of IMS

After performing the necessary recovery, you can reactivate the particular component that you shut down or detached. For example, if you take a database offline because of an I/O error and then recover it, you can again make it available to applications.

Use the appropriate IMS command to restart the component. Most often, this command is some form of the /START or UPDATE command. For example, use /START DATABASE or UPDATE DB START(ACCESS) to restart a database or /START PROGRAM to restart an application program.

A list of restart commands and their corresponding shutdown commands is shown in the following table:

Table 18. Commands for restarting components of IMS

Restart command	Corresponding shutdown command	Affected components
/OPNDST	/CLSDST	VTAM session
/START	/STOP /IDLE /DBDUMP /DBRECOVERY	Areas, databases, VTAM logons, lines, logical terminals, VTAM nodes, OLDS, programs, regions, external subsystem connection, transactions, users
/RSTART	/IDLE	Lines, physical terminals, logical links, VTAM nodes
/RELEASE	/HOLD	Conversations
/UNLOCK	/LOCK	Databases, logical terminals, nodes, programs, physical terminals, transactions

Table 18. Commands for restarting components of IMS (continued)

Restart command	Corresponding shutdown command	Affected components
UPDATE	UPDATE	Areas, data groups, databases, MSC resources, and transactions.

Restarting IMS

You can restart IMS after you have completed any necessary recovery. You can perform a normal, emergency, or emergency with OVERRIDE restart.

- *Normal restart* initializes the IMS subsystem without reference to any previous execution of IMS (a cold start), or restarts a system that was terminated using a checkpoint command (a warm start).
- *Emergency restart* initializes an IMS subsystem after a system failure. During an emergency restart, IMS does the following:
 - Closes the OLDS from the WADS
 - Resets each active transaction to its last sync point
 - Resets each active region (BMPs, and CCTL or ODBA threads) to its last sync point

You do not need to restart connected ODBA application programs or CCTLs. An IMS shutdown or failure simply disconnects from the ODBA application or CCTL and generally does not affect it.

- Restores the databases to the last sync point
- Restores local message queues to their condition at the time of failure

An IMS failure does not generally affect shared message queues.

- If dynamic resource definition is enabled in the IMS system, the runtime resource definitions are restored by processing the checkpoint and X'22' logs

You must manually restart regions and BMPs. However, when restarting a batch message processing region, IMS automatically determines the last BMP sync point if you specified LAST as an EXEC parameter. Otherwise, IMS indicates what the last BMP sync point was so that you can specify it during restart.

If you have an IMS DBCTL standby environment, you can send the /ERESTART command to the pre-initialized DBCTL subsystem. Doing this is faster than starting the IMS job and waiting for the DBCTL-ready message before sending the /ERESTART command.

- *Emergency restart with OVERRIDE* is necessary after a failure in which DBRC is unable to mark the SSYS record in the RECON data set as abnormally terminated, such as a power, CPC, z/OS, or DBRC failure.

Automatic restart reduces MTO intervention and makes restart faster. With automatic restart, IMS automatically chooses the appropriate restart command (either /NRESTART or /ERESTART). If restart processing abnormally terminates before initial checkpoint, the appropriate automatic restart command is the same type (either /NRE or /ERE) as the aborted restart. The operator does not enter a restart command. Specify automatic restart by including AUTO=Y in the JCL.

You can implement normal restart procedures easily, and generally do not need to worry about it. However, an operator error could adversely affect system integrity in the following situations:

- Specifying cold start after a /CHECKPOINT FREEZE | DUMPQ | PURGE command when some messages are not completely processed
- Failing to supply the proper security options

Restriction: Do not change the IMS operating environment between shutdown and normal or emergency restart. For example, if you change from a nonshared-queues environment to a shared-queues environment, or change coupling facility structure names, IMS abends during normal or emergency restart. If you change the operating environment, you should cold start the IMS subsystem; or you can use IMS online change (/MODIFY command) to make changes while IMS is running.

Cold start

A cold start is required only when you first initialize IMS or after changing the nucleus during a system definition. You must also cold start IMS after scratching and reallocating shared-message-queue structures if these structures contained messages.

Request a cold start using the /NRESTART CHECKPOINT 0 or /ERESTART COLDSYS command. When you specify the COLDSYS keyword, the /ERESTART command initiates a cold start, not an emergency restart.

In a non-shared-queues environment, a cold start assumes empty message queues: if any messages exist, IMS discards them. In a shared-queues environment, IMS does not discard messages on shared queues during a cold start. If any affinities exist for VTAM generic resource groups, IMS deletes them. IMS never uses information from a previous shutdown as input to a cold start.

Optionally, you can format the IMS message queue data sets, WADS, and restart data set (RDS) during a cold restart. During a cold start, IMS loads all control blocks from the system libraries.

In a shared-queues environment, any locked messages remain locked after the IMS cold start until IMS (or other Common Queue Server client) unlocks or deletes them.

IMS systems that use dynamic resource definition can store resource definitions for application programs, databases, Fast Path routing codes, and transactions in a resource definition data set. During a cold start, IMS can be set up to automatically import the stored resource and descriptor definitions. If your IMS system is not configured to automatically import the resource definitions, you can issue the IMPORT command after the restart process is complete. The IMPORT command, as well as all of the other type-2 dynamic resource definition commands, require that the SCI and OM CSL components are started.

In order to recover changes to the runtime resource definitions across a cold start, you must export the resource definitions to a resource definition data set while IMS is up. If you do not, any changes made to resources and descriptors in the online system might be lost.

During a cold start, the partition status or access state will be copied from the HALDB master. If you export the HALDB master status to the RDDS or the repository, the status of the master and its partitions will be obtained from the RDDS or the repository.

If an error occurs while trying to access a RDDS during the automatic import process, automatic import terminates abnormally. IMS either continues with the cold start process without any runtime resource definitions defined, or IMS stops the cold start process. If an error occurs during automatic import processing because of an invalid resource or descriptor definition, one of the following situations occurs:

- Automatic import processing terminates abnormally, and IMS cold start terminates abnormally with a U3397abend.
- Automatic import processing continues, and the resource in error is marked and given a not-initiated status (NOTINIT).

The action IMS takes depends on the values specified for the RDDSERR and the IMPORTERR parameters in the DFSDFxxx member.

Warm start

A warm start is the most common way to reinitialize IMS, and is the recommended way of restarting after a controlled shutdown of IMS. You should attempt a warm start only if IMS terminated in an orderly manner. You can specify a warm start by using the /NRESTART command.

Before starting IMS and entering the restart command, you must determine whether IMS should:

- Reload MSDBs
- Format specific system data sets during restart
- Activate specific security functions during restart

The following situations require you to reallocate and reformat specific system data sets after a controlled IMS shutdown:

- When a message queue data set is full (indicated by messages DFS206, DFS207, or DFS208) and IMS was shut down by an internally-generated /CHECKPOINT DUMPQ command.
- When a write error occurred on the restart data set (RDS) during checkpoint or restart (indicated by message DFS3127I).
- When you need to change the amount of space allocated to the WADS or a write error occurs on the WADS during normal operation. You should change the amount of space allocated to the WADS only after a controlled IMS shutdown. You can replace a WADS that encountered a write error during normal restart (/NRE) or emergency restart (/ERE). In any case, when IMS uses a new WADS, you must format it during restart.
- When read or write errors occur on one of the MSDB data sets during normal operation (indicated by messages DFS2718I or DFS2722I).

For other error situations that require reallocation and formatting of some or all of the system data sets, you must use the /ERE command.

During a warm start, IMS:

- Reestablishes the status of the control region using control blocks that were logged at termination
- Optionally restores local message queues if IMS logged them (you issued a /CHE DUMPQ | PURGE command)
- Restores in-doubt units of work so they can later be resolved

Unlike emergency restart, you might have to manually release any transactions on the suspend queue by:

- Re-executing the transaction by issuing a `/START TRAN trannname` command. This command releases a specific transaction.
- Issuing a `/DEQUEUE SUSPEND` command. This command releases all transactions on the suspend queue.

When using XRF with VTAM Multinode Persistent Session (MNPS), IMS operates in an XRF environment and must use two ACBs: the APPLID ACB and the MNPS ACB. During a controlled shutdown, IMS closes the MNPS ACB and the APPLID ACB. Closing the MNPS ACB terminates MNPS tracking.

After a normal restart that follows this controlled shutdown, VTAM no longer maintains persistent sessions. Therefore, when IMS opens the MNPS ACB, it does not need to recover any previous sessions.

Related concepts:

 XRF and VTAM generic resources (System Administration)

Emergency restart

IMS requires an emergency restart whenever it terminates without a controlled shutdown. When restarted, IMS restarts from the point of failure. You can initiate an emergency restart using the `/ERESTART` command.

With an emergency restart, IMS:

- Backs out DL/I in-flight units of work
- Applies committed but unwritten DEDB changes to the database
- Retains any in-doubt units of work for IMS subsystems connected to a CCTL or to Recoverable Resource Management Services (RRMS), the z/OS syncpoint manager.
- Resolves any unfinished External Subsystem Attach Facility (ESAF) units of work with external subsystems
- Releases locked messages in a shared-queues environment
- Uses checkpoint and X'22' log records to recover MODLBKS resources

Before starting IMS and entering the emergency restart command, you need to know the following:

- Whether system data sets must be reallocated and formatted during restart.
- Whether IMS has cleaned up its resources. You can check for message DFS627I or DFS627W to determine the previous ending status of the IMS resource termination manager. These messages tell you if IMS successfully completed resource cleanup processing. Successful resource cleanup guarantees IMS restart without requiring a z/OS restart.

After an IMS failure, you can use the ESAF indoubt notification exit routine DFSFIDN0 to determine in-doubt units of work during restart of the failed IMS. You can resolve the work outside of IMS.

During an emergency restart, IMS restores the system to the status it had at the last system checkpoint. Databases and areas reflect their content at the last sync point for each dependent region. The local message queues reflect their content at the time of the failure. And the system has just taken a new checkpoint. When you

restart the dependent regions, IMS reschedules the previously processing programs. IMS also releases any transactions on the suspend queue.

After an IMS failure when XRF is used with VTAM MNPS, VTAM continues to maintain persistent sessions until the timer specified by the PSTIMER= keyword (specified in DFSDCxxx IMS.PROCLIB member) expires. If there is no XRF takeover, then VTAM might continue to maintain active persistent sessions when the MNPS ACB is opened after emergency restart. If this occurs, IMS issues the command SETLOGON OPTCD=NPERSIST, which tells VTAM to drop all persistent sessions being maintained. This is consistent with emergency restart processing when the XRF USERVAR= keyword is used, which also has no session recovery. IMS then issues the command SETLOGON OPTCD=PERSIST to start the normal session persistence for all new sessions on the MNPS ACB.

If an emergency restart fails, you do not have to cold start the entire IMS subsystem. The following table shows the various emergency restart cold start commands and how they affect an IMS system.

Table 19. Effects of emergency restart commands on an IMS subsystem

Emergency restart command	Effect on the IMS DB or DBCTL subsystem	Effect on the IMS TM or DCCTL subsystem
/ERESTART COLDBASE	Cold restart	Emergency restart
/ERESTART COLDCOMM	Emergency restart	Cold restart
/ERESTART COLDSYS	Cold restart	Cold restart

- The /ERE COLDBASE command performs a cold start of the DB portion of an IMS DB/DC subsystem.
If you use this command, you are responsible for the recovery of the databases. IMS does not redo committed DEDB updates, and does not back out in-flight updates for DL/I databases. IMS identifies and stops databases that have in-doubt data or that need backout or recovery. You can backout in-flight DL/I data by running the Database Batch Backout utility, for which you should close (and, optionally, archive) the OLDS.
- The /ERE COLDCOMM command performs a cold start of the TM portion of an IMS DB/DC subsystem.
This command initializes the message queues, recovers DEDBs, reloads MSDBs, and backs out in-flight changes to DL/I databases. At the same time, IMS maintains all existing in-doubt data.
- The /ERE COLDSYS command allows you to cold start both DB and TM.
This command is essentially a combination of both /ERE COLDBASE and /ERE COLDCOMM, but it does not read the OLDS. Therefore, closing the OLDS is essential. The processing described for both the /ERE COLDBASE and /ERE COLDCOMM commands also pertains to /ERE COLDSYS.

Of the three forms of emergency restart command, you will likely use /ERE COLDCOMM most often.

If an emergency restart (/ERE) fails and a subsequent emergency restart (/ERE COLDBASE or /ERE COLDCOMM) also fails, you must issue the /ERE COLDSYS command to restart IMS. You must also close (and optionally, archive) the last OLDS that was used online before performing this /ERE COLDSYS restart because the log must be available in case database recovery is necessary.


Restriction: Do not add a Fast Path DBD to the active ACBLIB between an abnormal termination and an emergency restart. If you add a Fast Path DBD to the active ACBLIB after abnormal termination of IMS, it is not accessible after the emergency restart. If you make changes to a Fast Path DBD in the active ACBLIB after the last warm start or online change (/MODIFY), an emergency restart may fail with abend U0168.

If an IMS database encounters an I/O error and the Extended Specified Task Abnormal Exit (ESTAE) routine is not able to execute, then IMS does not create an EEQE and does not lock the error block. Consequently, an IMS emergency restart does not detect the bad block. In this situation, you might need to use the OVERRIDE keyword on the /ERE command or set the abnormal termination flag in the RECON data set for the failing subsystem.

Related concepts:

“Emergency restarts and dynamic resource definition” on page 78

Related reference:

 ESAF indoubt notification exit routine (DFSFDN0) (Exit Routines)

Defer of BMP backout programs to speed emergency restart

After a system-wide failure, IMS backs out the interrupted online programs as it restarts. This backout must be complete before regular work can continue. You can request postponement of BMPs by specifying the NOBMP keyword on the /ERESTART command.

This backout hastens restart if either of the following conditions is true:

- A BMP had run a long time since its last commit point and needs large quantities of data backed out
- The program has a low priority and could have its restart postponed with no trouble

If you postpone BMP backout, you must run the Database Batch Backout utility some time before you rerun the BMP, or before you put the affected database back online (in case other programs need it).

If you register the affected database with DBRC, DBRC prevents programs from using that database before a backout is performed.

IMS restart and global resource status in an IMSplex that uses RM

If you are operating in an IMSplex environment that uses the Resource Manager (RM) component of the Common Service Layer, you need to understand how restarting an IMS system affects both the global and local status information of the IMS resources that the restarting IMS system shares with the rest of the IMSplex.

The first IMS system to start in an IMSplex sets the PLEXPARM values that the IMS system knows about in RM. Subsequent IMS systems that start in the IMSplex obtain the PLEXPARM values by reading the PLEXPARM entry in RM. Subsequent IMS systems either use the PLEXPARM entry in RM if it has a value or update the PLEXPARM value in RM if they introduce a PLEXPARM variable.

When an IMS system restarts, if global status is maintained for a resource type, the type of restart determines how the global status is applied at the restarting IMS. For example:

- For cold starts, the global status of databases, areas, and transactions takes precedence.
- For warm starts, the restarting IMS system first applies the local status saved in the log records and then applies the global status stored in RM only if it changed while the IMS system was down.
- For emergency restarts, the restarting IMS system first applies the local status saved in the log records and then applies the global status stored in RM only if it changed while the IMS system was down.
- For COLDSYS emergency restarts, the global status of databases, partitions, areas, and transactions takes precedence.
- For COLDBASE emergency restarts, the global status of databases, partitions, and areas take precedence.
- For COLDCOMM emergency restarts, the global status of transactions takes precedence.

Any error reading the global status of resources in RM during IMS initialization results in a message DFS3500W that states that the maintenance in RM of the global resource status is disabled for the specified resource type.

Before attempting to restart IMS again, take one of the following actions:

1. Cancel IMS or shutdown IMS after restart completes. After the error condition reported in the DFS3308W message is fixed, restart IMS. IMS will be synchronized with RM status, and no additional action is required.
2. Allow the IMS restart to complete, then take the following actions:
 - Issue the appropriate UPDATE IMS SET(PLEXPARM) command to enable global status at this IMS.
 - Issue the appropriate QRY TRAN|DB|AREA SHOW(GLOBAL) commands to see the resources with global status.
 - Issue appropriate command with the SCOPE(ACTIVE) parameter to ensure that the IMS is synchronized with RM status.

Related concepts:

“Starting IMS systems and global command status” on page 91

How to perform a subsystem restart with the z/OS Automatic Restart Manager

You can use the z/OS Automatic Restart Manager (ARM) to restart a subsystem or a job after a failure of the subsystem or environment in which it is running. IMS supports the z/OS Automatic Restart Manager in the DB/DC, DBCTL, DCCTL, and XRF environments.

The following subsystems also use the z/OS Automatic Restart Manager:

- The IRLM subsystem
- The CQS subsystem
- The CSL Open Database Manager (ODBM)
- The CSL Operations Manager (OM)
- The CSL Resource Manager (RM)
- The CSL Structured Call Interface (SCI)

IMS batch environments and IMS utilities do not support the z/OS Automatic Restart Manager.

After a failure, the z/OS Automatic Restart Manager will restart only the IMS control region and the CQS subsystem (each independently of the other). Other regions (such as DLISAS and DBRC) are in turn restarted by the IMS control region, and IMS dependent regions are typically restarted using automation.

In an XRF environment, the active IMS subsystem is automatically de-registered from the z/OS Automatic Restart Manager while the alternate IMS subsystem is in its tracking phase. Thus, the active subsystem is not be restarted after a failure; instead, the XRF alternate takes over as expected.

The z/OS Automatic Restart Manager ignores AUTO=N in the IMS startup parameters.

If an IMS or CQS subsystem abends before it completes restart, it is de-registered from the z/OS Automatic Restart Manager so that the abend can be fixed before attempting restart again.

If an IMS or CQS subsystem that is registered with the z/OS Automatic Restart Manager is canceled by the z/OS operator (using the CANCEL command), that IMS subsystem is not restarted by the z/OS Automatic Restart Manager unless the operator includes the ARMRESTART keyword on the command.

Related concepts:

 The z/OS Automatic Restart Manager (ARM) (System Administration)

BMP restart after system failure

IMS does not automatically restart BMPs after a system failure. You must start them yourself after IMS has completed its restart processing. The way in which you restart a BMP depends on whether its last checkpoint records exist in the OLDS.

When the last checkpoint records are in the OLDS, you can specify one of four values during restart:

- The 8-byte checkpoint ID
- The 14-byte time stamp ID IIIIDDDHHMMSSTHMIJU+0000 from message DFS0540I, where:
 - IIII is the region ID
 - DDD is the day of the year
 - HHMMSSTHMIJU+0000 is the time in hours, minutes, seconds, and tenths, hundredths, thousandths, ten-thousandths, hundred-thousandths, and millionths of a second, plus the local time zone offset
- The value LAST

Specify this value as an EXEC parameter. When you specify LAST, you must not change the BMP's job name, PSB name, or program name; IMS uses these names to locate checkpoint information for the BMP.

Restarting a BMP using XREST

If CKPTID=LAST is coded, IMS can automatically locate the last checkpoint taken by the BMP under most conditions if all of the following requirements are met:

- The BMP is restarted with the same job name, the same PSB, and the same program name used when it ended abnormally.

- The BMP restarts on the same IMS system that it ended abnormally on, and this system was not cold-started after the BMPabend occurred.
- If the restart checkpoint is no longer available on an OLDS data set, the SLDSREAD Logger function must be available. In this case, the checkpoint is on an archived SLDS data set.

If these requirements are met, an IMSLOGR DD statement that points to the log data set is not required. Coding an IMSLOGR DD statement and using IMS automatic checkpoint location are mutually exclusive processes. If an IMSLOGR DD statement is specified, IMS does not use the automated process to locate the checkpoint, but only searches the IMSLOGR concatenation. If the required checkpoint is not found in the concatenation, the BMP ends abnormally withabend code U0102.

If these requirements are not met, an IMSLOGR DD statement that points to the correct log data set is required. Automatic checkpoint location is restricted to CKPTID=LAST. If you use a checkpoint value, you must code the appropriate IMSLOGR DD statement.

Important: If you do not specify CKPTID=LAST, the data updated between the specified checkpoint and the actual last checkpoint taken has been committed by IMS.

If the program issued XRST and CHKP calls, you can restart it from the last checkpoint by specifying the checkpoint ID in the JCL and supplying, as input, the SLDS created in the program's previous execution.

Restarting batch jobs

Before you restart a DL/I batch job, ensure all database changes have been backed out. IMS might have backed out the database changes dynamically if a pseudo-abend occurred and you specified **BKO=Y** in the JCL and the data set resides on DASD. Otherwise, run the Database Batch Backout utility.

To restart a batch job, specify the ID of its last checkpoint in the JCL and supply, as input, the SLDS created in the program's previous execution. The input is specified on the IMSLOGR DD statement.

The batch job that is performing the restart must have the same job name as the original. If you do not use the same job name for the restart, IMS cannot find the checkpoint and the job fails with a U0102 error.

Reconnecting CCTLs or ODBA application programs

IMS DB system shutdown or failure generally does not cause the termination of a CCTL or ODBA application program, but it does disconnect the CCTL or ODBA application program. The CCTL or ODBA application program can automatically reconnect after an IMS system termination, or a CCTL operator can manually reconnect it.

Related reference:

 CCTL exit routines (Exit Routines)

Chapter 3. Monitoring IMS

You can monitor IMS by using various IMS commands, utilities, and tools.

Related concepts:

 Tools for detailed monitoring (System Administration)

Monitoring the system

In order to gather problem determination and performance information, you need to monitor the status of the system on a regular schedule. For example, to determine if you should start an extra message region, you might monitor the status of the queues during peak load.


You can determine the current status of the system by issuing `/DISPLAY STATUS` or `QUERY` commands, specifying the appropriate keywords. You can monitor the status of the IRLM by using the z/OS command `MODIFYIrImlproc,STATUS`.

You can use traces enabled by commands such as `/TRACE`, `UPDATE TRAN`, and `UPDATE MSLINK` to help diagnose system operation problems. These commands turn on or off various IMS traces, which record use of IMS control blocks, message queue and I/O line buffers, save area sets, and MSC links.

IMS records trace information on the IMS log unless you request that the traces be recorded on an external trace data set. You can also use the `/TRACE` command to trace locking activities and to start and stop the IMS monitor.

You can use the Operations Manager (OM) audit trail log to see the commands, the command responses, and the selected system messages that are routed through OM. When OM audit trail logging is enabled, the OM audit trail log is saved to log stream in the z/OS System Logger, as specified on the `AUDITLOG=` keyword of the `IMSPLEX` parameter in the OM initialization parameters `PROCLIB` member (`CSLOIxxx`). You can use TSO `SPOC` to view the OM audit trail log.

Related concepts:

 IMS Trace Facility (System Administration)

 IMS Monitor (System Administration)

 Collecting and interpreting IMS monitoring data (System Administration)

“User access problems” on page 151

Monitoring IMS Connect connections

IMS manages TCP/IP connections through IMS Connect, which serves as the TCP/IP gateway to IMS. The IMS Connect function can also be viewed as a TCP/IP socket server.

TCP/IP connections to IMS Connect can be classified into the following types:

- IMS DB client connections, where the client is accessing IMS DB.
- IMS TM client connections, where the client is accessing IMS TM.
- IMS Control Center connections, where the IMS Control Center is used to submit IMS type-2 commands to an IMSplex from a TCP/IP connection.

- IMS-to-IMS TCP/IP connections, where one IMS Connect instance connects to another IMS Connect instance for communications between two geographically remote IMS systems.

In addition to the TCP/IP connections that can be made to IMS Connect, you must also consider the connections between IMS Connect and the IMS server. IMS Connect runs in a separate address space from IMS and connects with certain components of IMS by using different communications methods.

For IMS DB client support, IMS Connect uses the Structured Call Interface (SCI) component of the IMS Common Service Layer (CSL) to connect to the IMS Open Database Manager (ODBM) component of CSL. ODBM communicates directly with IMS DB through the IMS Open Database Access (ODBA) interface.

IMS Connect also uses SCI for direct communications with IMS when IMS-to-IMS TCP/IP connections are used to support Multiple Systems Coupling (MSC) TCP/IP links.

For IMS TM client support, IMS Connect uses the z/OS cross-system coupling facility to connect to the Open Transaction Manager (OTMA) component of IMS.

To monitor and verify the status of each type of the connection managed by IMS Connect, you might need to use different commands or facilities.

The following subsections describe only some of the commands that you can use to check the status of IMS Connect connections.

Related concepts:

“MSC TCP/IP link operations” on page 273

Checking port TCP/IP connection status

You can check the status of all of the TCP/IP connections on a port.

To see the status of all of the TCP/IP connections on a port:

Issue any one of the following IMS Connect commands:

- In the IMS type-2 command format, QUERY IMSCON TYPE(PORT)
- In the WTOR command format, VIEWPORT
- In the z/OS MODIFY command format, QUERY PORT

Related reference:

 [QUERY IMSCON commands \(Commands\)](#)

 [IMS Connect WTOR commands \(Commands\)](#)

 [IMS Connect z/OS commands \(Commands\)](#)

Checking client TCP/IP connection status

You can check the status of the TCP/IP connections for one or more IMS Connect clients.

To check the status of TCP/IP connections for specific clients:

Issue the command QUERY IMSCON TYPE(CLIENT) in the IMS type-2 command format. No equivalent WTOR or z/OS MODIFY command exists for QUERY

IMSCON TYPE(CLIENT); however, the WTOR VIEWPORT command and the z/OS MODIFY QUERY PORT command provide some client-related connection information.

Related reference:

➞ [QUERY IMSCON commands \(Commands\)](#)

➞ [IMS Connect WTOR commands \(Commands\)](#)

➞ [IMS Connect z/OS commands \(Commands\)](#)

Checking remote IMS Connect TCP/IP connection status

You can check the status of TCP/IP connections with a remote IMS Connect instance.

To check the status of a TCP/IP connection to another instance of IMS Connect:

Issue any of the following IMS Connect commands:

- In the IMS type-2 command format, QUERY IMSCON TYPE(RMTIMSCON) or QUERY IMSCON TYPE(SENDCLNT)
- In the WTOR command format, VIEWRMT
- In the z/OS MODIFY command format, QUERY RMTIMSCON

Related reference:

➞ [QUERY IMSCON commands \(Commands\)](#)

➞ [IMS Connect WTOR commands \(Commands\)](#)

➞ [IMS Connect z/OS commands \(Commands\)](#)




Checking IMSplex member SCI connection status

You can check the status of connections between IMS Connect and SCI within an IMSplex.

IMS Connect uses SCI to communicate a number of different IMSplex members for a number of different purposes. IMS Connect provides different commands to check the status of an IMSplex and the status of the IMSplex members.

- For IMSplex status, issue any of the following IMS Connect commands:
 - In the IMS type-2 command format, QUERY IMSCON TYPE(IMSPLEX)
 - In the WTOR command format, VIEWIP
 - In the z/OS MODIFY command format, QUERY IMSPLEX
- For IMS connection status for MSC TCP/IP link support, issue any of the following IMS Connect commands:
 - In the IMS type-2 command format, QUERY IMSCON TYPE(MSC)
 - In the WTOR command format, VIEWMSC
 - In the z/OS MODIFY command format, QUERY MSC
- For ODBM connection status for IMS DB access, issue any of the following IMS Connect commands:
 - In the IMS type-2 command format, QUERY IMSCON TYPE(ODBM)
 - In the WTOR command format, VIEWIA
 - In the z/OS MODIFY command format, QUERY ALIAS

Related reference:

-  [QUERY IMSCON commands \(Commands\)](#)
-  [IMS Connect WTOR commands \(Commands\)](#)
-  [IMS Connect z/OS commands \(Commands\)](#)

Checking XCF data store connection status for IMS TM clients

You can issue variations of the IMS and IMS Connect commands to determine the status of XCF data store connections for IMS TM clients.

In the following examples, the XCF connection between IMS Connect and OTMA is defined by the following DATASTORE configuration statement:

```
DATASTORE=(ID=DSNAME, MEMBER=ICONNAME, TMBER=IMSNAME, GROUP=GRPNAME...)
```

- From IMS, use the following commands to check the status of IMS Connect:

- /DIS OTMA

When IMS Connect is ready for use, the output of the /DIS OTMA command appears as shown in the following example:

```
GROUP/MEMBER  XCF-STATUS  USER-STATUS  SECURITY TIB  INPT SMEM  DRUEXIT  T/O  ACEEAGE
GRPNAME
-IMSNAME      ACTIVE      SERVER      FULL*
-ICONNAME     ACTIVE      ACCEPT TRAFFIC
```

* CHECK, FULL, NONE or PROFILE depending on the OTMA Security setting (for example, enter /SEC OTMA NONE on the MVS system console to turn off RACF security for IMS OTMA clients). FULL is the default setting for OTMA security at IMS startup.

- /DISPLAY TMBER *tmembername* TPIPE

For example, the command input /DISPLAY TMBER CLIENT1 TPIPE ALL might produce the following output:

```
MEMBER/TPIPE  ENQCT  DEQCT  QCT STATUS
CLIENT1
-TPIPE1       0       0       0 TRA,STO
-TPIPE2       2       2       0 TRA,STO
-TPIPE3       1       0       1 TRA,STO
*94165/170756*
```

- From IMS Connect, you can use the following WTOR-format commands to display status:

VIEWHWS

Displays information about data stores, ports, and IMSplex.

Tip: If you use the SUMMARY option for the VIEWHWS command, you can display IMS Connect information without the detailed output for each socket.

VIEWDS

Displays information about data stores.

Alternatively, you can issue the following IMS type-2-format commands through the OM API to display the status:

QUERY IMSCON TYPE(CONFIG)

Displays the status and activities of IMS Connect. Unlike the VIEWHWS command, individual resources such as port and data store are not displayed.

QUERY IMSCON TYPE(PORT)

Displays the status and activities of one or more ports defined to IMS Connect.

If you specify the SHOW(CLIENT) keyword, an additional line is displayed for each client associated with the specified port.

QUERY IMSCON TYPE(DATASTORE)

Displays the status and activities of one or more data stores defined to IMS Connect.

- Enter the following IMS commands if you fail to receive a response to a request message sent from an IMS TM client (or to check the readiness of the host data store):

/DIS A REG

You can use this command to verify that the dependent region where your host application runs is properly configured and ready to accept messages:

REGID	JOBNAME	TYPE	TRAN/STEP	PROGRAM	STATUS	CLASS
1	Job1	TP			WAITING	1, 2, 3, 4

/DIS TRAN TranName

You can use this command to verify the class and status of a transaction and whether or not a transaction is currently queued for processing:







TRAN	CLS	ENQCT	QCT	LCT	PLCT	CP	NP	LP	SEGSZ	SEGNO	PARLM	RC
TRANNAME	2	1	1	65535	65535	8	8	8	0	0	NONE	0

QCT is the number of transactions that are currently queued. ENQCT includes transactions that have been dequeued (processed), as well as those that are currently on the queue.

- If you receive an IMS Connect error message on either the z/OS system console or in the response message at the client when you issue IMS or IMS Connect commands, the format of IMS Connect message numbers is HWSxnnnn, where the *x* represents an alphabetic character and *nnnn* represents a 4-digit number.
- IMS Connect requires that all active clients have unique client names. IMS TM Resource Adapter creates a unique CLIENTID, which identifies each request that an application makes to execute an IMS transaction. If you are using TCP/IP clients that access IMS TM other than IMS TM Resource Adapter, you must ensure that those clients each use a unique client name. This client name value is displayed for a client in the "CLIENT=" or "ORIGIN=" fields. When IMS Connect receives an initial request from IMS TM Resource Adapter with a blank Client ID, IMS Connect generates a unique Client ID to be used for that TCP/IP persistent socket. IMS Connect returns this Client ID to IMS TM Resource Adapter on any reply message.

Important: There are OTMA restrictions on IMS commands.

Related reference:

-  (Communications and Connections)
-  /DISPLAY OTMA command (Commands)
-  /DISPLAY TMEMBER command (Commands)
-  QUERY IMSCON commands (Commands)
-  IMS Connect WTOR commands (Commands)
-  IMS Connect z/OS commands (Commands)

IMS system log utilities

The system log data sets are a basic source for statistics about the processing performed by the online system. Individual log record types contain data that can be analyzed in many ways. For example, you can select and format all activity pertaining to a specified user ID, or about IMS pools.

IMS provides several utilities to assist with extracting log records from the system log. These utilities also assist with reducing and merging data that is spread across several log data sets.

In addition to the following utilities, IMS Performance Analyzer for z/OS provides a comprehensive suite of reports to help you manage the performance and resource utilization of your IMS systems.

Related concepts:

-  IMS Performance Analyzer for z/OS overview

File Select and Formatting Print utility (DFSERA10)

You use this utility to examine message segments or database change activity in detail. This utility prints the contents of log records contained in the OLDS, SLDS, or the CQS log stream. Each log record is presented as one or more segments of 32 bytes. The printed output gives each segment in both character and hexadecimal formats.

You can specify selection criteria for a subset of the records rather than printing all records. You can also specify a starting record number and the number of records to process. You can use an exit routine to customize the selection and formatting of the log records.

Although you can use the File Select and Formatting Print Program to copy entire input logs, you can more conveniently use the Log Archive utility (DFSUARC0). You use one or more SLDSs as input and specify a user data set as output. Also, you need to specify DBRC=N0 in the EXEC statement to prevent DBRC from making entries in the RECON data set about your backup log. Making backup copies of the system log data sets can be useful to obtain an alternative input source for statistics and other monitoring activities occurring in parallel with production use of the system log.

Log Transaction Analysis utility (DFSILTA0)

In an IMS DB/DC or DCCTL environment, you can use this utility to collect information about individual transactions, based on records on the system log.

Many events are tabulated in the Log Analysis report produced by this utility, including total response time, time on the input queue, processing time, and time on the output queue.

You can select a start time for the report tabulation; analysis begins at the first checkpoint after the start time. To control how much transaction activity is tabulated, you can specify an interval (in minutes) of elapsed time from the start time before the utility ends the tabulation, or you can relate the activity reported to a number of IMS checkpoints.

You can retitle a Log Analysis report or change the sequence in which the detailed transaction lines are printed. You can sort by transaction code or by any of the fields in the report. You can also suppress printing so that the output is stored on a DASD data set.

Using this utility, you can create an output data set, in system log format, that is a copy of all or part of the input system logs. By having a copy of the system log, you can monitor system activity without impacting the use of the OLDS for recovery.

Related reference:

 Log Transaction Analysis utility (DFSILTA0) (System Utilities)


Statistical Analysis utility (DFSISTS0)

In an IMS DB/DC or DCCTL environment, you can use this utility to produce several summary reports. You can use these reports to obtain actual transaction loads and response times for the system. The statistics produced are dependent on the input system log data sets.

The following set of reports is produced:

- Telecommunication line and terminal (distributed traffic over 24-hour day)
- Transaction (distributed activity over 24-hour day)
- Transaction response
- Messages queued but not sent (listing by destination and by transaction code)
- Program-to-program messages (listing by destination and by transaction code)
- Application accounting
- IMS accounting

Related reference:

 Statistical Analysis utility (DFSISTS0) (System Utilities)

Gathering performance-related data



IMS provides the DB Monitor and the IMS Monitor that gather and format IMS performance-related data and record this data on a statistics log.

The DB Monitor is available to IMS batch systems; it can monitor the activity between application programs and databases. The IMS Monitor is available to IMS online systems. In addition to performing all of the functions of the DB Monitor, the IMS Monitor can track and record information about activities occurring in the IMS control region and data communication activities.

IMS uses the statistics produced by each Monitor to generate reports. The report programs run offline and print reports that summarize and categorize IMS activities.

The DB Monitor records performance data during execution of an IMS DB batch system. The DB Monitor can be active during the entire execution of an IMS batch job, or you can stop and restart it from the system console.

Related concepts:

-  DB Monitor reports (System Administration)
-  IMS Monitor reports (System Administration)

Activating and controlling the DB Monitor

To activate the DB Monitor, specify MON=Y in the PROC statement of the batch job. When you submit the job, IMS uses parameter substitution to update the PARM field of the EXEC statement with a Y in the appropriate position.

To stop the DB Monitor, the system console operator can use the MODIFY*jobname*,STOP command. Message DFS2215A displays on the system console when the Monitor is inactive.

To reactivate the DB Monitor, the console operator can use the MODIFY*jobname*,START command.

Message DFS2216A displays on the console when the Monitor is active again.

Logging the data from the DB Monitor

IMS records the data produced by the DB Monitor on either the batch log or a separate DB Monitor log defined by the //IMSMON DD statement. You use the //IMSMON DD statement in the batch procedure to control where data is logged.

- To store the Monitor records on the batch log, either include a //IMSMON DD DUMMY statement or omit the //IMSMON DD statement entirely.
- To store the Monitor records on a separate DB Monitor log, include a valid //IMSMON DD statement.

If, for any reason, IMS cannot open the DB Monitor log data set specified on the //IMSMON DD statement, IMS displays message DFS2217I on the system console. Batch execution continues, but the Monitor is inactive.

If the DB Monitor log device encounters I/O errors, IMS displays message DFS2219I on the system console. Batch execution continues, but the Monitor is inactive.

If you want to stop the Monitor and force an end-of-volume for the DB Monitor log, use the MODIFY*jobname*,STOPEOV command. When you use the STOPEOV keyword, the batch region does not continue executing until the z/OS mount request for a new data set is satisfied.

Note: If you enter an incorrect job name on the MODIFY command, z/OS issues an error message. If you make some other error while entering the MODIFY command, IMS issues message DFS2218I, followed by either message DFS2215A or message DFS2216A.

Data recording and the IMS Monitor

The IMS Monitor records performance-related data during execution of the IMS online subsystem. When a significant event occurs within IMS, IMS passes relevant data to the IMS Monitor when activated. The IMS Monitor formats a record describing the event, including the time stamps, and logs the event.

Activating and controlling the IMS Monitor

If you create a DFSDCMON member in the IMS.SDFSRESL data set, you do not need a DD statement in the IMS, DBC, or DCC procedures for the IMS Monitor because IMS dynamically allocates and deallocates the IMS Monitor data set. You start and stop the IMS Monitor with the `/TRACE` command.

Otherwise, to activate the IMS Monitor, you must include a DD statement (using IMSMON as the data set name) in the IMS, DBC, or DCC procedures to specify the IMS Monitor log data set. When you include this DD statement, the IMS Monitor becomes available. If the IMS Monitor is available but inactive, processor usage is unaffected.

Monitoring using the `/TRACE` command

In addition to starting and stopping the IMS Monitor, you can specify the types of events to be monitored using the `/TRACE` command. You can use the `/TRACE` command to monitor various activities.

Some of the activities that you can monitor with the `/TRACE` command include:

- Telecommunication line and logical link activity
- Scheduling and termination events
- Activity between application programs and message queues
- Activity between application programs and databases (full function and Fast Path)

You can also use the `/TRACE` command to limit the monitoring to:

- Particular databases, partitions, or areas
- Particular dependent regions
- A particular interval of time

Log data from the IMS Monitor

Due to how your IMS Monitor log data set is stored, you should process the log after you stop the IMS Monitor and before restarting it again.

If the IMS Monitor log data set is on tape, IMS issues a tape mount request each time you start the IMS Monitor, and IMS rewinds the tape each time you stop the IMS Monitor. If the IMS Monitor log data set is on DASD, IMS uses the same data set each time you start the IMS Monitor.

Recommendation: Start the IMS Monitor and allow it to run for a period of time, and then stop it to write a “snapshot” of current activity on the IMS Monitor log. You must stop the IMS Monitor before you take a shutdown checkpoint in order for the report program to produce usable output.

I/O errors on the IMS Monitor log data set

If a permanent I/O error occurs on the IMS Monitor log data set, IMS stops the IMS Monitor and issues message DFS2202. In this situation, you cannot restart the IMS Monitor until you restart IMS because IMS does not close the IMS Monitor log data set until you shut down IMS.

If the problem that caused the error has not been corrected when you restart IMS, you should specify a different volume or unit for the new execution.

Monitoring IRLM

You can use the `MODIFYirlmproc,STATUS` command to obtain information about IRLM activity.

This command provides:

- The IMS IDs of IMS subsystems using this IRLM
- The number of locks that are held and waiting for each subsystem on this IRLM
- Identification of this IRLM: its subsystem name and IRLM number

Tracing IRLM activity

The IMS Monitor does not collect IRLM trace activity. IRLM uses the z/OS component trace (CTRACE) facility. You can use the `TRACE CT` command to run various types of sublevel traces.

DBM Trace interactions with the identified DBMS.

EXP Trace any exception condition.

INT Trace member and group events other than normal locking activity.

SLM Trace interactions with the z/OS locking component.

XCF Trace all interactions with the z/OS cross-system coupling facility services.

XIT Trace only asynchronous interactions with the z/OS locking component.

Chapter 4. Shutting down IMS

This approach is a common sequence for shutting down the entire online system.

1. Stop communications based on your IMS environment.

Option	Description
Stop data communications	For an IMS DB/DC or DCCTL environment
Disconnect from the CCTL	For an IMS DBCTL environment

2. Stop dependent regions.
3. Stop MSC links.
4. Stop OTMA.
5. Stop the control region.
6. For an IMS DB/DC or DBCTL environment, stop the IRLM.
7. For a shared-queues environment, if CQS has not shut down, shut it down.
8. For an IMSplex system using CSL, shut down the CSL manager address spaces (OM, RM, and SCI).

The command used to shut down the control region also forces termination of data communications and the dependent regions if they have not already been terminated in an orderly way, and can also tell the CQS subsystem to shut down.

Stopping Transaction Manager

You can use the `/STOP DC` command to prevent new users from logging on to VTAM terminals. This command does not terminate existing sessions.

You can use the `/STOP LINE` and `IDLE LINE` commands to disable devices such as printers, spools, punches, disks, readers, and tapes. You can use the `/PURGE TRAN` or `UPDATE TRAN START(SCHD) STOP(Q)` command to prevent existing VTAM terminal users from entering new transactions.

To terminate a VTAM session, you can use one of the following commands, or you can wait for the end user to terminate the session.

- `/STOP NODE`
- `/CLSDST NODE`
- `VTAM VARY`
- `/IDLE NODE`

You can use this command if you already entered a `/STOP DC` or `/CHECKPOINT` command or you can wait for the end user to terminate the session.

Recommendation: Use the `/CLSDST NODE` command to terminate a VTAM session instead of the `/STOP NODE` command. The `/STOP` command requires that you subsequently issue a `/START` command before you or a user can initiate a session. You can initiate a session using the `/OPNDST` command.

To shut down VTAM, use the `HALT NET` command (abbreviated as `Z NET`). If you issue a `Z NET` command with no keywords, you terminate VTAM, and IMS

permits no new logons, but existing sessions continue. The existing VTAM sessions must be ended by the terminal user, the VTAM operator, or the MTO. If you issue the Z NET,QUICK command, VTAM does not allow new sessions to be established or any additional input or output operations to take place, and IMS terminates existing sessions.

Stopping APPC

You use the /STOP APPC command to prevent new users or programs from allocating APPC conversations with the IMS APPC BASE LU.

The /STOP APPC CANCEL command prevents new users or programs from allocating APPC conversations with any IMS APPC LU. In both cases, the conversation is rejected with a TP_Not_Available_No_Retry return code by APPC/MVS.

Stopping OTMA

You use the /STOP TMEMBER command to suspend input from one or more specified OTMA client transaction members (tmembers).

The connection remains active, transactions already being processed are unaffected, conversational transactions already in progress can complete normally, and all of the OTMA protocol commands can still be processed by OTMA.

If you want to shut down OTMA completely, use the /STOP OTMA command.

Stopping dependent regions

You can use the /STOP REGION command to terminate dependent regions. You also terminate dependent regions when you shut down the control region using a /CHECKPOINT command.

Shutting down the IMS control region

Use the /CHECKPOINT FREEZE|DUMPQ|PURGE command to shut down the IMS control region. You should provide guidelines for the MTO on selecting the right kind of shutdown, either specifying the FREEZE or PURGE keyword.

Using the FREEZE keyword causes the quickest shutdown, and the PURGE keyword causes the slowest shutdown.

- /CHECKPOINT FREEZE|DUMPQ|PURGE immediately terminates sessions for all logical units as follows:

FREEZE

Immediately after current input/output message

DUMPQ

After blocks have been written to the log

PURGE

After all queues are empty

- /CHECKPOINT [FREEZE|DUMPQ|PURGE] QUIESCE allows all network nodes to complete normal processing before IMS shuts down.

If dynamic resource definition (DRD) is enabled for your IMS system, you must ensure that changes to your runtime resource and descriptor definitions are saved to the resource definition data set.

The following table illustrates what IMS does in response to each type of checkpoint command.

Table 20. System actions during checkpoint.

System action	FREEZE	DUMPQ ¹	PURGE
Stop terminal input	X	X	X
Process transactions		X ²	X
Free message regions	X	X	X
Purge real storage message queue to disk	X	X	
Log checkpoint data	X	X	X
Notify the CCTL	X		X
Wait for the CCTLs to disconnect from their resource managers	X		X
Write MSDBs to checkpoint data set	X	X	X
Write CHKPT ID to master terminal	X	X	X
Purge database buffer pool	X	X	X
Terminate message regions	X	X	X
Send queued output	X	X	X
Wait for BMP's sync point	X	X	
Wait for BMP's completion			X
Send shutdown message to terminals		X	X
Dump message queues to log		X	
Close databases	X	X	X
Close queues	X	X	X
Close log	X	X	X
Terminate IMS control region	X	X	X

Notes:

1. The /CHECKPOINT DUMPQ command is not valid for an IMS DBCTL environment.
2. Applies to Fast Path messages only.

Recommendations:

- Have the MTO broadcast a message to all users about 10 minutes before entering the IMS shutdown command. This lets end users, except those involved in conversations, terminate their work in an orderly way.
- If you use the IMS monitor, stop it with the /TRACE SET OFF MON command before using the /CHECKPOINT command.
- Allow active VTAM terminals to complete processing before the IMS shutdown begins.

After the shutdown process has begun, you can use the /DISPLAY SHUTDOWN STATUS command to list the following information:

- Communication lines and terminals that still contain active messages
- BMPs and CCTL threads that are still active

To speed up the shutdown process, you can use the `/IDLE` command to stop I/O operations on specified lines.

In the IMS DBCTL environment, IMS notifies the CCTL that IMS is shutting down. In response, the CCTL must disconnect from the IMS DBCTL subsystem. The CCTL can control how long its threads process before it disconnects, for example, it could let all threads reach a sync point first.

If IMS fails to shut down, and if logging resources are available, you must force IMS to terminate.

Exceptions:

- If no logging resources are available, the OLDS must be archived before you can force shutdown.
- IMS does not shut down if the CQS subsystem is unavailable while you are running in a shared-queues environment.

Related concepts:

“Forced termination of IMS” on page 126

Shutting down IMS using the `/CHECKPOINT` commands

The way that IMS shuts down depends on which form of the `/CHECKPOINT` command you use. You can choose to specify the `FREEZE`, `DUMPQ`, or `PURGE` keyword, and each keyword shuts down IMS differently.

The forms of the `/CHECKPOINT` command are:

- `/CHECKPOINT FREEZE`
- `/CHECKPOINT DUMPQ`
- `/CHECKPOINT PURGE`

Shutting down IMS using `/CHECKPOINT FREEZE`

The fastest way to terminate IMS in an orderly way is to use the `/CHECKPOINT FREEZE` command. Using the `/CHECKPOINT FREEZE` command, you allow IMS to complete input and output messages that are in transit. Then IMS:

- Stops each communication line
- Terminates VTAM sessions
- Terminates message regions as soon as the messages currently being processed have completed
- Returns status code XD to batch message regions when they next issue a DL/I checkpoint call

In a DBCTL environment, IMS notifies the CCTL of the `/CHECKPOINT FREEZE` command so the CCTL can disconnect from its connected resource managers. The CCTL controls how long its CCTL threads can continue processing before the disconnect occurs. For example, the CCTL might let all its threads reach a sync point.

During shutdown, IMS also closes message queue data sets and databases, and writes checkpoint data to the OLDS, just as for a simple checkpoint. IMS issues message DFS994, which contains the last checkpoint ID, and closes the OLDS. IMS updates the checkpoint ID table to reflect the shutdown, and writes it to the restart data set (RDS). Then IMS closes the RDS and terminates the IMS control region.

During normal shutdown, IMS tells CQS to shut down. If you do not want CQS to shut down, specify the NOCQSSHUT keyword on the /CHECKPOINT command. If you are shutting down a single IMS subsystem in a sysplex in order to activate a system definition, change a parameter specification, or for reasons that are not related to the CQS subsystem, you do not need to shut down the CQS address space. By leaving the CQS address space active after IMS has shut down, CQS can continue to participate in structure checkpoint, overflow, and rebuild events.

Recommendations:

- Ensure that the CQS address space is active during structure checkpoint processing. During a structure checkpoint, z/OS deletes log records that are no longer needed for structure recovery, thus reclaiming log space. If CQS is not active during the structure checkpoint, it is possible that the CQS restart log records will be deleted. In this case, the next restart for the CQS address space must be a cold start. You must respond COLD to WTOR message CQS0032A. If you choose to leave CQS up, you can reduce the number of times that CQS needs to cold start due to log record deletion.

If CQS is active when IMS is started, IMS reconnects to the existing CQS and does not need to wait for CQS initialization and restart to complete.

- Although using the /CHECKPOINT FREEZE command is the fastest way to accomplish an orderly shutdown, you should not use it regularly for normal system termination. When you use the /CHECKPOINT FREEZE command, end users might not receive responses for extended periods of time (until you restart IMS). IMS also rejects Fast Path input messages during shutdown. If during restart, IMS must load the message queues from the log to the message queue data sets (a BUILDQ), you might need to access previously archived OLDSs (to isolate the last DUMPQ or SNAPQ checkpoint).

Shutting down IMS using /CHECKPOINT DUMPQ

The /CHECKPOINT DUMPQ command requests an immediate shutdown. Using the /CHECKPOINT DUMPQ command, you allow IMS DB/DC or IMS DCCTL to complete input and output messages that are in transit. Then IMS:

- Stops each communication line
- Terminates VTAM sessions
- Terminates message regions as soon as the messages currently being processed have completed
- Allows Fast Path message-driven regions to process all received messages before terminating
- Returns status code XD to batch message regions when they next issue a DL/I checkpoint call

In a nonshared-queues environment, IMS writes the contents of the message queues to the system log along with checkpoint data. In a shared-queues environment, IMS does not dump the shared queues during shutdown.

The cost of executing a /CHECKPOINT DUMPQ command over a /CHECKPOINT FREEZE command is the time it takes to dump the queues; this time varies depending on the system design, transaction volumes, and activity at a specific time of day.

Shutting down IMS using /CHECKPOINT PURGE

Using the /CHECKPOINT PURGE command is the most time-consuming method of terminating IMS. As the result of the /CHECKPOINT PURGE command, IMS stops each input communication line as soon as it receives any messages that are in transit. IMS then:

- Processes all messages in the input queue, if possible (if the transaction and program are not stopped)
- Transmits all output, if possible (if the line and terminal are not stopped)
- Stops output communication lines after it has terminated all active regions and has sent all possible output messages
- Terminates VTAM sessions

IMS writes any unprocessed input messages and any untransmitted output messages to the system log along with checkpoint data.

For an IMS DBCTL environment, the /CHECKPOINT PURGE command operates exactly like the /CHECKPOINT FREEZE command, except that the IMS DBCTL system waits for BMPs to complete their processing.

Forcing an acceleration of the shutdown process

All three variations of full shutdown (/CHECKPOINT FREEZE | DUMPQ | PURGE) allow transactions that had already been started to complete, and allow BMP programs to run to their next checkpoint. These commands might also allow CCTL threads that had already been started to complete. In some cases, you might have to wait a long time before IMS finally shuts down. For example, a conversational transaction or CCTL might be waiting for an end user to enter a response; or a long BMP program might not take any checkpoints.

If necessary, you can force shutdown (although some work might be lost). You can use various forms of the /DISPLAY or QUERY command to discover what work is still running. Then, you can use other commands to force shutdown (for example, the /EXIT, /IDLE LINE, or /STOP REGION commands).

If all else fails, you can bring down IMS by issuing a z/OS MODIFY IMS,DUMP command or MODIFY IMS,STOP command.

Attention: Only use the z/OS MODIFY IMS commands as a last resort, because you must immediately enter an emergency restart (/ERESTART) command and shut down again before you can run batch jobs.

Shutting down an IMS system that uses dynamic resource definition

If dynamic resource definition (DRD) is enabled and you configured your system to automatically import resource and descriptor definitions from a resource definition data set (RDDS) or an IMSRSC repository during a cold start, ensure that any online changes to the runtime resource and descriptor definitions are reflected in a system RDDS or repository.

Any changes made to online resources and descriptors that are not saved to a system RDDS or repository are lost when you perform a cold start or when you use the /ERESTART command with the **COLDsys** parameter to perform an emergency restart.

If needed prior to shutting down IMS, initiate an export of resource and descriptor definitions to a system RDDS in one of the following ways:

- A system programmer can configure IMS to automatically export resource definitions when the /CHECKPOINT command is issued by setting up automatic export to an RDDS.
- Use the EXPORT command to export runtime resource and descriptor definitions from the online IMS system to a resource definition data set.

If needed prior to shutting down IMS, export resource and descriptor definitions to a repository by issuing the EXPORT DEFN TARGET(REPO) command.

Related concepts:

 Overview of the IMSRSC repository (System Definition)

 Exporting resource and descriptor definitions to an RDDS (System Definition)

Related tasks:

 Exporting resource and descriptor definitions to an IMSRSC repository (System Definition)

Related reference:

 /ERESTART command (Commands)

 /CHECKPOINT command (Commands)

Session termination

Terminating a session releases the terminal from its current logical connection to IMS (the VTAM application program). Terminating a session also makes the terminal available for sessions with other VTAM application programs, or you can terminate communications altogether.

Sessions can be terminated by the IMS MTO, the VTAM network operator, or the user at the terminal. There are two types of session termination:

Orderly termination

The terminal is allowed to complete normal processing before the session is terminated.

Immediate termination

Forces the terminal to terminate the session unconditionally.

Because these two session termination methods have different consequences, you must develop specific procedures for different situations.

Orderly termination

To initiate an orderly termination of the network, use the IMS /CHECKPOINT [FREEZE | PURGE | DUMPQ] QUIESCE command. When you specify the QUIESCE keyword, IMS sends a VTAM shutdown command to all terminals and waits until they all complete normal processing.

After all terminals indicate that shutdown is complete, IMS issues VTAM CLSDST macro instruction to cause VTAM to send the UNBIND command to all terminals. This command releases the terminals from their sessions with IMS, and prohibits additional data transmission.

During an orderly termination, the IMS MTO can terminate the network unconditionally, rather than wait for the orderly termination to complete, by initiating an immediate termination.

Immediate termination

To initiate an immediate termination of the network, use the IMS `/CHECKPOINT FREEZE | PURGE | DUMPQ` command. When you enter this command, without the `QUIESCE` keyword, IMS issues the VTAM `CLSDST` macro instruction to cause VTAM to send the `UNBIND` command to all terminals. This command releases the terminals from their sessions with IMS, and prohibits additional data transmission.

You can use the `/CLSDST` or `/STOP` commands to terminate stations or portions of a network selectively. These commands cause IMS to issue the VTAM `CLSDST` macro instruction for the specified stations. The `/STOP` command also prevents additional sessions from being established until you issue a `/START` command for the terminal.

The VTAM network operator can use the `VARY` command to terminate a session immediately. It is sometimes necessary for the VTAM network operator to use this command to terminate sessions that have an error preventing an I/O operation from completing.

The terminal user can use VTAM session control commands to terminate a session with IMS. For normal processing, when the user decides to terminate a session with IMS, the terminal should send the VTAM request-shutdown command. IMS completes any input or output currently in progress for that terminal and then issues the VTAM `CLSDST` macro instruction.

If the terminal detects an error condition from which it cannot recover, it can send the VTAM `terminate-session` command. VTAM releases the terminal from the session and notifies IMS.

Shutting down an IMS network

You can shut down the IMS network while shutting down IMS or without shutting down IMS. Use the `/CHECKPOINT` command to terminate the network and shut down IMS. Use the `/STOP DC` command to terminate the network only.

You can also use the VTAM `HALT NET` command to shut down the VTAM network.

If your network includes APPC, be aware of the following conditions:

- Allocated LU 6.2 conversations might prolong the shutdown process.
- Network shutdown does not prevent the application programs from establishing LU 6.2 conversations.
- LU 6.2-originated transactions, not scheduled but on the IMS queue, terminate with message DFS1970.

Terminating an ISC session from CICS

You can only use CICS control operator commands to terminate ISC sessions.

The CICS operator can release a session by using the `CEMT SET TERMINAL(termid) RELEASED | OUTSERVICE` command, where *termid* is the four-character session name (defined on the `DEFINE SESSIONS` command in the

CSD utility), or the TRMIDNT in the DFHTCT TYPE=TERMINAL program. The following list describes the effects of the keywords for the CEMT SET TERMINAL command:

RELEASED

The session terminates when active transactions complete (unless you also specify the PURGE or FORCEPURGE keywords) and the session is left in a VTAM between-brackets state.

When you specify RELEASED, CICS initiates an orderly termination between IMS and CICS. Although, from the CICS point of view, the session may appear to be in warm-start state, it is actually in cold-start state as a result of the VTAM SBI/BIS (Stop Bracket Initiation / Bracket Initiation Stopped) command flow.

OUTSERVICE

The session terminates when active transactions complete (unless you also specify the PURGE or FORCEPURGE keywords), the session is left in a VTAM between-brackets state, and no additional transactions can use the terminal. For a VTAM terminal, using this keyword also causes it to be released and the operator to be signed off, either immediately or when the current transaction has completed.

When you specify OUTSERVICE, CICS requests resynchronization when the session is re-initiated, that is, the session is left in a warm-start state. To initiate a session in cold-start mode, the CICS operator must use the CEMT SET TERMINAL (termid) COLDACQ command.

You can specify both RELEASED and OUTSERVICE together on the same command.

Restriction: The OUTSERVICE keyword is not applicable for LU 6.2 devices.

Recommendation: Although a CICS application program can issue an EXEC CICS DISCONNECT call to initiate orderly session termination, this approach is not recommended in a normal application. However, you can write an “operator control” application to issue this call.

Any messages about the session's termination are sent to transient data destination CSMT.

Related tasks:

“Connecting ISC sessions from CICS to IMS” on page 90

Stopping the IRLM

You can stop the IRLM from the system console through a version of the MODIFY or STOP commands.

```
MODIFYirlmproc,ABEND,NODUMP
STOPirlmproc
```

Recommendation: Terminate the IRLM subsystem only after all IMS subsystems connected to the IRLM have completed their processing.

Shutting down CQS

During normal shutdown in a shared-queues environment, IMS tells Common Queue Server (CQS) to shut down, but if it does not, you can use the z/OS STOP command to shut down the CQS subsystem. To shut down CQS, IMS can issue the CQSSHUT request, the CQSDISC request with the CQSSHUT=YES parameter command.

The IMS subsystem does not have access to the shared queues as long as CQS is down, and messages that are locked by that IMS subsystem are unavailable to all other IMS subsystems in the sysplex.

Normally, when a client disconnects from CQS by using the CQSDISC request and specifying the CQSSHUT=YES parameter, CQS shuts down after no clients are connected to it. In some cases, however, the CQS address space remains active, even when no clients are connected to it. This can happen under any of the following conditions:

- No client is connected to CQS when CQS is started.
- A client that was connected to CQS terminates abnormally, without issuing a CQSDISC request to disconnect from CQS, or issues a CQSDISC request with the CQSSHUT=NO parameter specified.

You can shut down a CQS address space that has no clients connected to it by issuing the z/OS STOP command and specifying the job name of the CQS address space.

P *cqsjobname*

cqsjobname is the job name of the CQS address space that you want to stop. If no clients are connected to a CQS, that CQS shuts down. If clients are connected to the CQS, the STOP command is rejected, and message CQS0300I is issued.

Related concepts:

 CQS clients and handling special events (System Programming APIs)

Shutting down an IMSplex

If you are running IMS as an IMSplex with CSL, and you are shutting down the entire IMSplex, you should stop the CSL manager address spaces (ODBM, OM, RM, and SCI) after you shut down the IMS control region.

If you are planning to restart the IMS control region, you do not need to stop the CSL address spaces. The new IMS control region will connect to the existing CSL address spaces when it initializes.

The following briefly summarizes shutting down the IMSplex:

1. Shut down the IMSplex components that participate in that IMSplex by issuing a /CHE FREEZE or similar command to the individual IMS components.
2. Shut down CSL using one of the following methods:
 - A z/OS MODIFY command
 - A z/OS STOP command for each component
 - A CSLZSHUT request to an SCI in the IMSplex
3. Shut down CQS by issuing a z/OS MODIFY command.

Shutting down the CSL

A Common Service Layer (CSL) is comprised of multiple address spaces. As such, you can shut down either an entire CSL or individual CSL manager address spaces.

You can shut down:

- A single CSL component, such as an instance of ODBM, OM, RM, or SCI
- A CSL, including all of its components, on a single z/OS image
- A CSL, including all of its components, that spans an IMSplex across multiple z/OS images

The entire CSL should be shut down only if no IMSplex components are connected to any CSL managers. You should first issue an IMS /CHE FREEZE or similar command to terminate all the IMSplex components that might be connected to CSL prior to shutting down CSL. If a CQS address space has more than one client, it might not shut down, and you might need to shut it down after all the clients have terminated.

If you shut down an instance of ODBM, OM, or RM and there are other instances of ODBM, OM, or RM active in the IMSplex, IMSplex members can still participate in IMSplex activities; however, an IMSplex member can communicate only with the SCI with which it is registered. If that SCI is shut down, any IMSplex members on the same z/OS image as that SCI cannot communicate with other IMSplex members until that SCI is restarted.

Attention: If you are running Automatic RECON Loss Notification (ARLN), you must keep SCI active (even after IMS is shut down) until all jobs that need access to the RECON data sets have completed.

If you are shutting down only some of the CSL managers to, for example, apply maintenance, you can shut down the individual CSL managers using the STOP command. You do not have to shut down the other IMSplex components if you are only shutting down a CSL manager to apply maintenance.

To shut down the entire CSL on one z/OS image, a program can issue a CSLZSHUT request to SCI that specifies the entire CSL to shut down. Or, you can issue a MODIFY command for one CSL. The MODIFY command shuts down the CSL on the z/OS image that is associated with the SCI that receives the command. The CSL managers quiesce work and then terminate. This command is used to terminate the CSL on a single z/OS image in an orderly manner.

If you shut down the whole IMSplex, a CSLZSHUT request can be issued to SCI that specifies that the entire CSL is to be shut down. Or you can issue the MODIFY command for the entire CSL. The MODIFY command shuts down the CSLs on all z/OS images in a single IMSplex that is associated with the SCI.

Recommendation: Use the CSL SHUTDOWN command with the z/OS MODIFY command interface to shut down the CSL, rather than stopping individual components.

You can shut down the CSL by either:

- Issuing the z/OS STOP command to individual CSL manager address spaces.
- Issuing the CSL SHUTDOWN command using the z/OS MODIFY command interface.

- Using the CSLZSHUT request in a CSL application programming interface (API).

Related reference:

 CSLZSHUT: shutdown request (System Programming APIs)

Shutting down the CSL using z/OS commands

You can shut down the CSL as one unit by issuing the CSL SHUTDOWN command to any SCI in the IMSplex with the z/OS MODIFY command interface. You can stop individual modular units in the IMSplex by issuing the z/OS STOP command to the address space you want to stop.

To shut down a CSL on one z/OS image, issue the z/OS MODIFY command as follows:

```
F scijobname,SHUTDOWN CSLLCL
```

This command shuts down the CSL on the z/OS image associated with the SCI that receives the command. Use this version of the command to shut down the CSL on a single z/OS image.

To shut down an entire IMSplex, issue the z/OS MODIFY command as follows:

```
F scijobname,SHUTDOWN CSLPLEX
```

This command shuts down the CSL managers on all z/OS images in a single IMSplex associated with the SCI that receives the command.

Note: To shut down the CSL managers using the SHUTDOWN CSLPLEX command, a local SCI is required. If you issue the SHUTDOWN CSLPLEX command on a system without an active SCI, the CSL managers will not shut down.

In each of these examples, scijobname is the name of the SCI in the CSL. After it receives the command, SCI notifies other CSL managers (ODBM, OMs, and RMs) to stop, and then SCI stops. If clients are currently connected to any CSL manager and were not first stopped with a /CHE FREEZE or other command, message CSL0300I is issued, work is quiesced, and then the CSL manager stops.

Shutting down the CSL ODBM

You can shut down ODBM by using one of several methods.

The options to shut down ODBM include:

- The system operator can use a z/OS STOP command:

```
P odbmjobname
```

odbmjobname is the job name of the ODBM address space to stop.

- Issue the CSL SHUTDOWN command, which can shut down either a CSL on one z/OS image or an entire IMSplex.
- Issue the CSLZSHUT request. CSLZSHUT is a programming interface that an assembler program can issue to shut down one or more CSL address spaces.

If no clients are connected to ODBM, ODBM shuts down. If clients are connected to ODBM, message CSL0300I is issued, and ODBM quiesces in-flight work. After all work is quiesced, the ODBM address space terminates.

Before shutting down an ODBM, consider the reasons for shutting down and how shutting down ODBM can impact other IMSplex members.

Related reference:

 CSLZSHUT: shutdown request (System Programming APIs)

Shutting down the CSL OM

Before shutting down an OM, consider the reasons for shutting down and how shutting down OM can impact other IMSplex members.

Recommendation: Although you can shut down OM by itself, you should shut down OM by shutting down the CSL as one unit.

To shut down the OM, do one of the following:

- The system operator can use a z/OS STOP command:

P omjobname

omjobname is the job name of the OM address space to stop.

- Issue the CSL SHUTDOWN command.
- Issue the CSLZSHUT request.

If no clients are connected to OM, the OM shuts down. If clients are connected to OM, message CSL0300I is issued, and OM quiesces in-flight work. After all work is quiesced, the OM address space terminates.

Related reference:

 CSLZSHUT: shutdown request (System Programming APIs)

Shutting down the CSL RM

Before shutting down a Resource Manager (RM), consider the reasons for shutting down and how shutting down RM can impact other IMSplex members. There are multiple ways to shut down RM.

Recommendation: Shut down RM by shutting down the CSL as one unit.

To shut down the RM, do one of the following:

- The system operator can use a z/OS STOP command: P rmjobname rmjobname is the job name of the RM address space to stop.
- Issue the CSL SHUTDOWN command.
- Issue the CSLZSHUT request.

If no clients are connected to RM and no IMSplex-wide process steps are in progress (or no IMSplex-wide processes are in progress for an RM with no resource structure defined), RM shuts down. If clients are connected to RM, message CSL0300I is issued, and RM quiesces in-flight work. RM waits for process steps to timeout and returns the process step responses to RM clients.

For an IMSplex defined with no resource structure, RM terminates any IMSplex-wide process that is in progress and issues message CSL2210I for each terminated process. To avoid this premature termination of the process, complete all IMSplex-wide processes before RM shutdown is attempted.

If the IMSRSC repository is enabled, RM termination disconnects from any repository RM is managing and deregisters from the Repository Server (RS) address space.

Related concepts:

- ➞ Overview of the IMSRSC repository (System Definition)
- ➞ CSL RM, IMS, and Repository Server termination (System Administration)

Related reference:

- ➞ CSL SHUTDOWN command (Commands)
- ➞ CSLZSHUT: shutdown request (System Programming APIs)

Related information:

- ➞ CSL0300I (Messages and Codes)
- ➞ CSL2210I (Messages and Codes)

Shutting down the CSL SCI

Before shutting down the Structured Call Interface (SCI), consider the reasons for shutting down and how shutting down the SCI can impact other IMSplex members. There are multiple ways of shutting down SCI.

Recommendation: Shut down SCI by shutting down the CSL as one unit.

To shut down the SCI, do one of the following:

- The system operator can use a z/OS STOP command:

```
P scijobname
```

scijobname is the job name of the SCI address space to stop.

- Issue the CSL SHUTDOWN command.
- Issue the CSLZSHUT request.

If no clients are connected to SCI, the SCI shuts down. However, if there are registered members on the local z/OS image, message CSL0300I is issued, and SCI does not process any new requests or messages from local members. After all in-flight requests have completed or timed out, the SCI address space terminates.

Related reference:

- ➞ CSLZSHUT: shutdown request (System Programming APIs)

Forced termination of IMS

You normally use the /CHECKPOINT command to shut down IMS. However, in certain error situations such as a control region loop, you might need to force termination of IMS. In this case, use the z/OS MODIFY command, and be sure to request a dump of the IMS control region.

You can also use the z/OS CANCEL command, specifying the DUMP keyword, to terminate IMS. If you need to enter this command more than once, give each CANCEL time to complete before entering it again.

Attention: Using the z/OS CANCEL command more than once might stop IMS from completing termination tasks in the proper order and from releasing z/OS system resources. If IMS does not complete termination, you might need to IPL your system before restarting IMS.

Related concepts:

“Control region failures” on page 132

“Shutting down the IMS control region” on page 114

Related reference:

 MODIFY command

Related information:

 CANCEL command

Offline dump formatter

Offline dump formatting can reduce the duration of IMS outages by increasing the time required for an online IMS subsystem to terminate. You should be able to restart online subsystems faster as long as z/OS remains operational.

Any DL/I batch job can also use offline dump formatting after calling for a SYSMDUMP.

IMS can produce a machine-readable dump for later offline formatting. This can be an SDUMP, SYSMDUMP, standalone dump, a dump produced using the DUMP command, or any other machine-readable dump, such as a z/OS SVC dump.

Producing a dump using the z/OS MODIFY command

Use the z/OS MODIFY command to create a machine-readable dump rather than a z/OS or JES CANCEL command.

The CANCEL command often produces undesirable online IMS dump formatting for canceled address spaces in addition to the requested dump for offline formatting.

Producing a dump using the z/OS DUMP command

The z/OS operator can request a console dump at any time, whether or not IMS is terminating, by entering the DUMP COMM=description command.

z/OS responds with the following message:

```
xx IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
```

To which the z/OS operator must then reply with the following command:

```
xx,JOBNAME=(imsname,dbrcname,dliname,ir1lname),  
SDATA=(PSA,NUC,SQA,RGN,CSA,TRT)
```

Unless you dump all specified areas, the IMS Offline Dump Formatter might not be able to format all desired IMS subsets of the dump.

Producing a dump using the standalone dump (SADMP)

The z/OS operator can take a standalone dump of the entire system if z/OS fails. The operator receives a message when the SADMP completes.

This includes:

- Suspending operation of the system.
- Setting the address of the device containing your high-speed SADMP program, which you have generated to dump real and virtual areas.

To dump virtual areas, code PROMPT on the AMDSADMP macro.

- Performing a STORE STATUS.
- Initiating the IPL of the standalone dump program. This program prompts you for:

```
AMD001A TAPE=____
AMD011A TITLE=_____...
```

- Answering the prompts for dumping virtual storage areas:

```
AMD059D ENTER DUMP OPTIONS, 'LIST', OR 'END'.
> dump CSA,ASID('IMSjobname','DL/Iname','DBRCname','IRLMname')
AMD059D ENTER DUMP OPTIONS, 'LIST', OR 'END'.
> end
AMD010I PROCESSING ASID=  ASCB=  JOBNAME=*MASTER*
```

- Waiting until the standalone dump program completes. The following messages signal completion:

```
AMD005I  REAL DUMP DONE
AMD025I  VIRTUAL DUMPING COMPLETE FOR CSA
AMD010I  PROCESSING ASID=0001
:
AMD023I  VIRTUAL DUMP COMPLETE - 00
```

Recommendation: Do not press the EXTERNAL INTERRUPT key before the dump completes because this could prevent the IMS Offline Dump Formatter from producing a formatted dump later.

Keeping dump data sets available

z/OS console and IMS master terminal operators must keep the system dump data sets available to avoid losing dumps of the IMS subsystem. z/OS operators should transfer and process IMS-produced SDUMPs using the IEBGENER program according to your installation procedures.

If you produce IMS dumps using SYSMDUMP DD statements, the MTO should transfer and process the dump data sets according to your installation procedures.

The z/OS operators should be aware of the following IMS messages that relate to the offline dump process:

```
DFS3906I  DFSDUMP FAILED BECAUSE ALL SYSTEM DUMP DATA SETS ARE FULL
DFS3906A  REPLY "S" TO SKIP, OR "U" TO RETRY AFTER CLEARING
          A DUMP DATASET.
```

In this case, if a dump is required, reply U after clearing one or more dump data sets. If a dump is not required, reply S.

```
DFS3907I  DFSSUMP FAILED BECAUSE A DUMP IS IN PROGRESS
DFS3907A  REPLY "S" TO SKIP, OR "U" TO RETRY AFTER THE CURRENT
          DUMP COMPLETES.
```

In this case, if the dump is required, reply U after you receive the following message:

```
IEA911E  COMPLETE/PARTIAL DUMP ON SYS1.DUMPxx.
```

If a dump is not required, reply S.

```
IEA793A  NO DUMP data sets AVAILABLE FOR DUMP=nnn BY JOB (imsproc).
USE THE DUMPPDS COMMAND OR REPLAY D TO DELELTE THE DUMP
```

In this case, the dump is written to the dump data set after clearing one or more dump data sets. If a dump is not required, reply D.

Database resource adapter storage

You can use the database resource adapter (DRA) to access an IMS DB subsystem. In this instance, the DRA code actually runs in the CCTL region or z/OS application region, not in any of the IMS DB regions.

A shutdown of the IMS DB subsystem with a dump does not produce any information about the DRA.

Chapter 5. IMS failure recovery

When a failure occurs, you must perform the correct recovery and restart procedures to ensure system integrity. You can use the IMS control region Extended Specified Task Abnormal Exit (ESTAE) routines to clean up resources after an IMS abend.

The reason for the abnormal termination and the results of cleanup processing have a great effect on the subsequent restart process. The processing sequence of these routines is as follows:

- The storage management task ESTAE routine writes recovery data to the restart data set (RDS) and then closes the RDS.
- The IMS control task ESTAE routine terminates all dependent regions and issues message DFS629I.
- All ESTAE routines attempt to purge the OLDS log buffer and to write a termination log record.
- The ESTAE routines tell DBRC to sign off abnormally.
- The ESTAE routines tell IRLM to quit, and either retain or release locks, based on whether DBRC indicates that databases have been updated.
- The ESTAE routines disconnect XRF alternate subsystems from the z/OS availability manager to perform I/O prevention.

In addition, the IMS resource cleanup module cleans up resources whenever IMS terminates. It releases the common storage area (CSA), closes DBDSs, deletes the subsystem interface control block, closes the VTAM Access-method Control Block (ACB), and issues message DFS627I.

If you use the DL/I address space (you specified LSO=S in the IMS procedure), IMS issues message DFS603I to indicate that cleanup is successful.

z/OS system failures

If IMS terminates because of a z/OS failure or a hardware or power failure, you must IPL z/OS and then restart IMS using an /ERESTART command.

Database I/O errors, occurring as a result of, or at the same time as, a hardware failure, can leave damaged data on DASD. If IMS terminates abnormally after this kind of an I/O error before IMS can record the I/O error on the IMS log and in DBRC, IMS does not know that there is damaged data. And if IMS does not know that there is damaged data, your data integrity is compromised.

If you experience or suspect I/O errors that lead to these kinds of failures so that it is necessary to use the /ERE OVERRIDE command, check and, if necessary, recover all databases allocated for dynamic backout or DEDB Redo during emergency restart.

If APPC/MVS fails, restart APPC/MVS and issue the IMS /START APPC command to establish the interface between IMS and APPC/MVS.

Control region failures

Software errors causing an abend or a loop, or failures relating to system data sets (including the log data sets) can cause control region failures. After determining and correcting the cause of the failure, you can restart the system with the /ERE command.

If the failure results in a loop in the control region, you might need to force termination of the region with the z/OS MODIFY command.

When the IMS control region fails, APPC/MVS terminates LU 6.2 conversations with the DEALLOCATE abend.

When the IMS control region fails, the CQS subsystem to which it is connected does not shut down. If you need to shut down the CQS subsystem, use the z/OS STOP or CANCEL command.

Related concepts:

“Forced termination of IMS” on page 126

Emergency restart failures

Emergency restart backs out incomplete changes from programs that were active at the time of failure. However, damaged logs or message queues might prevent emergency restart from completing. If this occurs, you might be able to perform an emergency restart from a previous checkpoint or by specifying FORMAT ALL on the /ERESTART command.

If you are running an IMS DB/DC system, and an emergency restart fails, you do not have to cold start the entire system. You can use the commands /ERE COLDCOMM and /ERE COLDBASE. These commands cold start part of the system while restarting the other part: /ERE COLDCOMM cold starts the DC part and restarts the rest, while /ERE COLDBASE cold starts the DB part and restarts the rest.

If you are running an IMS DBCTL or DCCTL system, and an emergency restart fails, you must cold start the entire system using the /ERE COLDSYS command.

Re-establishing database integrity

If no emergency restart succeeds, an appropriate response depends on your understanding the reason for the failure. Although it might not work exactly as written for your installation, use the following procedure to reestablish database integrity.

1. Run the Log Recovery utility (DFSULTR0) to close the OLDS (CLS mode). If the utility needs additional logs, concatenate them in your JCL, oldest first, and rerun the utility.
2. Be sure that your log is correct and that your databases are consistent with the log.
3. Produce an SLDS using the Log Archive utility (DFSUARC0).
4. Forward recover all updated DEDBs to make them consistent with the log using the Database Recovery utility (DFSURDB0).

Recommendation: If the recovered area utilized the shared VSO option, ensure that there are no failed-persistent XES connections to the CF structure used by

the area. If necessary, use the z/OS SETXCF command to delete the connections prior to cold start of the systems involved.

5. For full-function databases, run the Batch Backout utility (DFSBB000) to back out programs (PSBs) that were active at the time of failure. Determine these from the output of the Log Recovery utility.
6. Update the RECON data set, if necessary, to reflect the current status of the system.
7. If you use MSDBs, rebuild the MSDBINIT data set from the last used MSDBCP1 and MSDBCP2 data sets and the last SLDS.

At this point, database integrity should be reestablished. If you are using DBRC or data sharing, additional actions might be required before you perform a cold start.

System data set failures

If an IMS system data set fails because of an unrecoverable I/O error, you must scratch, reallocate, and reformat the data set before you restart the system using the /ERESTART command.

Message queue data set failures

Two types of problems can occur with message queue data sets: they can run out of space, or an I/O error can occur. In either case, you need to scratch and reallocate the problem data set, and increase its size if necessary.

Note: This topic does not apply to a DBCTL environment or to a shared-queues environment.

IMS automatically reconstructs the contents of the message queues when you restart IMS and specify the BUILDQ and FORMAT keywords on the restart command.

Restriction: IMS can only reconstruct the message queue data sets from the initial cold start or from a SNAPQ or DUMPQ checkpoint.

To recover message queues during IMS restart, use the /NRE BUILDQ or /ERE BUILDQ command. If you also want to reinitialize the message queue data sets (that is, reformat them with null records), use the FORMAT keyword on the restart command.

To use the /NRE BUILDQ command, you must have previously shut down IMS using a /CHECKPOINT DUMPQ | PURGE command. You can use the /ERE BUILDQ command to recover messages from a prior SNAPQ checkpoint (/CHECKPOINT SNAPQ).

If an emergency restart fails, you must initiate an emergency cold start using the /ERE COLDSYS command. An emergency cold start does not perform any database recovery, so you must close and archive the last OLDS.

If an IMS restart fails, you can use the IBM IMS Queue Control Facility for z/OS (QCF) to recover the message queues if you use one of the following cold start commands after the failed restart: /ERE CHECKPOINT 0, /ERE COLDCOMM, or /ERE COLDSYS. With QCF, you can select messages from the OLDS or an SLDS and requeue them to the message queues after IMS restarts (cold start). QCF provides recovery modes that analyze and select the messages to be requeued.

Related reference:

 [IMS Queue Control Facility overview](#)

Other system data set failures

You should make periodic backup copies of IMS system data sets so that if an error occurs, you can recover your data or recreate it from a backed up copy.

If an error occurs on an IMS data set, such as IMS.ACBLIBx or IMS.FORMATx, you need to recover it. If you have made periodic backup copies of the data set, you can use the latest copy. If you have not changed the content of the data set since making the copy, you can use it as is. If you have changed the content of the data set since making the copy, for example, you have added ACB generations since copying the ACBLIB data set, you need to redo the changes before the backup copy can be considered up-to-date.

If you have not made backup copies of the data set, you need to recreate it. For example, if you lose the active format library (IMS.FORMATx) and you have no backup copy of it, you need to rerun Message Format Services and redo all your format definitions.

If you have inactive data sets from an online change, you can recover those data sets involved with online change (MODBLKS, ACBLIB, and FORMAT). In this case, you must reapply any changes made to the system data sets.

IMS does not offer any specific utilities or commands to recover system data sets.

RECON data set recovery

If an I/O error occurs on a RECON data set, DBRC tries to use an available spare data set. DBRC then copies the healthy RECON data set to the spare data set, and then activates the spare data set.

When other subsystems using the failed RECON data set complete their processing, you can delete and redefine the error data set so it can be used as a spare. If, however, you want to analyze the RECON error, you should allocate new space for the RECON data set, rather than deleting and redefining it.

If DBRC cannot locate a spare data set, all currently executing jobs continue to process using the RECON data set in single mode. DBRC allows new jobs to start with only one RECON data set if you specify the STARTNEW keyword on the INIT.RECON or CHANGE.RECON command (or the DSPAPI FUNC=COMMAND COMMAND=INIT.RECON or DSPAPI FUNC=COMMAND COMMAND=CHANGE.RECON API requests).

Attention: Do not allow jobs to start with only one RECON data set because it jeopardizes the integrity of the system.

Restoring RECON data sets if both are unusable

It is unlikely that both RECON data sets would be unusable. However, if this situation occurs, you must back up, delete, and redefine your RECON data sets, make an image copy of all applicable DBDs, and close any open, out-of-date OLDSs.

To restore and resynchronize the RECON data sets with the databases:

1. Stop all jobs that require access to the RECON data set.

2. Optional: If you can access both RECON data sets, use the VSAM access method services REPRO command to back them up. This step is optional, but recommended.
3. Use the VSAM access method services utility to delete and redefine your RECON data sets.
4. Issue the access method services REPRO command to restore one of the RECON data sets.
5. Issue the REPRO command to restore the other RECON data set from the first.
6. Issue the DBRC LIST.RECON command or the DSPAPI FUNC=QUERY API request to list one of the RECON data sets. Based on the list, determine which DBDSs IMS updated since you made the backup in Step 2. If you cannot determine which DBDSs have been updated, assume that all have been updated.
7. Issue the DBRC CHANGE.IC command (or the DSPAPI FUNC=COMMAND COMMAND=CHANGE.IC API request) with the INVALID keyword to mark all image copy records in error for all applicable DBDSs in Step 6.
8. Make an image copy of all applicable DBDSs from Step 6.
9. Issue the DBRC BACKUP.RECON command (or DSPAPI FUNC=COMMAND COMMAND=BACKUP.RECON API request) to make a backup copy of the RECON data sets.
10. Close any open, out-of-date OLDSs using the NOTIFY.PRILOG command (or DSPAPI FUNC=COMMAND COMMAND=NOTIFY.PRILOG API request) to clean up the new RECON data set.

The RECON data sets are now restored and resynchronized with the databases.

Note: If DBRC does not manage a large number of databases, it might be easier to use the following procedure:

1. Stop all jobs that require access to the RECON data set.
2. Define new RECON data sets.
3. Initialize these RECON data sets.
4. Register the environment (always keep a backup copy of the most recently initialized, but not yet used, RECON data set available).
5. Take image copies of all databases.

Log errors

You can experience three types of errors during logging. A write error on the OLDS, a read error on the OLDS or SLDS, and a read or write error on the WADS can occur.

OLDS write error

When an OLDS has a write error, IMS stops the OLDS (or a pair of OLDSs) in error. No specific MTO interaction is required. When an OLDS is no longer needed for dynamic backout, IMS dynamically deallocates a stopped OLDS.

If you are using single logging, IMS switches to the next OLDS. If the error causes only two OLDSs to remain available, IMS terminates. If an error occurs and no other OLDS is available, IMS abends with code U0616.

If you are using dual logging and an error occurs (either to one or both volumes in the pair), IMS switches to the next OLDS pair and continues. What happens next depends on what is specified on the DEGRADE control statement of the DFSVSMxx member.

DEGRADE=YES

Specifies that IMS will continue to use all good pairs of OLDSS until only two pairs remain. At this point, IMS switches (degrades) to single-logging mode, using whatever good data sets are left from each pair of OLDSS. When only two good OLDSS remain, IMS terminates (as it does in normal single-logging mode). DEGRADE=YES is the default.

DEGRADE=NO

Specifies that IMS will terminate when only two good pair of OLDSS remain.

Recommendation: Use the /START OLDS command to allocate new pairs of OLDSS.

OLDS or SLDS read error

When read errors occur on the OLDS or the SLDS during emergency restart, IMS abends. When read errors occur on the OLDS during dynamic backout, backout fails.

WADS error

When a WADS has an I/O error, IMS switches to another WADS if one is available. If none is available, processing continues without a WADS. IMS maintains log-write-ahead protocols by truncating log buffers. As with OLDS errors, no specific MTO interaction is required. You should scratch and reallocate any WADS with a write error after you shut down IMS.

Log error recovery

Problems can occur with any of the log data sets (OLDS, SLDS, WADS, restart data set (RDS), and RLDS). Your actions depend on the log, the type of error, whether single or dual logging was in effect when the error occurred, and what processing was being performed when the error occurred.

Log Recovery utility (DFSULTR0)

The Log Recovery utility produces a usable log data set from an OLDS or SLDS that contains read errors or was not properly closed.

The Log Recovery utility has four modes of operation:

- CLS mode, to close an OLDS
- DUP mode, to create:
 - Create an interim log containing error records
 - Create a closed batch SLDS containing an end-of-file mark
- REP mode, to:
 - Read the interim log
 - Replace the error records with data that you specify
 - Create a new log
- PSB mode, to produce a report of active PSBs from the input log

|

When IMS opens the input logs, DBRC validates them, ensuring that the data set name and volume serial numbers agree with those in the RECON data set. For log recovery, DBRC requires that the entire log data set be specified as input to the IMS Log Recovery utility, not just selected volumes. For a closed SLDS, DBRC requires only the last volume as input to the Log Recovery utility.

When IMS opens the output logs, it calls DBRC to do the following:

- Create an IPRIOLDS record (and ISECOLDS record if you are using dual logging) using the time stamp of the input log and the data set name from the job file control block of the output log. Applies to DUP mode.
- Create new OLDSs from interim OLDSs. Applies to REP mode.

DBRC issues the GENJCL.CLOSE command to generate a job to run the Log Recovery utility. When you issue this command, specify the subsystem ID of the IMS subsystem that created the OLDS to be closed. You can also specify the OLDS to be closed. If you do not specify an OLDS, DBRC closes the most recent, open OLDS. You can issue the GENJCL.CLOSE command using the Recovery Control utility or as an IMS command.

Related reference:

 Log Recovery utility (DFSULTR0) (System Utilities)

OLDS recovery

IMS automatically closes the OLDS during normal shutdown or during emergency restart. You must close an OLDS before you can archive it or use it as input to any utility.

You must close the OLDS using the Log Recovery utility in either of the following circumstances:

- When an emergency restart fails and, rather than performing another emergency restart, you perform a cold start
- When IMS does not close the OLDS because IMS detected a write error (single logging only)

The Log Recovery utility recovers the OLDS from the following types of errors:

- An I/O error while reading the input log data set
- An error in the log record
- A sequence error in the log record, the log block, or the OLDS write time stamp

To recover the OLDS, the Log Recovery utility does the following:

- In CLS mode, the utility closes an input OLDS from information in the WADS if IMS cannot close it because of a system failure. Otherwise, the utility closes the OLDS from the OLDS used immediately after the OLDS in error. The utility uses the immediately prior OLDS, if any, to establish a base point for close processing, using the last block sequence number.

If dual logging is in effect, IMS must close both OLDSs. If IMS successfully closes the logs, you do not need to use the Log Recovery utility.

- In DUP mode, the utility reads an OLDS and duplicates all readable records on a log.
- In REP mode, the utility reads the interim log created in DUP mode, copies good log blocks, and replaces error blocks with good ones based on information that you specify in control statements. The output log data set is a usable OLDS.

If the OLDS being recovered has not been closed (DBRC shows a stop time of zero), you must use the output from REP mode as input to CLS mode.

SLDS recovery

The Log Recovery utility closes a system log data set (SLDS) created by an IMS batch job. You must close an SLDS before you can use it as input to any utilities or to an IMS restart.

The Log Recovery utility recovers the SLDS from the following types of errors:

- An I/O error while reading the input log data set
- An error in the log record
- A log record sequence error

To recover an SLDS, the Log Recovery utility reads the SLDS based on the mode that it is in:

- In DUP mode, the utility reads an SLDS and duplicates all readable log records onto an interim log. When the utility encounters the first error or end of file (EOF), it stops copying and closes the SLDS. When you specify a nonzero error count, DUP mode writes error blocks and error ID records on the interim log until either the error count reaches the number you specify or EOF is reached. The utility uses the interim log containing these error blocks and error ID records as input to REP mode.

You can close an SLDS by using DUP mode and specifying the error count as zero. REP mode is not required.

- In REP mode, the utility reads the interim log created by DUP mode, copies good log blocks, and replaces error blocks with good ones based on information that you specify in control statements. The output log data set is a usable SLDS.

WADS or RDS log recovery

You cannot recover a write-ahead data set (WADS) or a restart data set (RDS). If a problem occurs with either type of data set, you must scratch, reallocate, and format it during an IMS restart.

Dependent region failures

Two types of dependent region failures can occur: an application program failure or a region controller failure. In each case, if an LU 6.2 conversation established by the application program is running in the failing dependent region, APPC/MVS terminates the conversation with the DEALLOCATE abend.

- Application program failures, in which the application program abends but the region is still active.
- Region controller failures, in which the region fails. This kind of failure can occur, for example, when the program loops.

Application program failures

When an IMS application program abends, IMS issues message DFS554I (and in some cases message DFS555I) and then backs out all database changes made by the failing program to the last checkpoint, or to last point at which the program was scheduled. You must use the /START PROGRAM and /START TRANSACTION commands to restart the program and transaction after correcting the cause of the failure.

If an error occurs during database backout and restartable backout does not correct the problem, follow the procedures for database recovery before reactivating the program and transaction.

Region controller failures

One symptom of a dependent region failure is an increase in the number of transactions on the message queue. If the region appears to be looping, stop the region with the `/STOP REGION ABDUMP` command. If the region does not stop, enter the `/STOP REGION CANCEL` command.

You can monitor the status of your regions with the `/DISPLAY ACTIVE`, `/DISPLAY Q TRANSACTION`, and `/DISPLAY Q BALGRP` commands.

z/OS issues messages IEF450I and IEF404I to indicate an abend in a BMP, MPP, or IFP region. If these messages are not followed by a DFS554I message, issue the `/STOP REGION ABDUMP` command so that IMS recognizes that the region has terminated and backs out any database updates.

After you determine and correct the problem that caused the region to loop or that caused the abend, use the `/START` command to start the transaction and application program again.

Restart of BMP regions

If the application program uses the IMS restart call (XRST), you can restart the program from its latest checkpoint. You can identify this checkpoint from the DFS681I, DFS395A, or DFS682I messages.

DFS681I

This message lists the checkpoint ID and the PSB name, and is issued when the program takes a checkpoint.

DFS395A

The Database Batch Backout utility issues this message when it cannot restart.

DFS682I

This message lists the checkpoint ID and the program name, and is issued after emergency restart or after dynamic backout following an abend.

Specify the checkpoint ID value or LAST in the PARM field of the JCL used for restarting the BMP.

To restart a BMP after a cold start of IMS, use the IMS Log Analysis utility to scan for X'37' log records that contain BMP restart information for the desired BMP (job name, PSB name, program name). The last X'37' log record printed contains the information needed to restart the BMP region.

Restart of Fast Path regions

If a Fast Path message-driven application program or DEDB online utility program abends, you can restart the Fast Path region immediately. You can restart the region with the `/START REGION` command.

You should provide the MTO with documentation of the names of the procedures that are started using the `/START REGION` command because the procedure name must be given with this command. You should keep this documentation current because these regions require system resources, such as virtual storage and database buffers.

Database failures


IMS issues message DFS0451I or DFS0451A when there is an error in an IMS database. Every time IMS encounters a DL/I I/O error, IMS creates an extended error queue element (EEQE) that identifies the block or VSAM control interval in error.

When IMS closes a database, it automatically retries read and write errors on DL/I databases. If successful, forward recovery of the database is not required. Otherwise, forward recovery is eventually required. It might be possible to defer recovery to a more convenient time. Deferring recovery does not inhibit scheduling access or updating.

Using DEDB multiple area data sets also allows application programs to continue when I/O errors exist. For DEDB I/O errors, IMS issues messages DFS2571, DFS2572, DFS3712, and DFS3713. If a DEDB area is not available, the application receives an FH status code.

IMS maintains I/O information and buffer images across restarts. IMS does this by recording the EEQEs in DBRC, notifying all systems that are using the IRLM, and, during initialization and checkpoint, logging EEQEs and virtual buffers to the OLDS.

Related concepts:

 Database failures (Database Administration)

Database recovery

The primary methods for recovering a database after a failure are *forward recovery* and *backward recovery* or backout. Two methods of database recovery are possible for each type of recovery.

Forward recovery involves reconstructing the database from a copy of the database made prior to the database failure. Forward recovery uses the information you have been keeping, such as image copies, logs, and so forth, and reapplies it to the database copy. It is based on the notion that if you knew what the data was like at one time and you know what you have done to it since then, you can process the data to return the database to the state it was in just before it was lost.

IMS supplies two methods for the forward recovery of databases:



- The Database Recovery utility (DFSURDB0).
- If you have the Database Recovery Facility or Online Recovery Service products installed, you can use the /RECOVER command.

Backout allows you to remove incorrect or unwanted changes from existing information or work without rebuilding from a prior copy of the database.


There are two types of database backout:

- Dynamic backout, in which IMS automatically backs out changes to a database, usually after an application program error.
- Database batch backout, in which the Batch Backout utility (DFSBB000) removes database changes made by IMS batch jobs and online programs.

Related concepts:

-  Recovery of databases (Database Administration)
-  Database backout (Database Administration)

Related reference:

-  Database Recovery utility (DFSURDB0) (Database Utilities)

Recovering from a network failure when the remote terminal stops responding

If IMS stops responding to a remote terminal, you might be able to recover the session if the IMS master terminal is still active. The MTO receives a message that indicates why the terminals are not operational if they have not been stopped by the /CLSDST, /STOP, or /PSTOP commands. Correct the error and issue a /START or /RSTART command to restore operation for the terminals.

You can use the various forms of the /DISPLAY command to list the status for communication lines, terminals, nodes, and ETO dynamic users.

- Use the /DISPLAY LINE command for communication lines
- Use the /DISPLAY LINE PTERM command for terminals
- Use the /DISPLAY NODE command for nodes
- Use the /DISPLAY USER command for ETO dynamic users

For a VTAM terminal session, you might only need to use the /OPNDST NODE command to restart it.

If you have a hung node, you can free it with a /CLSDST FORCE command, if the following conditions exist:

- A VTAM display indicates that no session exists.
- The output from a /DISPLAY NODE command indicates that a CID exists, the node is connected, and the node is not idle.

If a session does exist, terminate it with the VTAM command VARY INACT,FORCE, but use caution, because this command terminates all parallel sessions.

If a static node or ETO dynamic user is hung in Fast Path input response mode, as indicated by a “RESP-INP-FP” status in the output of the /DISPLAY NODE or /DISPLAY USER commands, you can reset the node or user with the appropriate /STOP and /START command sequence. For static nodes, issue /STOP NODE and /START NODE. For ETO dynamic users, issue /STOP USER and /START USER.

If a static node or ETO dynamic user is hung in full-function input response mode, as indicated by a “RESP-INP” status in the output of the /DISPLAY NODE or /DISPLAY USER commands, you can reset the node or user with the /DEQUEUE command to discard response-mode output so that the /RSTART command can reset terminal response mode.

If all terminals are operational and appear to function, then it is likely that an application program has terminated abnormally, or is in a loop, or has a resource conflict that prevents its scheduling. To find the cause of the problem, use one of the following commands: /DISPLAY ACTIVE, /DISPLAY PROGRAM, /DISPLAY TRAN, or QUERY TRAN.

If you use `/DISPLAY ACTIVE` and continually see a program listed as active in a MPP or BMP region, that program is probably in a loop. For a program that is continually active, check that its PROCLIM value (specified in the PSB) is correct. Use the `/STOP REGION ABDUMP` command to terminate the region, or use the `/STOP` command to stop programs and transactions. You should be careful when using the `/STOP REGION CANCEL` command because in certain instances it can cause an abend of the IMS control region.

CPI Communications failures

For a CPI Communications driven LU 6.2 conversation, IMS TM is not involved and places no restrictions on your choice of either committing or backing out updates.

Session failure

In the event of a session failure, a CPI Communications driven application program can provide full integrity by issuing commit (SRRCMIT) or a backout (SRRBACK) calls. If an LU 6.2 session fails during an LU 6.2 conversation, you can end the conversation or continue processing.

If the session fails before the transaction completes phase one of the two-phase commit sync-point protocol, IMS takes action as follows:

- For both standard and modified standard DL/I application programs, if the original conversation is asynchronous, IMS commits all database changes, sends all output messages asynchronously, and deallocates any established conversations.
- For standard DL/I application programs:
 - If the conversation is synchronous and a user destination exit routine is provided, the routine determines whether to abort or commit. If it commits, IMS sends all output messages asynchronously and commits all database changes. If it aborts, or if no user destination exit routine is provided, IMS aborts the transaction and backs out all changes.
- For modified standard DL/I application programs:
 - If the user destination exit routine commits, IMS commits database changes and inserts all messages to the IMS message queue asynchronously.
 - If the original conversation is synchronous and the user destination exit routine aborts, IMS aborts the transaction, backs out all database changes, and discards all messages inserted to the IMS message queue (unless the insert used an express PCB).

If the session fails after the transaction starts phase two of the two-phase commit, IMS sync-point processing continues for both standard and modified DL/I application programs, despite a session failure.

System failure

When the IMS system fails during a CPI Communications driven LU 6.2 conversation, IMS determines whether to resolve in-doubts for IMS-protected resources. Examples of IMS-protected resources are IMS DB databases, DB2® for z/OS databases, and IMS TM messages.

For standard and modified standard DL/I application programs, if the IMS system fails before the transaction completes phase one of the two-phase commit sync-point protocol, IMS backs out changed data during IMS restart. This backout

includes all discardable transactions that were processing at the time of the system failure. IMS requeues non-discardable transactions so they can be processed.

If the transaction completes phase one of the two-phase commit, IMS resolves in-doubts during IMS restart. If only IMS resources are affected, IMS commits changes. If DB2 for z/OS resources are affected, IMS tells DB2 for z/OS to commit or abort the changed data, as appropriate.

Recovery processing for CPI Communications driven application programs

If your application programs require recovery assistance, they should use the implicit API provided by IMS rather than the CPI API. Implicit support allows application programs that do not normally use LU 6.2 protocols to use LU 6.2 devices.

The implicit API support is the original IMS DL/I API (using *xxx*TDLI calls, where *xxx* represents the programming language used).

IMS provides no recovery processing for application programs that use the explicit CPI application programming interface (API). IMS discards all messages for and from CPI Communications driven application programs at IMS restart, regardless of their state at the time of failure. Application program designers should use the SAA resource recovery resynchronization functions for explicit CPI application programs.

Related reference:

 DL/I calls for IMS and CICS (Application Programming)

MSC VTAM message resynchronization and recovery

For IMS systems linked by VTAM, IMS maintains the integrity of messages and control blocks during system and session failures. IMS recovers sessions automatically if both half-sessions save resynchronization information.

Because a logical link is broken after an abend caused by a system failure, user action is required for recovery. Performing a normal cold start destroys session resynchronization information, so you should use the /ERESTART command to restart IMS.

During emergency restart, IMS restores the message queues. After the link is reestablished, IMS resynchronizes messages between half-sessions.

In contrast, the message queues are lost during a cold start because IMS cannot restart the link at the point at which the system failed. Messages that were lost must be resubmitted.

After you reestablish a VTAM link after a session fails, you must reassign as primary the half-session originally assigned as primary (the one from which you issue the /RSTART command or the type-2 UPDATE MSLINK(*linkname*) START(COMM) command).

A z/OS system can run more than one IMS subsystem at a time. One IMS subsystem communicates with the other subsystem using real-storage-to-real-storage communication. z/OS provides the real-storage-to-real-storage connection primarily for system protection and testing purposes.

If a physical link between systems fails, the MTO should reestablish communication through an alternate physical link:

1. Use the /MSASSIGN LINK command or the type-2 UPDATE MSLINK(*linkname*) SET(MSPLINK(*msplinkname*)) to reassign the logical link to an alternate physical link. This command must be entered on both systems.
2. Use the /RSTART LINK command or the type-2 UPDATE MSLINK NAME(*linkname*) START(COMM) command to start the logical link.

CQS failures and structure failures

The CQS subsystem automatically recovers from most types of failures. Because it is registered with the z/OS Automatic Restart Manager, if the CQS subsystem fails, z/OS restarts it. If one of the structures in the coupling facility fails, CQS initiates a structure rebuild (using data from the CQS log) to recover the structure.

If you lose connection with the last remaining coupling facility, IMS stops sharing the message queues and stops to work because there are no queues available for input or output. After the connection is restored, CQS reconnects to the structures and rebuilds them if necessary, and IMS resumes work.

In the case of a resource structure, a coupling facility list structure, CQS does not support structure recovery, structure checkpointing, or overflow processing. To provide optimum performance, CQS does not log changes to resource structures. To aid in recovering resources, you can define resource structures in the CFRM policy to duplex them automatically.

In the case of a nonrecoverable queue structure (RECOVERABLE=NO), CQS does not support structure recovery. When a nonrecoverable queue structure fails, CQS begins recovery processing, and cold starts the newly-allocated structure (initializes it to empty). All active CQSs that were connected to the failed structure then abend with an ABENDU0373. You should restart the CQSs, at which time, the CQSs performs resync processing with any connected clients. You can manually restart CQS, use the z/OS ARM feature (ARMRST=Y), or use your own automation. When CQS restarts after the failure of a nonrecoverable queue structure, all data in the previous (failed) structure is lost.

CQS log recovery

CQS uses a z/OS log stream for its recovery log. You cannot recover if the last offload data set in the z/OS log stream is full and CQS fails. Be sure to schedule CQS system checkpoints frequently enough to prevent your offload data sets from filling. You can define groups of offload data sets for the system log, and z/OS uses each group in turn so that you do not run out of data sets. CQS can recover your system by using these data sets.

CCTL failures

CCTL failures commonly include CCTL region failures, in which a program is not stopped, and when a program is stopped

CCTL region failure

If a CCTL fails, it normally disconnects from the IMS DBCTL system and has no effect on the IMS DBCTL system itself. The CCTL threads are terminated immediately or after they complete their current DBCTL request. These thread terminations appear as DFS554I messages. The programs are not stopped.

What action IMS takes for database changes made (units of recovery), depend on the sync-point state of the thread. If a CCTL fails, a U113 abend of the DBCTL control region can occur.

Related reference:

 CCTL exit routines (Exit Routines)

CCTL thread failure

A CCTL thread failure can occur if the CCTL thread makes a database resource adapter (DRA) request, and IMS ends abnormally during that request.

When this happens, IMS issues a DFS554I message, and the CCTL indicates the failure to its operator. The program is stopped; you can start it again using the /START command after the problem is resolved. If a CCTL thread fails within the CCTL itself, IMS takes no action.

A thread can also fail because of an abend while the thread is within IMS or the DRA. For example, when the IMS DBCTL system fails, all CCTL threads end abnormally in the DRA with a U002 abend. What happens to the UOR associated with the CCTL thread is determined by its sync-point state. If the failure occurs while UOR is in-flight, IMS backs the UOR out; if the UOR is in-doubt (in the middle of the sync-point process), the UOR remains in-doubt.

Related concepts:

 Unit of recovery (Application Programming)

CCTL thread looping

Use CCTL commands to determine if the CCTL is looping or if it has issued a DBCTL request and is waiting for a response. Refer to your CCTL documentation if the CCTL is looping. Otherwise, issue the /DISPLAY command to find the region with the recovery token that corresponds to the CCTL task that is waiting for a response from DBCTL. Then use the /STOP REGION command to force that thread to end.

DBCTL failures

A termination of a IMS DBCTL system does not cause connected subsystems to terminate; the subsystems are simply left without DBCTL services. If any of the DBCTL address spaces (DBC, DBRC, or DLISAS) fails, all of its address spaces are terminated.

The state of the DRA is determined by how the CCTL responds to notification of the IMS DBCTL system failure.

Normally, you terminate the IMS DBCTL system using a /CHECKPOINT FREEZE command, but you can also use the z/OS MODIFY command to force termination, particularly if an IMS shutdown or failure results in a loop. If you use the MODIFY command and want a dump of the address space, use the DUMP keyword; otherwise use the STOP keyword. The DBCTL control region terminates with a U0020 abend, and IMS issues the following messages: DFS628I and DFS629I. These messages tell you that an abend is scheduled and what the job name is.

After determining the cause of the failure and correcting it, restart the IMS DBCTL system using the /ERESTART command. The CCTL cannot reconnect to the IMS DBCTL system until the /ERE command completes. The CCTL has several options on how to respond to a DBCTL failure.

When a CCTL connects to an IMS DBCTL system after a restart, they resynchronize in-doubt UORs automatically. Resynchronization means that the CCTL requests an action (commit, abort, or forget) and the IMS DBCTL system processes that request. If the commit or abort process fails, IMS takes the same actions as in the DB/DC environment. If the commit process fails for systems with DEDBs, IMS sends message DFS2282I to the DBCTL operator. For all failures, IMS removes the EEQEs because the UOR is no longer in-doubt.

If a CCTL requests resynchronization of an in-doubt UOR, and the DBCTL subsystem has no knowledge of that UOR, IMS returns code 218 to the CCTL, and sends message DFS2283I to the DBCTL operator. The DBCTL operator can use the /DISPLAY command to display the in-doubt UORs known to the IMS DBCTL subsystem.

Related concepts:

 Recovery in an IMS DBCTL environment (System Administration)

Related reference:

 CCTL exit routines (Exit Routines)

IRLM failures

When an IRLM fails, IMS subsystems using the IRLM cannot continue normal operation. IMS terminates active programs using the IRLM with abend 3303 and puts their messages on a suspend queue.

For wait-for-input programs using the IRLM, IMS either terminates them with abend 3303 on their next database call or gives them a QC status code. IMS inhibits subsequent program scheduling until IMS reconnects with the IRLM.

Before IMS can back out the IRLM failure, every dependent region with an intent to use a database must terminate its threads. IMS cannot terminate or reconnect to the IRLM until all dependent regions disconnect from the IRLM. Therefore, you must abnormally terminate any region that has not terminated (such as those programs that are waiting in an for a timer expiration or a for a reply to a Write-to-Operator with Reply (WTOR) message).

Restart the IRLM using the z/OS START command. Then, reconnect IMS subsystems to the restarted IRLM using the z/OS MODIFY`jobname`RECONNECT command. IMS automatically dequeues messages on the suspend queue.

Recovery with data sharing

In a data-sharing environment, you must extend IMS recovery procedures by protecting surviving subsystems from database records that contain incomplete changes left by a failing subsystem. IMS recovers a shared database by restoring the database from the latest image copy.

You can use the merge function provided with the IMS Database Change Accumulation utility (DFSUCUM0) to merge the log data sets from the subsystems that updated the database.

Apply the records from the utility (in the output data sets) for a forward recovery. If only one subsystem in a data-sharing environment updated the database, IMS does not need to use the sequence numbers and the Database Change Accumulation utility, because it has only one set of log data sets to manage.

IMS data sharing also introduces other challenges. Two IMS Internal Resource Lock Managers (IRLMs) might lose communication with each other because:

- An z/OS cross-system coupling facility (XCF) might fail.
- An IRLM might fail, leaving its IMS subsystems running.
- A z/OS system might fail, bringing down the IRLM, VTAM, DBRC, and IMS subsystems running under it.

In a data-sharing environment, IMS must ensure data integrity after failure of one or more components of data sharing.

The IRLM and DBRC work together with IMS to protect shared databases before and after a failure leaves incomplete changes in these databases. This protection remains after a failure until failed subsystems complete backout of incomplete changes.

Related reference:

 Database Change Accumulation utility (DFSUCUM0) (Database Utilities)

DBRC and protecting data

DBRC provides database-level (level 1) data sharing and database-level protection after a failure occurs. DBRC prevents conflicts in database allocation that could compromise database integrity. DBRC uses the RECON data set to record information about each database authorization it grants to IMS subsystems, and to indicate whether a subsystem is updating a database.

DBRC removes this indication when any of the following occur:

- The subsystem terminates normally.
- Dynamic backout completes after a failure.

If an IMS subsystem fails during an update, the RECON data set still indicates this incomplete update. DBRC does not grant additional database authorizations unless the IRLM retains locks that protect the incomplete changes.

IRLM and protecting data

The internal resource lock manager (IRLM) can protect your data if an IMS subsystem fails.

If the IRLM is present and remains running, it retains locks on incomplete changes at the database record level after an IMS subsystem fails. The IRLM removes these locks when dynamic backout completes.

If the IRLM fails while its IMS subsystems continue to run, the subsystems stop all work in progress, dynamically back out all incomplete database changes, and call DBRC to relinquish their authorizations to shared databases. After you restart the IRLM, the subsystems resume processing.

If two or more IRLMs are active during block-level data sharing, and one IRLM fails with or without the z/OS system that it runs under, the surviving IRLMs and their IMS subsystems reauthorize with DBRC the use of all shared databases.

In this case, because the surviving IRLMs know about the block-level locks (that protect the incomplete changes made by the failed subsystem), they permit their IMS subsystems to continue block-level sharing of the affected databases.

Fast Database Recovery (FDBR) regions

In a sysplex data-sharing environment, if one IMS subsystem fails while it holds locks on data in shared databases, the other sharing IMS subsystems must wait for the failed IMS to restart and release its locks. To reduce the amount of time that the sharing subsystems must wait, you can use IMS Fast Database Recovery (FDBR) regions.

FDBR regions monitor an IMS subsystem and can automatically recover database resources (shared databases and areas) if that IMS subsystem fails.

An FDBR region tracks a single IMS subsystem, so you must set up a separate FDBR region for each IMS subsystem in the sysplex that you want FDBR to track.

An FDBR region does not have to run in the same z/OS system as the IMS subsystem it tracks, but both the FDBR region and the IMS subsystem must be in the same z/OS cross-system coupling facility (XCF) group. The FDBR region also must have access to the following IMS data sets:

- ACBLIBx
- MODBLKSx
- MODSTATx
- OLDS
- Restart data set (RDS)
- RECONx
- SDFSRESLx
- WADS

FDBR is similar to XRF in that it operates in distinct phases:

- Surveillance
- Recovery process
- Post-recovery

Restrictions:

- You cannot use an FDBR region with the following types of subsystems: DCCTL, Extended Recovery Facility (XRF), or DBCTL standby.
- While using a tracking system such as FDBR, you must not use a buffer manager that is different from that of the active system.

You do not have to register a CSL with an FDBR region. If you do register a CSL with an FDBR region, it gives you the following benefits:

- a more complete picture of the IMSplex when looking at it from the QUERY IMSPLEX command
- a monitor program can watch the status of FDBR using the notifications that SCI sends when an IMSplex member joins and leaves the IMSplex

FDBR surveillance

During the surveillance phase, the FDBR region monitors database activity by reading the IMS log. It monitors all activity for shared databases and areas, and can monitor activity for nonshared areas.

An FDBR region also monitors RSR-related log activity at an active site.

Restriction: An FDBR region does not track MSDB activity. It also does not track databases at an RSR tracking site. If the IMS subsystem that the FDBR region is tracking fails, the FDBR region automatically begins recovery for all tracked databases and areas. If the IRLM for the tracked IMS subsystem fails, or if there is no log activity from the tracked IMS subsystem for a specified length of time, the FDBR region issues a message telling the operator to begin FDBR recovery.

FDBR recovery process

If the tracked IMS subsystem fails, the FDBR region automatically begins recovery. During automatic recovery, the FDBR region uses the MVS™ Availability Manager to determine when I/O is complete for the databases and areas it will recover.

During recovery, the FDBR region does the following for all tracked databases and areas:

- Allocates and opens database data sets
- Backs out updates for full-function databases
- Performs REDO processing for DEDB areas

After recovery is complete, the FDBR region tells IRLM to release retained locks for the recovered databases and areas. The FDBR region then issues a message that recovery is complete.

If the tracked IMS subsystem's z/OS cross-system coupling facility (XCF) fails, the FDBR region stops tracking and issues a message. If the FDBR region's XCF fails, the region abnormally terminates.

To begin recovery manually, issue the z/OS `MODIFY fdbbrproc, RECOVER` command. For manual recovery, the FDBR does not wait for I/O to complete because it assumes the I/O is already complete.

FDBR post-recovery

When recovery is complete, all shared databases and areas are available to other sharing IMS subsystems in the sysplex. You can restart the failed IMS subsystem at any time.

During IMS restart, IMS does the following to complete recovery:

- Writes updates for shared VSO areas to DASD
- Recovers nonshared DEDB areas that were not tracked by the FDBR region
- Recovers MSDBs

After the IMS subsystem is running, you can reestablish tracking by the FDBR region by issuing the z/OS `START fdbbrproc` command.

Commands for using FDBR

Use the commands for Fast Database Recovery (FDBR) to start and terminate tracking, obtain status, and initiate recovery. For each command, `fdbrproc` represents the FDBR procedure name.

Starting tracking:

Normally, you start FDBR with a JCL job, submitted during system start, after the tracked IMS is started.

You can also issue the z/OS START command to start FDBR tracking of an IMS subsystem, as follows:

```
S fdrproc
```

Obtaining status:

You use the z/OS MODIFY command to obtain Fast database recovery (FDBR) status.

```
F FDR1,STATUS
```

This command returns the following information:

```
DFS000I: PHASE:   TRACKING   LOG-TIME:  17:48:41   FDR1
DFS000I: ACT-ID:  SYS3      GROUPNAME: FDRSYS3    FDR1
DFS000I: TIMEOUT: 060 SEC   AREA01:   NORECOV    FDR1
DFS000I: SVS0OPEN: SERIAL   FPBUFF:    LOCAL     FDR1
```

The returned information includes:

PHASE

This field indicates the phase of the FDBR region. You can display the following phases:

- INIT (initialization phase)
- TRACKING (tracking phase)
- RECOVERY (recovery phase)

ACT-ID

The IMS ID that FDBR is tracking. If the FDBR region is in initialization phase, N/A is shown.

GROUPNAME

z/OS cross-system coupling facility group name used for XCF monitoring. If the FDBR region is in initialization phase, N/A is shown.

TIMEOUT

XCF timeout value. If the FDBR region is in initialization phase, N/A is shown.

LOG-TIME

Time associated with the log record currently being processed by FDBR region. The length of time the FDBR lags behind the IMS in reading the log is the difference between the current time, as shown by the time stamp and the log time. If the FDBR region is in initialization phase, N/A is shown.

AREA01=

This field shows the share level 0|1 DEDB area recovery option as follows:

NORECOV

NORECOV option is specified in DFSFDRxx IMS.PROCLIB member.

RECOV

RECOV option is specified in DFSFDRxx IMS.PROCLIB member.

Initiating recovery:

Use the z/OS MODIFY command to initiate FDBR.

You can use this command if FDBR determines a time out status during z/OS cross-system coupling facility (XCF) or log surveillance. The command is not accepted if FDBR is being initialized or is already in a recovery process.

The following command starts FDBR immediately: `F fdrproc,RECOVER`

Terminating tracking:

Use the z/OS MODIFY command to terminate FDBR.

The following command stops FDBR tracking activity:

`F fdrproc,TERM`

The command is not accepted if FDBR is being initialized or is in a recovery process.

Stopping the FDBR region:

You can use z/OS MODIFY commands to stop the FDBR region functions.

1. The following command ends FDBR functions without producing a dump:
 - a. `F fdrproc,STOP`
FDBR ends with a return code of 0200.
2. The following command ends FDBR functions and produces a dump:
 - a. `F fdrproc,DUMP`
FDBR ends with return code of 0020.

User access problems

The task of determining and responding to user access problems can be done by the IMS MTO, the network support group, or a user liaison group.

If your installation has a user liaison group that helps end users with problem determination, you need to explain what information that group should obtain from the end user before they contact the IMS MTO. This information includes a description of the problem's symptoms (such as poor response time or no response at all), details of any system or application program error messages received, terminal status lights, and so forth. Be sure to give the user liaison group guidelines on whether to call the network support group, application support group, or IMS MTO.

After they have gathered all information from the end user, they should develop a plan for determining the cause of the problem and, if possible, correcting it. Typically, user problems can be classified as:

- Poor system response
- No response, or terminal keyboard locked
- Unexpected output (system or application program error messages)

If the general system response is poor, the MTO should follow the procedures described in "Monitoring the system" on page 103.

If the end user is getting no response at all or has a locked keyboard, the problem might be in the network or in the status of the application program resources required by the transaction. For example, an OLDS might need to be archived to

free additional logging resources. Use one of the following commands to determine if any of the required resources are unavailable:

- /DISPLAY STATUS
- /DISPLAY OLDS
- QUERY DB SHOW(STATUS)

If the end user is receiving application program error messages or invalid output, the MTO should follow your procedures for contacting the appropriate application support group. You should define the documentation that MTO need to provide to the support groups, including a dump of the message region (if one was produced), or exact details of the original transaction the user entered. For IMS or VTAM system messages, the MTO should follow the appropriate error recovery procedures.

Related concepts:

“Monitoring the system” on page 103

Recovery of IMS using Remote Site Recovery (RSR)

Generally, Remote Site Recovery (RSR) does not change or interfere with normal IMS operations. Most RSR functions and operations are automatic. RSR operational activities at the active site are minimal. Operations at the tracking site are different from those at the active site.

Related concepts:

 [RSR overview \(System Administration\)](#)

Preparing for tracking

Global service groups (GSGs), service groups (SGs), and registered databases are defined in the RECON data set by using DBRC utility commands. You can define multiple GSGs in a single set of RECON data sets. (A set of RECON data sets consists of the two active copies and a spare copy).

For a particular GSG, a unique set of RECON data sets exists for each SG. That is, the active and tracking SGs have their own RECON data sets.

The RECON data sets of different SGs need not be identical, either physically or logically. The GSG, both of its SGs, and all of the covered databases must be defined in each set of RECON data sets. In each set of RECON data sets, one of the SGs is defined as the local SG.

The GSG and the SGs are defined by using the DBRC INIT.GSG and INIT.SG commands. The initial role (active or tracking) of an SG is specified when the SG is defined. IMS subsystems sign on to DBRC specifying the GSG; after a GSG has been defined, the RSR environment can be initialized. IMS subsystems send or receive log data depending on the role of the local SG, active or tracking.

After a GSG is defined, DL/I databases and DEDB areas for which coverage is required can be assigned to the GSG. The database or area must be defined in the tracking RECON data set and an initial image copy must be sent to and received at the tracking site.

Recommendations: Define each covered (tracked) database and area in the tracking RECON data set before the tracking site receives any log data for these covered databases and areas. RSR supports the IMS shared-queues environment for

active IMS subsystems. A tracking subsystem cannot share IMS message queues with other IMS subsystems, and abend if it is started with the shared queue options.

Define all IMS subsystems that participate in sharing IMS message queues as part of the same service group. If the active subsystems share the IMS message queues, they must continue to share them after a remote takeover. Therefore, you cannot change from a shared-queues environment to a nonshared-queues environment during a remote takeover.

Shadowing data to the tracking site

You can shadow the data to the tracking site for any given DL/I or DEDB area. You want to do this to backup your data, perhaps at another site.

For any given DL/I database or DEDB area, the following steps are required to begin shadowing data to the tracking site:

1. Prepare the tracking site.
 - Define the database or area in the tracking site's RECON data set
Use the DBRC INIT.DB and INIT.DBDS commands, specifying the GSG name, to define the database.
For a DL/I database, all data sets of the database must be defined in the tracking RECON using the same DD names that are used in the active RECON data set. For a DEDB, the areas are assigned to GSGs, rather than the DEDB.
Only the databases or areas covered by the GSG need to be defined in the tracking RECON data set. Non-covered databases or areas can be defined in the tracking RECON data set, but no log records are sent (if you use the default log filter exit routine, DFSFTFX0) to the tracking site for those databases or areas.
 - Allocate the tracking database or area data sets.
For database-level tracking, create the DBDSs or ADSs. The physical attributes (such as OSAM blocksize or VSAM CI size) of the tracking data sets should be the same as those of the active data sets, although the data set names can be different.
2. Start the active IMS subsystem.
Be sure to specify the GSG name and RSR(YES) in the DFSRSRxx member of IMS.PROCLIB.
3. Initiate coverage at the active site.
 - Assign the database or area to the GSG.
If the database is already defined in the RECON data set, specify the GSG name on the CHANGE.DB command for DL/I databases or the CHANGE.DBDS command for DEDB areas. The database must not be authorized to any IMS subsystem when you issue the CHANGE command.
If the database is not defined in the RECON data set, use the INIT.DB and INIT.DBDS commands, and specify the GSG name.
 - Create image copy data sets.
An image copy data set of each DBDS or covered area is required.
 - You can use any of the following: the IMS Database Image Copy utility (without the concurrent option), the Database Image Copy 2 (DFSUDMT0) utility (with the exclusive option), or some other method (for example, a pack dump).

- You cannot use a concurrent image copy, that is, a copy created by the Online Database Image Copy utility (with the concurrent option) or the Database Image Copy 2 (DFSUDMT0) utility (with the shared option).
 - Send the image copy data sets to the tracking site.
 - Begin updating the database at the active site.
Use the /START DB, UPDATE DB START(ACCESS), /START AREA, or UPDATE AREA START(ACCESS) command to allow the online IMS subsystems to use the database or area.
4. Start the tracking subsystem.
 5. Register the image copy data sets at the tracking site.
 6. Install the shadow database.
 - If you plan to use database-level tracking, the image copy data set must be applied to any DBTRACK-covered database or area. Use the database recovery utility to apply the image copy data sets to the shadow database data sets. Use the GENJCL.RECEIVE command to generate the database recovery utility job steps to apply the image copy data set to the shadow database data set.
 - For non-IMS image copy data sets, you must apply the image copy data set to the shadow database data set using the appropriate method (for example, pack restore). After you apply the shadow database data set, use the NOTIFY.RECOV command to inform DBRC that the image copy has been applied to the tracking database data set.
 7. Begin shadowing (database-level tracking only).
The simplest way to begin shadowing databases or areas is to use the /START DB or UPDATE DB START(ACCESS) command on the tracking subsystem. As part of the start command processing, all database change records created since the image copy was created are applied to the tracking database data sets. This process is called *online forward recovery*.
Initially, you can specify all shadow databases and areas on one /START or UPDATE command for the most efficient processing.

Related tasks:

“Initiating tracking of a DL/I database”

“Starting a new tracking subsystem” on page 190

Initiating tracking of a DL/I database

For remote-site recovery (RSR) tracing, you must initiate tracking of a DL/I database.

In this example, the database readiness level is DBTRACK. The DBD name of the sample database is IVPDB2. The database has a single DBDS with DD name DFSIVD2. The remote-site recovery environment has already been established: the GSG is defined, and active and tracking subsystems are communicating.

To initiate tracking of a database (or area):

1. Prepare the tracking site
2. Initiate coverage at the active site
3. Install image copy data sets at the tracking site
4. Begin database shadowing

Related tasks:

“Shadowing data to the tracking site” on page 153

Preparing the tracking site:

Sample JCL that shows a job stream that defines the database in the tracking RECON and allocates the shadow DBDS is shown.

```
//DEFTDB JOB ...
//DBRC EXEC PGM=DSPURX00,REGION=2048K
//SYSPRINT DD SYSPRINT=A
//IMS DD DSN=IMSIVP33.DBDLIB,DISP=SHR
//JCLOUT DD SYSOUT=(A,INTRDR)
//SYSIN DD *
INIT.DB DBD(IVPDB2) SHARELVL(3) TYPEIMS GSGNAME(MYGSG)
INIT.DBDS DBD(IVPDB2) DDN(DFSIVD2) -
          DSN(IMSIVP33.DFSIVD2) -
          NOREUSE RECOVPD(10) GENMAX(5)

/*
//ALLOC EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DEFINE CLUSTER(
          NAME(IMSIVP33.DFSIVD2)
          NONINDEXED
          FREESPACE(10 10)
          RECORDSIZE(2041 2041)
          SHAREOPTIONS(3 3)
          UNIQUE
          VOLUMES(IMS803)
          CYLINDERS(02)
          CONTROLINTERVALSIZE(2048)
        )
        DATA(
          NAME(IMSIVP33.DFSIVD2.DATA)
        )
/*
```

In the previous sample:

1. The first step calls the DBRC utility to execute the INIT.DB and INIT.DBDS commands to register the tracking database and its database data set. When this step completes, the databases are *covered* databases. In this example no predefined image copy data sets are used, so no INIT.IC commands are needed.
2. The second step calls IDCAMS to allocate the database data set.

Initiating coverage at the active site:

The database must not be authorized to any active or batch subsystems when coverage is initiated. Use the /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command to unauthorize the database from any online IMS subsystems. You can use the DBRC LIST.DB command or the DBRC QUERY API request to determine if any subsystem has authorized the database.

Initiate database coverage by:

1. Assigning the database to the GSG
2. Making an image copy of the DBDS
3. Sending the image copy data set to the tracking site

The following sample JCL shows a job stream that initiates coverage at the active site. This job calls the DBRC utility to assign the database to the GSG

(CHANGE.DB command) and to generate a second job stream (GENJCL.IC) to create both an image copy of the DBDS and the commands required to register the image copy at the tracking site.

```
//COVERDB JOB ...
//DBRC01 EXEC PGM=DSPURX00,REGION=2048K
//SYSPRINT DD SYSOUT=A
//IMS DD DSN=IMSIVP33.DBDLIB,DISP=SHR
//JCLPDS DD DSN=IMSIVP33.JCLPDS,DISP=SHR
//JCLOUT DD SYSOUT=(A,INTRDR)
//SYSIN DD *
        CHANGE.DB DBD(IVPDB2) GSGNAME(MYGSG)
        GENJCL.IC DBD(IVPDB2) MEMBER(TRKSETUP)
/*
```

The GENJCL.IC command uses skeletal JCL member TRKSETUP as shown in the following sample image copy skeletal JCL. The tracking setup member begins with an image copy utility skeletal JCL step. For the IC step, this example uses predefined cataloged image copy data sets. The second step uses the DBRC GENJCL.USER command to generate a NOTIFY.IC command, which will be used at the tracking site to register the image copy data set.

Recommendation: Use a command to receive the image copy data set.

```
//IC%STPNO EXEC PGM=DFSUDMP0,REGION=800K%PARMCIC
/*
/*          JCL FOR IMAGE COPY.
/*
//SYSPRINT DD SYSOUT=A
%SELECT DBDS((%DBNAME,%DBDDN))
%DELETE (%DBADSAV NE 'AVAIL')
//%DBADDN DD DSN=%DBDSN,DISP=OLD
%ENDDEL
%DELETE (%DBADSAV NE '')
//%DBDDN DD DSN=%DBDSN,DISP=OLD
%ENDDEL
%ENDSEL
//DATAOUT1 DD DSN=%ICDSN1,DISP=OLD
//SYSIN DD *
%ICSYSIN
/*
//***** END OF IMAGE COPY SKELETAL JCL MEMBER
//GENIC EXEC PGM=DSPURX00,REGION=2048K
/*
/*          JCL FOR GENJCL.USER TO GENERATE A NOTIFY.IC COMMAND
/*
//SYSPRINT DD SYSOUT=A
//JCLPDS DD DSN=IMSIVP33.JCLPDS,DISP=SHR
//JCLOUT DD DSN=TSOID.U.CNTL(REMOTEIC),DISP=OLD
//SYSIN DD *
        GENJCL.USER MEMBER(NTFYIC) -
        USERKEYS((%DBD,'%DBNAME'),(%DDN,'%DBDDN'))
/*
```

The following sample shows the skeletal JCL used to build the NOTIFY.IC and GENJCL.RECEIVE commands with the REMOTEIC member.

```
//NTFYIC EXEC PGM=DSPURX00,REGION=2048K
//SYSPRINT DD SYSOUT=A
//JCLPDS DD DSN=IMSIVP33.JCLPDS,DISP=SHR
//JCLOUT DD SYSOUT=(A,INTRDR)
//SYSIN DD *
%SELECT IC((%DBD,%DDN),LAST)
        NOTIFY.IC DBD(%DBD) DDN(%DDN) -
        RUNTIME(%ICTIME) -
%DELETE (%ICTYPE EQ 'SMSNOCIC' | %ICTYPE EQ 'BATCH')
```

```

                                STOPTIME(%ICSTOP) -
%ENDDDEL
                                USID(%ICUSID) -
                                ICDSN(%ICDSN) -
                                VOLLIST(%ICVOLS) -
                                %ICTYPE
                                GENJCL.RECEIVE DBD(%DBD) DDN(%DDN)
%ENDSEL
/*

```

Finally, you need to transmit the image copy data set to the tracking site using a file transfer program.

Installing image copy data sets at the tracking site:

When the image copy data set arrives at the tracking site, execute the NOTIFY.IC and GENJCL.RECEIVE job stream that was created at the active site. The resulting job generated by the GENJCL.RECEIVE command installs the image copy data set.

Beginning database shadowing:

Use the /START DB or UPDATE DB START(ACCESS) command at the active site to start the database and begin normal update activity.

RSR in an IMSplex

You can use RSR in an IMSplex. An RSR tracking system can connect to a Common Service Layer (CSL) to enable automatic RECON loss notification or to receive commands from the Operations Manager (OM).

If an RSR tracking system does not use the OM or automatic RECON loss notification, it does not require an IMSplex, so it does not require a CSL. In this case, the RSR tracking IMS is not able to receive commands from the OM, but can continue to use existing methods for command entry. In the event of a remote takeover, if the RSR tracking IMS is required to take over the workload of an IMSplex, a CSL must be available at the tracking site.

If an RSR tracking IMS connects with the CSL, it registers for IMS commands with the OM using its IMSID. The IMSplex that the RSR tracking IMS joins must be a different IMSplex from the one at the active site.

You cannot use RSR to recover an entire IMSplex, but you can use RSR to recover an individual IMS in an IMSplex.

Related concepts:

 Automatic RECON loss notification overview (System Administration)

Tracking subsystem and the RSR complex

After the Remote Site Recovery (RSR) environment (at both sites) has been initialized, tracking begins. The tracking subsystem uses DBRC to record information about the log data that it receives.

As log records are written by an active subsystem, they are also sent to the tracking subsystem, where they are written to SLDSs. These log data sets (tracking SLDSs) are recorded in the RECON data set as they are created. In general, the tracking subsystem receives multiple log streams from the active site, one stream for each active IMS subsystem. Likewise, the tracking subsystem produces multiple output SLDSs, one set for each log stream received.

HALDB Online Reorganization (OLR) cannot be initiated on an RSR tracking system. However, if you run OLR at the active site, the tracking system can track the updates involved in the OLR as it does for other databases.

Related tasks:

 Monitoring HALDB Online Reorganization (Database Administration)

Tracking subsystem log

In addition to the tracking SLDSs, the tracking subsystem has its own log. Like any online IMS subsystem, the tracking subsystem creates OLDSS and SLDSs. All of these logs are maintained by DBRC.

The tracking subsystem does not create its own RLDSS because the tracking subsystem does not update databases on its own, but only shadows those database updates done at the active site.

The tracking subsystem's OLDSS are recorded in PRIOLDS and SECOLDS log records, just as is done for active IMS subsystems. Because there are no RLDSS, there are no PRILOG or SECLOG records for the tracking subsystem.

The SLDSs created when the tracking subsystem's OLDSS are archived are recorded in tracking SLDS records (PRITSLDS and SECTSLDS). With the exception of these additional record types, tracking SLDS records have exactly the same format and content as the SLDS records (PRISLDS and SECSLDS) used at the active site. The PRITSLDS and SECTSLDS record types prevent confusion with the PRISLDS and SECSLDS records created for the active subsystem.

Managing the databases in an RSR environment

When an active subsystem requests update authorization for a database, DBRC determines if the database is to be covered by RSR (tracked). If it is, the requesting subsystem must belong to the GSG to which the database belongs. DBRC rejects the authorization request if a covered database belongs to a different GSG.

A database-level tracking (DLT) subsystem requires exclusive authorization to the shadow databases. If the database is not currently authorized, DBRC grants authorization to the tracking subsystem in exclusive mode. DBRC rejects the authorization request if the database is currently authorized (for example, to a database utility).

The tracking subsystem receives the database update log records from the active subsystem. For an RLT tracking subsystem, the log records are simply written to the tracking SLDS. For a DLT tracking subsystem, the update records are not only written to the tracking SLDS, but they are also applied to the shadow databases.

For database-level tracking, DBRC knows the current state of the shadow DBDSs, that is, DBRC knows which log records have actually been applied to the DBDSs. For example, if a remote takeover occurs while a tracking DBDS is offline, that DBDS is unusable (marked recovery needed) until all log records received for that particular DBDS while it was offline are applied.

Database update set

A database update set is the collection of all updates that are made to a database while the database is continuously authorized to one or more IMS subsystems. DBRC uses a database update set to control the tracking of the databases.

A database update set begins with the first update to the database after a period with no update activity for the database. This update can be the first update after starting or restarting the active subsystem or it can be the first update after bringing the database online after it has been offline.

A database update set ends when all IMS subsystems stop updating the database. For example, a `/DBRECOVERY DB GLOBAL` or `UPDATE DB STOP(ACCESS)` command (routed to all IMS subsystems in the IMSplex) ends a database update set.

Each database update set is assigned a unique update set identifier (USID). DBRC records the USID whenever a DBDS or area is allocated, and stores the USID in the database record in the RECON data set. The allocation records also contain the current USID.

DBDS and area allocation

The existence of an allocation record in the tracking RECON data set indicates that the tracking subsystem has received database change records for a DBDS or area (or that they may be on the way in the case of block-level data sharing). DBRC writes the allocation record to the tracking RECON data set before the tracking subsystem actually applies the updates to the shadow DBDSs.

When the tracking subsystem receives a begin-update or end-update log record for a particular DBDS, the tracking subsystem notifies DBRC. Depending on the type of record, DBRC:

- Writes an allocation record in the RECON data set indicating the start of update activity for the DBDS by the active subsystem that wrote the begin-update record.
- Updates the allocation record to indicate the end of DBDS updates. If the active subsystem later begins updating the DBDS again, DBRC creates a new allocation record.

The tracking subsystem informs DBRC that updates are complete when either of the following occur:

- The tracking subsystem receives an end-update log record
- The active subsystem terminates, and the tracking subsystem receives the end of the log and applies all database change records

If for some reason IMS is unable to apply all of the database change records, the tracking subsystem informs DBRC of the last record applied. DBRC saves this position in the allocation record, along with an indicator that updates are incomplete. Thus, DBRC ensures that the DBDS is not used while updates are pending. The tracking subsystem can query DBRC to determine when to begin applying updates for online forward recovery after database tracking resumes.

When DBRC is notified of an image copy, DBRC deletes extraneous allocation (ALLOC) records and updates active ALLOC records to the time of the first log volume necessary for recovery, based on the oldest image copy for the DBDS or area. It is possible on an RSR tracking subsystem for an ALLOC record to be deleted or updated before DBRC applies the changes, for example, if the database was never started or tracking was stopped (suspended). In this case, you need to recover the DBDS or area before you can restart tracking.

IMS Fast Database Recovery (FDBR) regions are not used at the tracking site. FDBR regions can track and recover databases and areas at active sites, but not at the tracking site.

Normal operations for an RSR active site

RSR site operations generally include working with the Transport Manager Subsystem (TMS), the IMS logger, and the High-Speed Sequential Processing (HSSP) support for RSR. Additionally, you might work with online change, data sharing, and DL/I batch operations.

Transport Manager Subsystem

You can use five commands to control the operation of the Transport Manager Subsystem (TMS): DEFINE, DISPLAY, SET, START, and STOP.

The TMS commands DISPLAY, START, and STOP, are used in normal TMS operations. You issue these commands from a SYSIN data set while starting the component or from a z/OS console, by using the z/OS MODIFY command.

TMS DISPLAY command

The TMS DISPLAY command provides information about the definition and status of Transport Manager resources, such as the GSG name, SG name, and system name. If you need information about IMS resources at the active or tracking site, use the IMS /DISPLAY or QUERY command.

The Transport Manager Subsystem provides only communication services for the IMS active and tracking subsystems, so it is unaware of either site's activities.

TMS START command

The TMS START command activates the Transport Manager Subsystem (TMS), isolated log sender (ILS) process, or the system, depending upon which keywords you use with the command:

- START TMS starts the TMS.
- START ILS starts an ILS task.
- START SYSTEM initiates conversations with the specified TMS subsystem.

TMS STOP command

The TMS STOP command stops and deactivates the system, conversation allocation retry, or ILS, depending upon which keywords you use with the command:

- STOP SYSTEM terminates conversations with the specified subsystem.
- STOP RETRY stops automatic retry of conversation allocations for a specified TMS subsystem.
- STOP ILS terminates an ILS process.

Related reference:

 [IMS Transport Manager Subsystem commands \(Commands\)](#)

IMS logger

The Transport Manager Subsystem (TMS) receives the contents of OLDS buffers each time I/O for the OLDS starts. The TMS sends the contents of the OLDS buffer to the log router at the tracking site.

If VTAM becomes overloaded (indicated by message DFS4020I), the logger suspends sending of OLDS buffers until the next OLDS switch. If other errors occur while sending records, the logger terminates the conversation.

If there is no current conversation and an OLDS is full, the logger tries to re-establish a conversation with the log router; otherwise, the logger resumes a suspended conversation.

Use the IMS /STOP SERVGRP command if you do not want the logger to automatically attempt to establish a conversation for an OLDS switch. Use the /START SERVGRP command to attempt a conversation connection before the next OLDS switch.

If a connection to the tracking site is lost, DL/I batch at the active site does not try to re-establish the connection. The tracking subsystem detects the missing data and uses ILS to obtain it.

The Log Filter exit routine (DFSFTFX0) enables you to control the amount of log data that is sent to the tracking subsystem. IMS supplies a default routine, which eliminates diagnostic data and block-padding data.

Online change at the RSR active subsystem: The tracking subsystem does not automatically shadow online changes made to active site subsystems. The tracking subsystem does receive the online change log records (unless you filter them out using the Log Filter exit routine).

If the RSR active subsystem is restarted at the remote site and the RSR active subsystem is to continue to use global online change at the remote site, the RSR active subsystem needs to have its own OLCSTAT data set in its own IMSplex.

If the RSR active subsystem is restarted at the remote site and the RSR active subsystem is to use local online change at the remote site, the RSR active subsystem needs to use MODSTAT data sets instead of the OLCSTAT data set.

Recommendation: Before making online changes to an active IMS subsystem, copy the changes to the inactive library at the tracking site. This ensures that a copy of the changes is available if a remote takeover occurs after an online change at the active site.

If you have a lot of online changes to make in a short period of time, be sure the online change log records are received and processed at the tracking site before committing the changes at the active site. The tracking system issues a message for each committed online change.

If an online change affects the physical structure of a database, such as changing the logical record length or block size, you should copy the changes to the tracking subsystem staging library before making online changes and restarting the database at the active site. Afterwards, perform online changes to the tracking subsystem.

Because the active site can change roles and become a tracking site after a remote takeover, you should maintain, at the active site, tracking libraries that reflect online change activity.

Data sharing

In a data-sharing environment, IMS maintains database integrity for shared, covered databases at the tracking site. Active subsystems perform the necessary activities for support of tracking databases participating in block-level data sharing.

DL/I batch at the RSR active site


DL/I batch jobs at the active site can update databases that are tracked at the tracking site. All of the updates made by the batch job must be transported to the tracking site to maintain data integrity.

Restriction: For RSR to transport DL/I batch log data to the tracking site, the batch job must use both logging and DBRC.

You can use the Log Filter exit routine (DFSFTFX0) to filter batch job logs.

If a connection to the tracking site is lost, DL/I batch at the active site does not try to re-establish the connection. The tracking subsystem detects the missing data and uses ILS to obtain that data.

Related reference:

 Log Filter exit routine (DFSFTFX0) (Exit Routines)

High-Speed Sequential Processing (HSSP) support for RSR

In a non-HSSP environment, IMS logs DEDB updates in the X'5950' log record. In an HSSP environment, IMS handles DEDB updates (not including backward references and sequential dependent updates) differently.

The different ways that IMS handles DEDB updates include:

- The updates are stored in an image copy data set.
- IMS writes a X'5947' log record, which contains a bit map corresponding to the updated VSAM control intervals (CIs).

The actual updated data is in the image copy data set and not on the log. For recovery, IMS uses the image copy data set as a log extension. When the X'5947' log record is processed, IMS uses the bit map to get the corresponding data from the image copy data set.

For HSSP with RSR, IMS logs both the X'5947' and X'5950' log records because the tracking site uses the X'5950' records for recovery. When recovery is required at the active, IMS uses the image copy.

In a non-RSR environment, you can choose whether or not to log both records by using a set options (SET0) control statement keyword, LOGALL, in the DFSCTL data set. IMS ignores this keyword in an RSR environment.

Online change at the RSR active subsystem

The tracking subsystem does not automatically shadow online changes made to active site subsystems. The tracking subsystem does receive the online change log records (unless you filter them out using the Log Filter exit routine).

If the RSR active subsystem is restarted at the remote site and the RSR active subsystem is to continue to use global online change at the remote site, the RSR active subsystem needs to have its own OLCSTAT data set in its own IMSplex.

If the RSR active subsystem is restarted at the remote site and the RSR active subsystem is to use local online change at the remote site, the RSR active subsystem needs to use MODSTAT data sets instead of the OLCSTAT data set.

Recommendations:

- Before making online changes to an active IMS subsystem, copy the changes to the inactive library at the tracking site.
This ensures that a copy of the changes is available if a remote takeover occurs after an online change at the active site.
- If you have a lot of online changes to make in a short period of time, be sure the online change log records are received and processed at the tracking site before committing the changes at the active site.
The tracking system issues a message for each committed online change.
- If an online change affects the physical structure of a database, such as changing the logical record length or block size, you should copy the changes to the tracking-subsystem staging library before making online changes and restarting the database at the active site.
Afterwards, perform online changes to the tracking subsystem.
- Because the active site can change roles and become a tracking site after a remote takeover, you should maintain, at the active site, tracking libraries that reflect online change activity.

Data Sharing

In a data-sharing environment, IMS maintains database integrity for shared, covered databases at the tracking site. Active subsystems perform the necessary activities for support of tracking databases participating in block-level data sharing.

DL/I batch at the RSR active site

DL/I batch jobs at the active site can update databases that are tracked at the tracking site. All of the updates made by the batch job must be transported to the tracking site to maintain data integrity.

Restriction: In order for RSR to transport DL/I batch log data to the tracking site, the batch job must use both logging and DBRC.

You can use the Log Filter exit routine (DFSFTFX0) to filter batch job logs.

If a connection to the tracking site is lost, DL/I batch at the active site does not try to reestablish the connection. detects the missing data and uses ILS to obtain that data.

Related concepts:

“Online change at the RSR tracking subsystem” on page 171

Data sharing

In a data-sharing environment, IMS maintains database integrity for shared, covered databases at the tracking site. Active subsystems perform the necessary activities for support of tracking databases participating in block-level data sharing.

DL/I batch at the RSR active site

DL/I batch jobs at the active site can update databases that are tracked at the tracking site. All of the updates made by the batch job must be transported to the tracking site to maintain data integrity.

Restriction: For RSR to transport DL/I batch log data to the tracking site, the batch job must use both logging and DBRC.

You can use the Log Filter exit routine (DFSFTFX0) to filter batch job logs.

If a connection to the tracking site is lost, DL/I batch at the active site does not try to re-establish the connection. The tracking subsystem detects the missing data and uses ILS to obtain that data.

Normal operations for an RSR tracking site

Operations at an RSR tracking site include log router tracking, database tracking (including HALDB), online change, DL/I batch, and IRLM.

Log router

The log router component of the tracking subsystem performs various tasks that manages the routing process.

During normal operations, the log router component of the tracking IMS subsystem does the following tasks:

- Accepts and starts conversations
- Receives and stores active IMS log data
- Supports data sharing
- Sends data to database trackers
- Archives SLDSs
- Initiates and participates in milestone processing
- Truncates logs

Accepting and starting conversations:

The log router, or an active-site component such as the logger or ILS, initiates conversation allocation.

As part of its initialization, each component tries to allocate a conversation with its “peer” (equivalent component performing the same or similar function) at the other site. If the peer is not yet initialized, the allocation fails, and the component waits for its peer to allocate the conversation at the other site or for a /START SERVGRP command to initiate the conversation.

Receiving and storing active IMS log data:

The IMS logger uses a Transport Manager Subsystem (TMS) conversation to pass active IMS log data to the log router.

During processing of the first buffer of log data, the log router opens a pre-allocated SLDS and registers it with DBRC. Thus, one pre-allocated SLDS always exists for each log stream at the tracking site. The log router writes the buffer's contents into the SLDS. When the first SLDS is full, the log router closes it and notifies DBRC. The log router then opens a new SLDS, tells DBRC about it, and writes log records into the new SLDS.

At any one time, the output from a DBRC LIST.RECON command might show multiple tracking SLDSs with no start and stop times; this is normal for the tracking subsystem. All tracking logs have start and stop times when appropriate. IMS ensures that when a remote takeover occurs, all tracking logs have start and stop times.

If you request automatic archiving by specifying the ARCHIVE keyword in the DFSRSRxx member, the closed SLDs are automatically archived.

Some types of logged information cause the log router to send requests to DBRC to ensure that the state of the databases is maintained in the RECON data set. This action occurs whenever the log router detects one of the following events in the log data received from the active subsystem:

- Begin update
- End update
- Batch backout
- Time stamp forward recovery

Related concepts:

“Automatic archiving of the tracking SLDs” on page 166

PRILOG record compression:

The tracking subsystem receives and manages logs received from active subsystems. The tracking subsystem compresses the PRILOG records in these logs, just as an active subsystem does.

However, this compression is independent of the compression done at active sites: an active subsystem's PRILOG record could fill up and need compression before the tracking subsystem's PRILOG record, or the tracking subsystem's PRILOG record could fill up first.

It is important that the tracking subsystem receive current log data, to ensure that data is not compressed by the active subsystem before the tracking subsystem receives it. You can adjust the log retention period for the active subsystem's RECON data set to reduce the risk of this happening.

Recommendation: Do not record image copies at the tracking site so that the time of the oldest image copy is later than the current routing time. Use the /DISPLAY TRACKING STATUS command to determine the current routing time (the ROUTED-LOG time).

Related concepts:

 Deleting unnecessary RECON records (System Administration)

Database level tracking:

For a database level tracking (DLT) subsystem, additional processing occurs after the log data is stored at the tracking site. The log router passes database change records to the individual components responsible for database tracking. These individual components apply the changes to those databases registered as DBTRACK.

Related concepts:

“Database tracking” on page 167

Block-level data sharing:

The log router merges database update records in creation-time sequence. If log records are missing (that is, a gap exists) from one or more active subsystems for any shared database, the log router continues presenting records to the database tracking components, up to the beginning of the gap in the log.

The log router stops all routing of the log data at the gap point, until the gap is filled. The tracking system continues accepting log data from the active site. After the gap in the log has been filled, routing resumes.

Automatic archiving of the tracking SLDSs:

You can request automatic archiving of the tracking system log data sets (SLDSs) by specifying the ARCHIVE keyword, with various parameters, in the DFSRSRxx IMS.PROCLIB member. You can use the GENJCL.ARCHIVE command to archive tracking SLDSs manually when you do not use automatic archiving. After a successful archive, IMS deletes the original SLDS. The archived copy is the only one available for restart, online forward recovery, or catch-up processing.

If you do not specify ARCHIVE in the DFSRSRxx member, the log router does not perform automatic archiving of the SLDSs and you must use offline utilities to archive these data sets.

When you specify ARCHIVE, you can optionally request that IMS create RLDSs by also specifying RLDSDEF in the DFSRSRxx IMS.PROCLIB member.

Recommendation: Archive SLDSs to a storage medium that is either slower or less expensive than the medium used for the SLDS itself. If you have space limitations for your non-archived data, do not delay automatic archiving.

Use the following keywords on the ARCHIVE statement in the DFSRSRxx IMS.PROCLIB member to control automatic archiving:

MINSLDS

Specifies the number of SLDSs to be archived during any one automatic archive operation for a particular active subsystem being tracked.

When the number of tracked SLDSs not needed to restart an active subsystem after a remote takeover reaches the MINSLDS value, automatic archiving begins. The MINSLDS value specifies the number of tracking SLDS to be copied into a single archive data set.

MAXSLDS

Specifies how many SLDSs are to be maintained before an automatic archive is scheduled.

When the log router has a number of SLDSs equal to or greater than the MAXSLDS value for a particular active subsystem being tracked, it starts an automatic archive. This archive is started even if the SLDSs contain data that is needed for restart of the new active subsystem after remote takeover.

MAXPROCS

Specifies how many automatic archive operations can run concurrently.

If the number of archive operations exceeds the MAXPROCS value, some of the archive operations are queued.

All batch logs (SLDSs) become eligible for automatic archive as soon as they are closed.

Examples:

1. You have nine SLDSs not needed for restart for a particular active subsystem, and you specify:
ARCHIVE(MINSLDS(3) MAXSLDS(8) MAXPROCS(2))

Two archive operations are started:

- Three SLDSs are archived because the MINSLDS value is three.
- Two jobs are started because the MAXPROCS value is two. The other job is enqueued.
- Automatic archiving begins because you have exceeded the threshold set by the MAXSLDS value.

After the SLDSs are archived, you have six SLDSs.

2. You have 16 SLDSs not needed for restart for a particular active subsystem, and you specify:

```
ARCHIVE(MINSLDS(6) MAXSLDS(15) MAXPROCS(3))
```

Three archive operations are started:

- Six SLDSs are archived because the MINSLDS value is six.
- Three jobs are started because the MAXPROCS value is three. The other three jobs are enqueued.
- Automatic archiving begins because you have exceeded the threshold set by the MAXSLDS value.

After the SLDSs are archived, you have 10 SLDSs.

Related concepts:

“Receiving and storing active IMS log data” on page 164

Related reference:

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

Milestone processing:

The tracking subsystem records tracking status information on its own log (OLDS) during operation to provide restart capability. The collecting and writing of this information is called *milestone processing* and is performed at user-specified intervals between checkpoints and during tracking subsystem shutdown.

The log records written during a milestone are used during tracking subsystem restart to determine the correct point on each log stream from which to resume tracking.

The database trackers participate in milestone processing. It is during milestone processing that database updates are actually written to the databases. You specify the milestone frequency using the MILEINTV keyword in the DFSRSRxx member of IMS.PROCLIB.

When IMS log routing is idle (that is, no new records are being routed) for five consecutive milestones, IMS ignores further milestone requests until routing activity resumes.

Database tracking

When the database tracker receives the first change record for a shadow database or area registered as DBTRACK, it must acquire exclusive authorization for and open the database or area.

After DBRC grants authorization for the database or area and the database or area has been opened, the database tracker applies log records from the active subsystems to the shadow database or area.

After DBRC grants authorization to the database tracker, no other IMS subsystem, including all database utilities, can access the database or area until the tracker relinquishes its authorization. The database or area remains open and authorized until an error occurs or until you issue a /DBRECOVERY or UPDATE command for the database or area.

The log router periodically calls the database trackers to perform milestone processing. During milestone processing, buffers that contain database changes are written to the shadow databases or areas. A /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command for a DL/I database does not complete until two milestones have occurred.

Online forward recovery (OFR) allows the database trackers to bring databases, which had been unavailable, back to the current tracking state. You can initiate online forward recovery by issuing one of the following commands:

- /START DB
- /START DATAGRP
- UPDATE DB START(ACCESS)
- UPDATE DATAGRP START(ACCESS)

OFR is also used when the tracking subsystem receives an initial or a more current image copy that needs updates applied to it to bring the database up to date.

Databases and areas must have a STOPPED status before the tracking subsystem accepts a /START or UPDATE command. The /DBRECOVERY or UPDATE command might be required in specific situations to give the database or area the STOPPED status. One such situation is if an RLT tracking subsystem is switched to database-level tracking (DLT) with databases and areas registered as DBTRACK. In this situation, you need to issue a /DBRECOVERY or UPDATE command for the databases and areas before starting them for online forward recovery processing.

When the tracking subsystem tracks several active subsystems, because updates from one active subsystem can arrive before updates from another, the log router can pass log data to the database trackers in a different order than which it was created. In order to maintain database integrity in a data-sharing environment, log router merges and routes the data in creation-time sequence. However, log router might concurrently route log data created by consecutively-executing active subsystems or batch jobs.

As a result, a database tracker can receive updates from one update set before receiving updates for a previous update set. In this case, the database tracker does not apply the updates for any of the update sets. Instead, the tracking subsystem stops the database or area and issues message DFS4045A ("OFR NEEDED" reason) or DFS2960A (reason code 10). In response, issue the /START DB, UPDATE DB START(ACCESS), /START AREA, or UPDATE AREA START(ACCESS) command for the affected database or area in order to allow online forward recovery to apply these updates.

Related concepts:

"Database level tracking" on page 165

Database commands for HALDBs at the tracking subsystem

When you issue a database command, the database name can be a shadow non-HALDB, a shadow HALDB partition, or a HALDB master. A command that specifies a HALDB partition operates like a command for a non-HALDB in that the states of the HALDB partition include STOPPED, NOTOPEN, or ALLOCS.

The following IMS commands for a High Availability Large Database (HALDB) at the tracking site have special considerations that do not apply for a non-HALDB:

- /DBRECOVERY DB
- /DISPLAY DB
- /START DB
- QUERY DB
- UPDATE DB START(ACCESS)

When you issue a /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command for a HALDB partition, the partition is affected as follows:

- Tracking of changes in the partition is suspended.
- The partition assumes the stopped state, as when you issue the /DISPLAY DB or QUERY DB command.

For the HALDB partition, the tracking-suspended state and the stopped state occur together, even though the two states are slightly different in concept.

When you issue a /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command for a HALDB master, the master and all of its partitions are affected as follows:

- Tracking is suspended for all of the partitions, but the partitions do not assume the stopped state because of the /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command (though they might be in the stopped state for some other reason).
- The master assumes the stopped state, as when you issue the /DISPLAY DB or QUERY DB command.

To start Online Forward Recovery (OFR) for the HALDB partitions after the master has been stopped, issue a /START DB or UPDATE DB START(ACCESS) command for the master. These commands start OFR for only the partitions that were not stopped individually. Note that a partition could be stopped individually by either a /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command or an error. To start OFR for such a partition, enter a /START DB or UPDATE DB START(ACCESS) command for the partition. The rule is that tracking occurs for a HALDB partition only when neither the partition nor its master is in the stopped state.

And note that because database data is stored only in the HALDB partitions and not in the master, tracking occurs only for the partitions.

Effects of issuing /DBR, /STA, or UPD commands on HALDB masters or partitions:

The results, based on different initial states, of issuing a /DBRECOVERY, /START, or UPDATE command for a HALDB master or a HALDB partition are listed.

In the situations illustrated in these tables, tracking is already suspended for partitions shown as STOPPED or NOTOPEN before the command is issued. When the master's initial state is shown as STOPPED, tracking is already suspended for all of its partitions.

Table 21. Results of issuing /DBR or UPD command for HALDB master.

Conditions before command		Conditions after command	
Master	Partition	Master	Partition

Table 21. Results of issuing /DBR or UPD command for HALDB master (continued).

Conditions before command		Conditions after command	
Not STOPPED	ALLOCS	STOPPED	NOTOPEN, but not STOPPED Tracking suspended
Not STOPPED	STOPPED, NOTOPEN	STOPPED	STOPPED, NOTOPEN (no change)
STOPPED	NOTOPEN, but not STOPPED	STOPPED (no change)	NOTOPEN, but not STOPPED (no change)
STOPPED	STOPPED, NOTOPEN	STOPPED (no change)	STOPPED, NOTOPEN (no change)

Table 22. Results of issuing /DBR or UPD command for HALDB partition.

Conditions before command		Conditions after command	
Master	Partition	Master	Partition
Not STOPPED	ALLOCS	Not STOPPED (no change)	STOPPED, NOTOPEN Tracking suspended
Not STOPPED	STOPPED, NOTOPEN	Not STOPPED (no change)	STOPPED, NOTOPEN (no change)
STOPPED	NOTOPEN, but not STOPPED	STOPPED (no change)	STOPPED, NOTOPEN
STOPPED	STOPPED, NOTOPEN	STOPPED (no change)	STOPPED, NOTOPEN (no change)

In the following table, the /START or UPDATE command is rejected because these commands are attempting to start a HALDB master that is not in the stopped state. When this occurs, message DFS4108A (reason code 52) is issued. No change is made in the condition of the master or its partitions.

Table 23. Results of issuing /STA or UPD command for HALDB master.

Conditions before command		Conditions after command	
Master	Partition	Master	Partition
Not STOPPED Command rejected	ALLOCS	Not STOPPED (no change)	ALLOCS (no change)
Not STOPPED Command rejected	STOPPED, NOTOPEN	Not STOPPED (no change)	STOPPED, NOTOPEN (no change)
STOPPED	NOTOPEN, but not STOPPED	Not STOPPED	ALLOCS OFR started.
STOPPED	STOPPED, NOTOPEN	Not STOPPED	STOPPED, NOTOPEN (no change)

In the following table, the /START or UPDATE command is rejected because these commands are attempting to start a HALDB partition that is not in the stopped state. When this occurs, message DFS4108A (reason code 52) is issued. No change is made in the condition of the partition or its master. Also, if the /START or UPDATE command attempts to start a HALDB partition whose master is in a stopped state after the command is issued, OFR is not done. When this occurs, message DFS4108A (reason code 53) is issued. The partition's stopped state is reset, but OFR is not started because of the master's stopped state.

Table 24. Results of issuing /STA or UPD command for HALDB partition.

Conditions before command		Conditions after command	
Master	Partition	Master	Partition
Not STOPPED	ALLOCS Command rejected.	Not STOPPED (no change)	ALLOCS (no change)
Not STOPPED	STOPPED, NOTOPEN	Not STOPPED (no change)	ALLOCS OFR started
STOPPED	NOTOPEN, but Not STOPPED Command rejected.	STOPPED (no change)	NOTOPEN, but not STOPPED (no change)
STOPPED	STOPPED, NOTOPEN	STOPPED (no change)	NOTOPEN, but not STOPPED OFR not done

Online change at the RSR tracking subsystem

Perform online change to the tracking subsystem using the normal /DISPLAY MODIFY or /MODIFY PREPARE and /MODIFY ABORT | COMMIT command sequence.

Usually, the tracking subsystem is not concerned with database definition changes that do not require taking the database offline at the active site. This means, for example, that adding segment types does not require a matching change on the tracking subsystem for correct database tracking.

The tracking system issues a message for each committed online change. However, processing for the online change can occur some time after you receive the message. The tracking subsystem ensures that online change information is not acted upon until the tracking system can be sure that the online change log information will not be lost due to log truncation.

Because the tracking-subsystem libraries can be supersets of, or subsets of, the active-subsystem libraries, you might have to copy selected changed blocks from the active- to the tracking-subsystem libraries, and run the Application Control Blocks Maintenance utility to produce a proper staging library. If you maintain a library at the active site that will become the tracking library when the active site becomes a tracking site, be sure to copy that library to the current tracking site.

Related concepts:

“Online change at the RSR active subsystem” on page 162

DL/I batch at the RSR tracking site

The tracking site receives batch log data and writes it to SLDSs. The tracking subsystem informs DBRC of the log data set and the identity of any covered databases for which database-update records appear on the log.

For a DLT tracking subsystem, the log records are also passed to the individual database trackers.

IMS handles batch logs and online logs in the same ways. When a remote takeover occurs, you must back out any uncommitted changes using the Batch Backout utility (DFSBB000) because remote takeover processing does not dynamically back out uncommitted changes from batch logs.

RSR termination at the active site

You can terminate the Transport Manager Subsystem (TMS) without terminating IMS or conversations associated with IMS. To terminate the TMS, use the z/OS STOP command.

Terminating the TMS has the following additional effects:

- The DBRC instance in the TMS address space stops.
- The ILS stops.
- The TMS informs components identified to it that it is stopping.

Terminate IMS active subsystems in the usual way by using the IMS commands /CHECKPOINT PURGE, /CHE FREEZE, or /CHE DUMPQ.

When DBRC and ILS stop, IMS deallocates all conversations between TMSs and closes the VTAM ACB. Existing TMS conversations between active and tracking IMS subsystems are not affected by shutdown of the TMS address space. IMS cleans up component-related resources as IMS components terminate.

Recommendation: Consider the TMS a never-ending job, closely allied to VTAM. If VTAM fails or is shut down, the TMS shuts itself down. Except during VTAM shutdown, you should not terminate the TMS while IMS RSR (Remote Site Recovery) work is active or anticipated.

Normal RSR termination at the tracking site

When RSR terminates at the tracking site, the log router, the DL/I database tracker, and the Fast Path database tracker all perform various operations as part of shutting down.

Log router

During normal termination, the log router:

- Deallocates all current conversations

The log router deallocates all conversations between the log router and its partners at the active site (the loggers and ILS in the global service group). The log router does not allocate or accept any new conversations after the start of normal shutdown processing. The log router also disconnects from the TMS.

- Processes all active IMS log data in the system

When normal shutdown processing begins, the log router can have log data that has been received but not yet stored and processed. In this case, IMS stores all active IMS log data held by the log router in the appropriate SLDSs, and closes and records these SLDSs in the RECON data set.

- Stops online forward recovery processing

When the log router begins its shutdown processing, it stops online forward recovery processing. When the tracking subsystem restarts, IMS resumes any online forward recovery processing where it left off, unless the databases were recovered offline.

- Stops automatic archive
IMS stops all automatic archive processes immediately during normal shutdown.
- Deletes unused, preallocated SLDSs
IMS deletes all unused, preallocated SLDSs during normal shutdown.
- Initiates a termination milestone
During normal shutdown, IMS records a termination milestone.

DL/I database tracking

During normal termination, the DL/I database tracker ensures that all queued work is processed and that all altered buffers are written to the shadow databases.

Fast Path database tracking

When the database trackers receive the termination milestone from the log router, IMS begins shutdown for Fast Path database tracking. The termination milestone is the last milestone processed by the database trackers. IMS writes all received and committed updates to the shadow databases before termination. IMS ignores any update log records received but not yet committed. The log router passes these ignored records to the Fast Path database trackers during the next restart of the tracking subsystem. These records could also be processed if a remote takeover occurs.

IMS restart without remote takeover (at the active site)

Starting the active site after a termination, without remote takeover, is identical to restarting a non-RSR (Remote Site Recovery) subsystem.

The Transport Manager Subsystem (TMS) does not maintain any information across a component failure, and does not have restart or recovery functions that apply to component or CPC failure.

IMS restart without remote takeover (at the tracking site)

The tracking subsystem supports both IMS emergency restart and normal restart for recovery from tracking subsystem disruptions. A normal or emergency restart of a tracking subsystem is **not** a remote takeover.

Therefore, restarting a tracking subsystem is much like restarting any other IMS subsystem.

To provide restart capability for the tracking subsystem, IMS performs milestone processing to record additional information on the tracking subsystem's log during normal operation. These log records contain tracking status information about each active log stream being processed to enable the tracking subsystem to continue tracking where it left off prior to restart. Milestone processing occurs at regular intervals between checkpoints and requires less processing than an IMS checkpoint.

When IMS log routing is idle (that is, no new records are being routed) for five consecutive milestones, IMS ignores further milestone requests until routing activity resumes.

A cold start of the tracker erases the tracking subsystem's tracking status. After a cold start, the tracking subsystem does not resume tracking for any log streams for which tracking had not completed before the restart.

Recommendation: Cold start the tracking subsystem only after first trying warm or emergency restart and only after taking the correct steps in preparation. In fact, for a cold start other than for the first one, you receive message DFS2942I, which gives you the opportunity to change the restart command.

To prepare for a cold start, you must first do the following:

1. Make a backup copy of the RECON data set after all active subsystems' logs have been transmitted, received, archived, and routed, and after all active systems have shutdown normally.

This backup copy becomes the restored RECON data set if the tracking subsystem must be cold started.

2. Delete all tracking SLDSs created after the RECON data set backup was created.

When the tracking subsystem is cold started with the restored RECON, IMS spends some time catching up with active logs; reapplying the changes to the databases, though redundant, presents no problem.

After the tracking subsystem is running, IMS restores normal tracking.

Log router

When you restart a tracking subsystem after the tracking subsystem terminates without a remote takeover, the log router restarts automatically.

If the tracking site was receiving log data from the active site when the tracking subsystem terminated, after restart processing is complete, the log router repositions itself on the log stream that was being processed prior to termination. The log router obtains the log position of all streams that were active at the time of failure from the tracking subsystem's log records.

log position

The point in the logs at which the log router needs to start routing log records to the tracking components. IMS deletes tracked logs that are unused at the time the tracking subsystem failed as part of restart processing.

DL/I database tracking

IMS restores the status shadow databases had prior to termination. For example, a database that was stopped remains stopped. After the tracking subsystem is restarted, IMS restarts online forward recovery processes that did not complete prior to termination.

Related concepts:

"Online forward recovery at the active site" on page 178

Fast Path database tracking

The log router reprocesses log records that had been written to the tracked logs prior to termination but that had not yet been committed to the database. After the tracking subsystem is restarted, IMS restarts online forward recovery processes that did not complete prior to termination.

Related concepts:

“Online forward recovery at the active site” on page 178

Online change

The ACBLIBx, FORMATx, and MODBLKSx data sets for the active site and their corresponding data sets at the tracking site should be identical before starting the tracking subsystem for the first time. After the first time, you can use IMS online change commands to maintain these data sets.

Related concepts:

 Making online changes (System Administration)

Related reference:

 INITIATE OLC command (Commands)

DRD and RSR

Dynamic resource definition (DRD) is supported on the IMS systems at the active site. DRD is not supported on the Remote Site Recovery (RSR) tracking system. When a remote takeover occurs, DRD is supported on the new active IMS systems.

Recommendation: Keep the resource definitions synchronized between the tracking IMS system and the active IMS system. Definitional changes (using DRD) made to resources at the active site are not automatically made to the resources at the tracking site. You must perform a MODBLKS online change to make similar changes to the resources on the tracking IMS. The MODBLKS online change must be done before or with the ACBLIB online change.

If the new active IMS system will have DRD enabled and will be warm or emergency restarted, the runtime resource definitions for database, program, routing code, and transaction resources and descriptors are rebuilt from the log records.

If the new active IMS system will have DRD enabled, will have automatic import enabled, and will be cold started, you must ensure, before starting the new active system, that the most current resource definition data set (RDDS) from the active site is available to the IMS that takes over at the remote site or that you build a copy of the RDDS at the remote site.

You can build a copy of the most current RDDS by using data set mirroring and the Extended Remote Copy (XRC) tracking function of RSR or by using the DRD utilities that are provided with IMS and are accessible through the IMS Application Menu. For example, you can use either the DFSURCL0 utility to recreate the most current RDDS from the X'22', X'4001, X'4004', X'4006', X'4007', X'4083', X'4098', and X'45FF' log records that are transported from the active system to the remote site, or use the DFSURCP0 utility to copy the contents of an existing RDDS into a new RDDS.

IMS abnormal termination in an RSR environment

If DL/I database tracking abnormally terminates, any tracking activity in progress also terminates. If Fast Path database tracking abnormally terminates, all tracking activity that is in progress terminates immediately.

In the case of pseudo-abends, instead of bringing down the entire tracking subsystem, IMS generates a SNAP DUMP and stops the specific database or area. This stop processing includes tracking termination for the database or area,

database or area close, de-authorization, and deallocation. Tracking then continues for other databases and areas not affected by the pseudo-abend.

RSR does not affect IMS subsystem abnormal termination at the active site.

Transport Manager Subsystem

If an abnormal termination occurs, the Transport Manager Subsystem (TMS) generates a dump using the SDUMP z/OS macro and notifies IMS components that the TMS is terminating.

IMS components that have existing conversations can continue to use them after the TMS terminates.

You can use the z/OS CANCEL or STOP commands to terminate the TMS abnormally.

If ILS abnormally terminates, any activity that was in progress also terminates, but it does not terminate the TMS.

Isolated log transport processing in the logs at the active site

Sometimes, IMS cannot send log data to the tracking site at the same time as it writes the data to the OLDS at the active site. Sending missing log data is called *isolated log transport processing*.

For example, link failures and other disruptions on the active subsystem can cause IMS not to send log data. These gaps in log data transmission must be filled. When the log router at the tracking site detects a gap in the log data, it sends a request to the isolated log sender (ILS) to send the missing log data.

ILS is an RSR component that is required at the active site. ILS runs in the Transport Manager Subsystem (TMS) address space. It is not required at the tracking site, but you can start it there in anticipation of a remote takeover.

Use the START ILS command to start ILS. You can issue it either from a SYSIN data set as an automated process used when you start the TMS subsystem, or from a z/OS console using the z/OS MODIFY command.

If the START ILS command is entered before the TMS conversation has been established between the active and tracking sites, a conversation between ILS and the Tracking IMS may not be connected. This could happen when START TMS and START ILS commands are entered closely together such as those commands that are from a SYSIN data set of the TMS at the active site. The tracking IMS always tries to establish the conversation with ILS during initialization, therefore any recovery action is not required in most cases. However, if the tracking IMS is brought up earlier than the TMS at the active site, the operator might need to enter a /STA SERVGRP or /STA ISOLOG command at tracking site so that the connection with ILS is established. Alternately, the STOP ILS and START ILS commands can be tried again for TMS at the Active site in order to start the ILS function successfully.

When started, ILS identifies to the TMS, allocates a conversation with the log router at the tracking site, and waits for work requests from the log router. When it receives work requests from the log router, ILS obtains data set information from DBRC and reads and sends the log data to the tracking site. However, if the missing log data is contained in the OLDS that IMS is currently using, ILS cannot send the data until IMS switches to the next OLDS.

Optionally, you can specify concurrent conversations, using the MAXCONV keyword of the START ILS command. ILS uses multiple conversations if there are multiple log data sets that need to be sent.

If there are multiple TMSs at the active site, you can start an ILS in each TMS address space. However, the log router establishes conversations with only one ILS at a time.

Isolated log transport processing in the logs at the tracking site

The tracking subsystem must detect gaps in the log data and to initiate gap fill processing by sending requests to ILS at the active site. Log data that follows a gap cannot be processed until the gap has been filled. You can determine whether gaps exist and monitor gap status by using the IMS /DISPLAY TRACKING STATUS command.

Catch-up processing

When an active IMS is unable to send log data to the tracking IMS (because of a repairable cause, such as network communication disruption) and the log router recognizes a gap in received log data, the log router contacts the isolated log sender at the active site and requests the missing log data. Later, these records will be presented to the tracking IMS systems in order to bring them up to date with the latest possible active IMS log data. This procedure is called catch-up processing.

Log router

Sometimes, log records that are currently being routed to the individual database trackers are not the same log records that are currently being received from the active subsystem. This creates a situation in which the trackers must catch up with the log router.

Catch-up processing is necessary in the following sorts of situations:

- Network disruptions (involving, for example, the Transport Manager Subsystem (TMS), VTAM, communications controllers, communication lines)
- Network congestion that temporarily causes the logger to suspend sending log data
- Failure of an active subsystem that is involved in data sharing
- General failure of the tracking subsystem

Catch-up processing executes in parallel with other log router processing in order to minimize the impact of catch-up processing on the rest of the subsystem.

As the data arrives at the tracking site during catch-up processing, it is placed in SLDSs by the log router, and each SLDS is registered with DBRC, just as if it had arrived during normal processing.

Catch-up processing can include the following steps:

- The log router records data from tracked SLDSs before a tracking subsystem failure
- The log router records data from one or more tracked SLDSs during catch-up processing

- The log router processes data from main storage buffers received recently from the active subsystem's logger

The active site does not perform catch-up processing.

Performance of catch-up processing when using tape

Because of the number of data sets involved in catch-up processing, and because speed of processing is critical, using tape media for tracked SLDs for a DLT system can result in poorer performance and some unnecessary activity (such as excessive tape mounts) that delays catch-up processing. If you choose to use tape media for your tracked SLDs, be sure to minimize any operational delay.

DL/I database tracking

Catch-up processing greatly increases the database tracking REDO processing rate. To allow the tracking subsystem to handle this increase, you need to define a large number of PSTs for database tracking.

Recommendation: For the tracking subsystem, define a number of PSTs equal to twice the total number of DL/I PSTs used in all of the active subsystems tracked by the tracking subsystem.

Online forward recovery at the active site

Online forward recovery (OFR) is the process by which RSR brings shadow databases and areas that are out-of-date back to the current tracking state. OFR is available only for database level tracking (DLT) subsystems and only for databases and areas tracked at the database-readiness level (DBTRACK).

Online forward recovery does not use an image copy to recover, as do offline recovery methods, nor does it use change accumulation data sets. Instead, OFR reads the log data on the tracking SLDs and applies the updates to the shadow databases. After OFR completes for a given database or area, IMS applies subsequent updates by the normal database-tracking process.

If you have some databases or areas at the tracking site that are defined at the recovery-readiness level (RCVTRACK), you can change them to the database-readiness level (DBTRACK) by doing the following:

1. Ship an image copy of the database from the active to the tracking site.
2. Use the DBRC CHANGE.DB or CHANGE.DBDS command to change the readiness level to DBTRACK.
3. Install the shadow database using the DBRC GENJCL.RECEIVE command.
4. Issue one of the following commands to start the database or data group:
 - /START DB
 - /START DATAGRP
 - UPDATE DB START(ACCESS)
 - UPDATE DATAGRP START(ACCESS)

These commands bring the databases and areas up to current tracking level using online forward recovery.

You can also use online forward recovery when:

- You need to recover several databases concurrently.

Important: You should always plan to run just one OFR process for all of your databases that need to be recovered, rather than initiating a separate OFR process for each database. Running one OFR process for all databases greatly reduces the overhead for the recovery (for example, the log is read just once rather than repeatedly for each database).

Only one OFR can process at a time. Each OFR initiated after the first processes in turn when the previous OFR completes, fails, or is terminated by a /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command.

- You need to bring the databases up to date prior to performing a planned remote takeover.
- You need to schedule regular off-shift work, including bringing RCVTRACK shadow databases defined up to date.

Depending upon how long the database or area has been offline, or how much log data must be processed for online forward recovery, you might choose to send a new image copy of the database or area from the active site before beginning OFR processing. IMS excludes log data created by the active subsystem SLDs before the image copy from the log input used for online forward recovery processing.

DBRC maintains the tracking status of each shadow database and area, and it knows when a shadow database or area is not up to date.

Use a display command (such as, /DISPLAY TRACKING STATUS, /DISPLAY DB OFR, /DISPLAY AREA OFR, QUERY DB STATUS(OFR), or QUERY AREA STATUS(OFR)) to check the progress of OFR. When OFR has caught up to current database tracking, you see a message (DFS4056I for full-function databases, DFS2989I for DEDB areas) that says the database is started.

To stop online forward recovery, use the one of the following commands to stop the database or area:

- /DBRECOVERY DB
- /DBRECOVERY AREA
- UPDATE DB STOP(ACCESS)
- UPDATE AREA STOP(ACCESS)

You can issue these commands for a single database or area. To issue these commands for a group of databases or areas, use the DATAGRP keyword. If part of a group is stopped during OFR, recovery continues for the other databases or areas in the group. If the entire group is stopped, online forward recovery stops for the entire group.

Other commands also stop online forward recovery, such as /CHECKPOINT FREEZE. OFR automatically restarts its processing when the tracking subsystem restarts.

The active site does not perform online forward recovery.

Related concepts:

“DL/I database tracking” on page 174

“Remote takeover processing at the tracking site” on page 187

“Fast Path database tracking” on page 174

Remote takeover

When a remote takeover occurs, the tracking site takes over the workload of the active site. DBRC is involved in remote takeover, both at the original active site and the new active site. A remote takeover can be either planned or unplanned.

To DBRC, the main effect of a remote takeover is to reverse the roles of the two service groups (SGs); the original tracking SG becomes the active SG, and the original active SG can later become the new tracking SG (if this is possible and appropriate).

Planned remote takeover

All work is transferred from the active site to the tracking site. You might use a planned remote takeover, for instance, to minimize system downtime during a hardware upgrade or site maintenance, or if you recognize that a disaster is imminent.

You initiate a planned remote takeover from the active site. The result of a planned takeover is an orderly shutdown of the active subsystems, and a known state of the subsystems at both sites.

Unplanned remote takeover

All work is transferred from the active site to the tracking site, just as with a planned remote takeover. An unplanned remote takeover can be used when a sustained outage in system availability occurs at the active site, as in the case of a natural disaster.

You initiate an unplanned remote takeover from the tracking site. The result is a known state of the subsystem at the tracking site, but the state of the original active site is unknown. When the takeover occurs, IMS cannot do an orderly shutdown of the original active site, so the log records for any transactions that are being processed at the original active site at the time of takeover might not reach the tracking site.

The remote takeover process at the new active site is the same regardless of the type of remote takeover. The tracking subsystem completes its work and terminates. After all log data is written to SLDSs and (for databases tracked at the database readiness level) applied to the databases, the tracking subsystem informs DBRC that tracking is complete. DBRC then updates the RECON data set to make the local SG an active SG so that active IMS subsystems can be started at the new active site to complete the remote takeover.

At the original active site, the remote takeover process depends on the type of remote takeover:

- For a planned remote takeover, the covered databases are in a known state and no action is required at the original active site.
When the last active subsystem terminates, DBRC updates the GSG record in the RECON data set, making the local SG the tracking SG. After this is done, you can start a tracking subsystem at the old active site.
- For an unplanned remote takeover, there is no remote takeover process. The active subsystems effectively disappear.

However, if the site is intact, you might want to make the original active site a new tracking site. Unfortunately, because the takeover was unplanned, the databases at the old active site are not in a consistent state with the new active site. To make the old active site into a new tracking site, you must transport image copies of all covered databases from the new active site. This process is the same as initializing any tracking site for the first time, with one exception; you must use the RESET.GSG command to remove all recovery information for the covered databases of the GSG from the original active RECON data set.

Related tasks:

“Starting a new tracking subsystem” on page 190

Related reference:

 RESET.GSG command (Commands)

Planned remote takeover

To initiate a planned remote takeover, issue the /RTAKEOVER command from at least one **active** subsystem. All databases must be in a consistent state for a planned remote takeover to succeed.

To ensure their consistency, be sure that all active IMS subsystems have shut down without error.

IMS rejects the /RTAKEOVER command if DBRC has a subsystem record that is marked as abnormally terminated. This record indicates that an online system abended but has not been restarted, or that a batch job has abended and batch backout has not been run.

If IMS rejects the command, you need to restart the failed online subsystem, run batch backout for failed batch jobs, or convert the planned remote takeover to an unplanned one.

One consequence of the /RTAKEOVER command is that no IMS subsystem can normally sign on to DBRC. DBRC still permits recovery signon in order to allow abnormally terminated subsystems to be restarted and shutdown.

The tracking site receives notification of the remote takeover through information in the transmitted log data from the active site. You must then stop other active online IMS subsystems in the same global service group (GSG) using either the /CHECKPOINT FREEZE or /RTAKEOVER command. Allow any active batch DL/I jobs to run to completion.

When all subsystem and batch termination completes, DBRC changes the SG's role at the active site from active to tracking. After the SG's role is changed to tracking, DBRC does not permit any active subsystems to be started at that site.

If you want to recover the message queues in a non-shared-queues environment at the tracking site after the remote takeover, issue one of the following commands on all active subsystems in the GSG before issuing the /RTAKEOVER command:

- /CHECKPOINT DUMPQ
- /RTAKEOVER DUMPQ

Issuing one of these commands minimizes restart time at the tracking site. If you do not use the DUMPQ keyword on either command, the new active subsystem will use a DUMPQ from a prior checkpoint.

Recommendation: Ensure that an ILS is running during planned remote takeover. If ILS is not running and the tracking subsystem did not receive all log data from the active loggers, the remote takeover will not complete.

Related concepts:

“Planned remote takeover and XRC tracking” on page 194

“Unplanned remote takeover” on page 183

“How to start active subsystems after remote takeover” on page 188

Related tasks:

“Performing a planned remote takeover”

Performing a planned remote takeover

A planned remote takeover allows you to transfer all work from the active site to the tracking site. You might do this, for instance, to minimize system downtime during a hardware upgrade or site maintenance, or if you recognize that a disaster is approaching or imminent.

At the beginning of this sample takeover, the active subsystem is at site 1 and the tracking subsystem is at site 2.

1. Issue the /RTAKEOVER command.

During the remote takeover, you see the following messages, in addition to the usual IMS shutdown messages. The messages shown below are displayed at either the active or tracking sites and are also displayed on either the system console or the IMS MTO. The active subsystem (at site 1) issues the following messages: DFS2939I, DFS058I, DFS4036I, DFS4024I, and DFS4078I. The tracking subsystem (at site 2) issues the following messages: DFS4124I, DFS2932I, DFS2913I, DFS2500I, DFS4125I, and DFS4126I. In the following sample message output, the active subsystem is known as IMSA in service group GLEN and the tracking subsystem is known as IMST in service group (SG) DAVE.

```
DFS2939I REMOTE SITE PLANNED TAKEOVER IN PROGRESS
DFS058I  hh mm ss RTAKEOVER COMMAND IN PROGRESS
```

```
DFS4124I PLANNED TAKEOVER REQUESTED
DFS2932I DATABASE UPDATES PRIOR TO SYSTEM  IMSA TAKEOVER HAVE BEEN ROUTED
```

```
DFS4036I CONVERSATION ENDING WITH SERVICE GROUP DAVE
DFS4036I CONVERSATION ENDED WITH SERVICE GROUP DAVE
DFS4024I STOP SERVGRP PROCESSING (INTERNAL) COMPLETE
```

```
DFS2913I CONVERSATION WITH IMS  IMSA TERMINATED: ACT SYS SHUTDOWN
DFS2500I DATASET  IMZnnnnn SUCCESSFULLY CREATED|DELETED|ALLOCATED
DFS4125I PLANNED TAKEOVER IN PROGRESS
```

```
DFS4078I ILS FOR GSG STL HAS LOST CONTACT WITH SG.SYSTEM: DAVE.TMST
```

```
DFS4126I TAKEOVER COMPLETE
```

2. List the RSR records in the RECON data set (optional step).

Check that DBRC has switched the SG values in the RECON data sets for IMSA at site 1 and IMST at site 2. For example, the RECON data set for IMST at site 2 should show the local SG as active.

The SSYS record should have been removed in the RECON data set associated with IMSA at site 1 and the RECON data set at site 2 should have a SSYS record for the IMSID of the yet-to-be-activated IMSA at site 2, with ABNORMAL TERMINATION set to OFF and TRACKER TERM set to ON.

3. Change the startup parameters for IMSA at site 2.

The tracking subsystem at site 2 had IMSID=IMST and TRACK=DLT; change these (or use PARM1= on the EXEC statement) to IMSID=IMSA (the IMSID of the active subsystem from site 1) and TRACK=NO.

Switch the APPLIDs for all subsystems. APPLID1, which always specifies the active's APPLID, should be changed to the 'site 2' active APPLID. APPLID3 should be changed to yet to be started site 1 tracking subsystem's APPLID. Also check for the presence of USERVAR in the DFSPBxxx member. If a USERVAR was specified for the active system before the takeover, the same USERVAR should be specified for the IMSA active subsystem at site 2.

4. Start the IMSA active subsystem at site 2 using the /ERESTART command.

The following messages are displayed on the new active subsystem:

```
DFS2946I  SYSID=IMSA HAS ASSUMED THE ACTIVE ROLE TO SERVICE GROUP DAVE
DFS994I   RSR TAKEOVER VIA *BUILDQ* EMERGENCY START COMPLETED
```

5. Change the startup parameters for the new tracking subsystem at site 1.

The IMSA active subsystem at site 1 had IMSID=IMSA and TRACK=NO. Change the IMSID to IMSx (or use PARM1= on the EXEC statement). The new tracking subsystem IMSID does not have to be the same as the old tracking subsystem IMSID, but in the examples below, IMST is used. Also, change TRACK to RLT or DLT, as required.

Switch the APPLIDs for all subsystems. APPLID1, which always specifies the active's APPLID, should be changed to the APPLID for IMSA at site 2. APPLID3 should be changed to the APPLID for the new tracking subsystem at site 1.

6. Start the new tracking subsystem by issuing the following command:

```
/NRE CHKPT 0 FMT ALL
```

You see message DFS2923I and many DFS2500I messages, in addition to the regular IMS startup messages.

7. Begin tracking by issuing a /SWITCH OLDS command in IMSA at site 2 or by issuing the /START SERVGRP command on the new tracking subsystem at site 1.

You see DFS4036I and DFS4024I messages as tracking begins.

8. Perform catch-up, if required.

If START ILS is included in the TMS startup member, catch-up processing is started automatically. Otherwise, issue F TMSA,START ILS(*ilsname*) on IMSA at site 2.

9. Perform online forward recovery, if required.

Issue one of the following sets of commands on the new tracking subsystem at site 1 to start the all databases and areas:

- /START DB ALL and /START AREA ALL
- UPDATE DB START(ACCESS) and UPDATE AREA START(ACCESS)

Related concepts:

"How to start active subsystems after remote takeover" on page 188

"Planned remote takeover and XRC tracking" on page 194

"Planned remote takeover" on page 181

Unplanned remote takeover

To initiate an unplanned remote takeover, issue the /RTAKEOVER UNPLAN command from the **tracking** subsystem.

Whenever possible, do the following before initiating an unplanned remote takeover:

- Terminate active subsystems.

- For an RLT tracking subsystem, shut it down and restart it as a DLT tracking subsystem, then let online forward recovery recover the databases prior to the unplanned remote takeover.
- Allow ILS processing to complete (to fill gaps in the log data).
After the /RTAKEOVER command has completed processing, IMS does not accept further ILS data, which can increase the likelihood that you will have truncated logs.

Log truncation

The process by which IMS removes log records from tracking SLDSs so that IMS does not apply invalid data to the databases after the unplanned remote takeover is complete.

The following situations can cause log truncation:

- If some log records for an active subsystem are missing at the tracking site, IMS truncates the records beyond the missing data.
- When log streams must be merged in creation-time sequence before routing to the tracking components, routing for all streams stops at the point that a gap is encountered on a log stream or when there is no more log data for a stream.

When an unplanned remote takeover occurs in these situations, IMS truncates the subsystems' logs at the point where log routing has stopped. The log router does not delete the truncated data, however, because you might be able to use it, even though IMS cannot.

After IMS truncates a log, IMS issues a message (DFS2936I) with the truncated SLDS names, including the log record ID of the truncation point. You can then look at the data in the log records offline and take appropriate action.

If you specify the NOREVERSE keyword on the /RTAKEOVER UNPLAN command, you must not reverse the remote takeover. The tracking subsystem saves and processes all data it has received from the active site, regardless of whether that data was committed (written to the OLDS) on the active subsystem. In this case, IMS issues message DFS4121I.

If you do not specify NOREVERSE keyword, the tracking subsystem discards any uncommitted data it has received from the active subsystem.

Related concepts:

"Planned remote takeover" on page 181

"How to start active subsystems after remote takeover" on page 188

Related tasks:

"Performing an unplanned remote takeover"

"Unplanned remote takeover and XRC tracking" on page 194

Related reference:

 /RTAKEOVER command (Commands)

Performing an unplanned remote takeover

An unplanned remote takeover allows you to transfer all work from the active site to the tracking site, just as a planned remote takeover does. You initiate an unplanned remote takeover when a sustained outage in system availability occurs at the active site, as in the case of a natural disaster.

The output shown assumes the active subsystem is known as IMSA in service group GLEN and the tracking subsystem is known as IMST in service group DAVE. In the following scenario, the current active subsystem (IMSA) is running at site 1 and the tracking subsystem (IMST) is running at site 2.

1. Start the remote takeover.

Issue the /RTA UNPLAN command on the tracking subsystem. You see the following messages on the tracking subsystem as IMSA shuts down. No messages are shown for the active subsystem because its state is unknown (and unimportant) for an unplanned takeover.

```
DFS058I hh mm ss RTAKEOVER COMMAND IN PROGRESS
DFS4123I UNPLANNED TAKEOVER IN PROGRESS
DFS4126I TAKEOVER COMPLETE
DFS994I *CHKPT ynnn/hhmmss**FREEZE*
```

2. Check the RECON data set for the tracking subsystem.

DBRC switches the SG records, for example, GLEN (the old active SG) is now the tracking SG and DAVE is set to ACTIVE and LOCAL.

3. Restart an active IMS subsystem at site 2.

To restart an active subsystem at site 2, perform the following steps (same as for a planned takeover):

- Change the startup parameters for the yet-to-be-started IMSA subsystem at site 2.

The tracking subsystem (IMST at site 2) had IMSID=IMST and TRACK=DLT; change these (or use PARM1= on the EXEC statement) to IMSID=IMSA (the IMSID of the active subsystem from site 1) and TRACK=NO.

Switch the APPLIDs for all subsystems. APPLID1, which always specifies the active subsystem's APPLID, should be changed to the APPLID of IMSA at site 2. APPLID3 should be changed to the APPLID of the tracking subsystem at site 1. Also check for the presence of USERVAR in the DFSPBxxx member. If a USERVAR was specified for the active system before the takeover, the same USERVAR should be specified for the new active system.

- Start an active subsystem at site 2 using the /ERESTART command.

The following messages are displayed on the new active subsystem:

```
DFS2946I SYSID=IMSA HAS ASSUMED THE ACTIVE ROLE TO SERVICE GROUP DAVE
DFS994I RSR TAKEOVER VIA *BUILDQ* EMERGENCY START COMPLETED
```

4. Start a tracking IMS subsystem at site 1.

To start a tracking subsystem at the new tracking site, perform the following steps:

- Stop the TMS on the no-longer active subsystem at site 1.

From the z/OS console, issue P TMSA.

The TMS has the RECON data sets allocated and you need to redefine and re-initialize them.

- Reallocate the RECON data sets and database data sets on the new tracking subsystem at site 1.

Remember to change the SG parameters (the SG is now local), and check the RCVTRACK | DBTRACK parameters in the INIT.DB and INIT.AREA commands.

- Ship the image copies and restore them.

Image copies of the database data sets that are going to be tracked must be made and shipped to the new tracking site. The image copies must be registered with DBRC, and if you are using database-level tracking you must generate jobs to restore them to shadow database data sets.

Be sure to issue one of the following sets of commands for all databases and areas on the new active subsystem:

- /DBRECOVERY DB ALL and /DBRECOVERY AREA ALL
- UPDATE DB STOP(ACCESS) and UPDATE AREA STOP(ACCESS)

When all the image copies have been received, issue one of the following sets of commands to start all databases and areas on the new active subsystem:

- /START DB ALL and /START AREA ALL
- UPDATE DB START(ACCESS) and UPDATE AREA START(ACCESS)
- Start a TMS on the new tracking subsystem.

From the z/OS console, issue S TMSAP,SUB=JES2.

Related concepts:

“How to start active subsystems after remote takeover” on page 188

“Unplanned remote takeover” on page 183

Related tasks:

“Unplanned remote takeover and XRC tracking” on page 194

Related reference:

 /RTAKEOVER command (Commands)

Canceling a remote takeover

After the tracking subsystem has successfully shut down (in preparation for remote takeover processing), you have an opportunity to change your mind about the remote takeover and reverse the process. This reversal can only be performed in the operational window following the /RTAKEOVER command (for either planned or unplanned remote takeover) and before any activity that makes the reversal impossible.

Any of the following actions commit you to the remote takeover and make reversal impossible:

- Using the NOREVERSE keyword on the /RTAKEOVER UNPLAN command
- Entering an /ERESTART or /NRESTART command on any IMS subsystem at the new active site
- Issuing a DBRC RESET.GSG command at the old active site
- Submitting a DL/I batch job at the new active site
- Submitting a DL/I Batch Backout utility job at the new active site

To reverse a remote takeover, you must issue the DBRC CHANGE.SG command with the NORTA keyword:

- At both sites for planned remote takeovers
- At the remote site for unplanned remote takeovers

The NORTA keyword indicates that the remote takeover is to be reversed and that the service groups are to continue in their current roles (at their current locations).

Before issuing the CHANGE.SG NORTA command, wait for a DFS4123 or DFS4125 message saying that the planned or unplanned takeover is in progress. Using this command before the active or tracking subsystem has acknowledged that the remote takeover is in progress does not guarantee that IMS has updated the RECON data set to reflect the proper state and that the active or tracking

subsystem recognize the reversal. If you use the CHANGE.SG NORTA command outside of the allowed remote takeover window, IMS issues message DSP0144I.

If you specify the NOREVERSE keyword on the /RTAKEOVER command, IMS does not prevent you from reversing the takeover. However, if you do reverse the takeover, integrity of the log and databases at the tracking site is unknown. IMS issues message DFS4122A to warn you of potential problems.

Remote takeover processing at the tracking site

During a remote planned or unplanned takeover, the tracking subsystem completes tracking and then tells DBRC to change the service group to an active role, before shutting down.

When this is complete, you can start the IMS subsystems that were running at the old active site at the new active site.

Log router

During remote takeover, the log router suspends processing and performs the following functions:

- Processes all available log data
- Deallocates all existing conversations with the active sites
- Writes all buffers received from the active sites to SLDs
- Routes log records for DBTRACK databases to the tracking components
- Halts any online forward recovery processing

After remote takeover, you need to run the appropriate recovery utilities for all databases for which online forward recovery did not complete. If you want to use OFR to recover the databases, delay the remote takeover until OFR completes.

- Truncates the log when necessary

For an unplanned remote takeover, some of the log data might be truncated so that updates that are out of sequence are not applied to the shadow databases.

DL/I database tracking

After a remote takeover is initiated, the DL/I database tracker processes all remaining update records routed from the log router, and applies all relevant database updates to the shadow databases.

IMS terminates any online forward recovery activity still in progress. Databases that were in the process of being recovered with OFR require batch recovery before they can be available for transaction processing.

Fast Path database tracking

After a remote takeover is initiated, the Fast Path database tracker stops processing just as if the tracking subsystem had been shut down.

Batch DL/I handling

After remote takeover completes, any batch DL/I work that was not completed or backed out at the old active site must be backed out at the new active site (the old tracking site) using the Batch Backout utility.

Restriction: Do not begin any Batch Backout jobs until you receive the DFS4126I message telling you that the remote takeover is complete. The information in the RECON data set is not complete until the tracking subsystem has successfully shutdown after processing the /RTAKEOVER command.

Related concepts:

“Online forward recovery at the active site” on page 178

How to start active subsystems after remote takeover

For each IMS online subsystem at the old active site, you need to start an equivalent IMS subsystem at the new active site. After a remote takeover, use an appropriate IMS restart command to start the new active subsystems.

Before starting your active subsystems at the new active site, you need to ensure that the remote takeover, planned or unplanned, completed successfully (as indicated by message DFS4126I). You also need to be sure that you have copies of all of the IMS libraries, including the online change libraries (ACBLIBx, FORMATx, and MDBLKs), at the new active site.

Table 25. Starting a new active subsystem

Takeover Type	To Save TM Messages	To Discard TM Messages
Planned Remote Takeover	/NRE BUILDQ	/NRE CHECKPOINT 0 or /ERE COLDCOMM or /ERE COLDSYS
Unplanned Remote Takeover	/ERE BUILDQ	/ERE COLDCOMM or /ERE COLDSYS

Restriction: In a shared-queues environment, you cannot rebuild the shared queues using the IMS restart commands. See “Sharing message queues after a planned remote takeover” on page 189 for information about restoring the shared-queues environment.

For unplanned remote takeover, you must ensure that the subsystem records are removed from the RECON data set for all occurrences of active site IMS subsystems. To do this, you must restart (/ERESTART) each one, or use the DBRC DELETE.SUBSYS command (along with any required database recoveries and back-outs).

You must allocate DBDSs and DEDBs for RCVTRACK databases and areas and you must recover them using the database recovery utilities or the change accumulation utilities. In this way, you can make these databases and areas ready for work at the new active site. Because these recoveries are generally done serially, recovery could take a long time.

DBTRACK databases and areas should be ready for work at the new active site with no further processing.

After remote takeover, all shadow database copies become master databases. However, conditions that reflect physical database and DBDS errors are not preserved because these errors did not occur at the old tracking site (the new active site). IMS obtains ADS DD names and data set names from the RECON data set rather than from the log.

The indirect list data sets and primary index data sets for all partitions of partitioned HDAM (PHDAM) and partitioned HIDAM (PHIDAM) databases need

to be rebuilt using the Index/ILDS Rebuild utility (DFSPREC0) before the partitions are ready for work at the new active site.

Handling in-flight data after an unplanned remote takeover

During the first startup after a remote takeover, IMS backs out in-flight data for databases that were both authorized to the tracker and registered as covered at the time of the remote takeover.

If in-flight data exists for a non-covered and non-authorized database, IMS does nothing for the database.

If in-flight data exists for a non-covered but authorized database, IMS issues message DFS981I indicating the database was stopped due to a remote takeover.

If in-flight data exists for a covered but non-authorized database, IMS schedules a deferred or restartable backout for the database, and issues message DFS981I indicating that the database is stopped due to remote takeover.

Handling in-doubt data after an unplanned remote takeover

IMS backs out any in-doubt data that exists after a remote takeover immediately after the first startup completes. Thus, it looks exactly like in-flight data to a new active subsystem, and the actions performed, depending upon authorization and covered status, are identical to those described for in-flight data.

The tracker tracks in-doubt extended error queue elements (EEQEs) to protect the data until the creating-system is started after a remote takeover. The data is thus protected from other subsystems. IMS only registers in-doubt EEQEs at the tracking site, because the in-doubt state is a logical condition that can exist at both sites. Write errors are physical conditions that do not usually exist at both sites. IMS only registers write-error EEQEs if they occur on the tracking subsystem.

Handling in-doubt data after a planned remote takeover

Only two actions are possible for in-doubt data at remote takeover:

- If the in-doubt data exists for a non-covered database, IMS does nothing for that database.
- If in-doubt data exists for a covered database, IMS schedules a deferred or restartable backout for the database, and issues message DFS981I indicating a backout failure.

Recommendation: Because during a cold start, IMS loses all traces of in-doubt data, do not attempt a cold start if there is the possibility that in-doubt data exists at the time of the remote takeover.

Because IMS maintains in-doubt data across warm and emergency restarts, IMS resolves in-doubt data by backing it out for covered databases only.

Sharing message queues after a planned remote takeover

To rebuild shared queues at the new active site (after a planned remote takeover), issue the /CQSET SHUTDOWN SHAREDQ ON command before issuing the /RTAKEOVER command on the active IMS. The /CQSET command tells the IMS Common Queue Server (CQS) to take a structure checkpoint during CQS

shutdown. CQS writes the structure checkpoint to the CQS structure recovery data set (SRDS), which can be used to rebuild the structures on another coupling facility (at the remote site, for example). You could also use the remote dual copy feature of 3990-6 DASD to make these data sets available at the remote site.

After the remote takeover, start active IMS systems at the remote site.

Restriction:

- You cannot use the /ERE BUILDQ command to rebuild the shared message queues. In a shared-queues environment, IMS ignores the BUILDQ keyword.
- You cannot rebuild shared queues after an **unplanned** remote takeover.

Related concepts:

“Planned remote takeover” on page 181

“Unplanned remote takeover” on page 183

Related tasks:

“Performing an unplanned remote takeover” on page 184

“Performing a planned remote takeover” on page 182

Starting a new tracking subsystem

When the new tracking subsystem has caught up and is tracking normally, you can initiate a planned or unplanned remote takeover.

Recommendation: To prevent the tracking subsystem from requesting excessive amounts of data, complete the following steps before reinstalling the tracking subsystem:

1. Shut down all active IMS subsystems.
2. Issue the DBRC DELETE.GSG command.
3. Issue the DBRC INIT.GSG and INIT.SG commands to redefine the GSG.
4. Issue the DBRC CHANGE.DB | DBDS command to register the databases and areas as covered.
5. Restart the active subsystems.

By following the previous steps, the IMS tracking subsystem does not request data created before the DELETE.GSG command was issued.

When the new tracking subsystem has caught up and is tracking normally, you can initiate a planned or unplanned remote takeover.

Sometime after a remote takeover, you might want to start a tracking subsystem at the old active site or at another site. If you want to use the old active site, the first thing to do after a remote takeover, especially an unplanned one, is to return the operating environment of the old active site to a stable environment, if possible.

After a planned remote takeover, you can start a tracking subsystem at the old active site (or other site) on any CPC that has access to the RECON data sets used by the old active sites. Although not required, you might have the tracker use the same SSID as any one of the old active subsystems. No extra steps are necessary to begin tracking the new active subsystems.

To start a new tracking subsystem at the old active site after an unplanned remote takeover, perform the following steps:

1. Issue the DBRC RESET.GSG command, or rebuild the RECON data set as if you were doing an IMS initial installation. If you use the RESET.GSG command, you also need to use the CHANGE.SG command to change the local role to tracking.
2. Obtain image copies of the databases to be tracked from the new active site, and install them as you did for the old tracking site.
3. Cold start a tracking subsystem, using the FORMAT ALL keywords on the /NRESTART command.

If you rebuild the RECON data set after an unplanned remote takeover, the new tracking subsystem asks for **all** log data that was created for the GSG. This includes log data tracked before the remote takeover. If you issue the DBRC RESET.GSG command, IMS only requests log data after the remote takeover. In either case, if the new active site has been running for a long time before you are able to establish a new tracking site, the tracking subsystem could ask for weeks or months of data.

Related concepts:

“Remote takeover” on page 180

Related tasks:

“Shadowing data to the tracking site” on page 153

Discontinuing use of RSR

If you need to discontinue Remote Site Recovery (RSR), you can still test the active IMS subsystem.

If it is necessary to discontinue testing of RSR, yet continue testing the active IMS subsystem, you can do either of the following (assuming the same IMS subsystem is to be used):

- Continue without the tracking subsystem, with no change of parameters.
You can continue without the TMSs as well.
- Restart the active subsystem with a new RSR member, specifying RSR(N0).
In this case, there will be authorization failures against the RSR-tracked databases when they are accessed. To avoid these failures, issue the DELETE.GSG command.

How to coordinate recovery with DB2 for z/OS in a data sharing environment

In data sharing environments where XRC transmits DB2 for z/OS data to a remote site, RSR provides support for coordinating disaster recovery for IMS and DB2 for z/OS. This support allows IMS and DB2 for z/OS logs to be synchronized and ensures that IMS and DB2 for z/OS databases are consistent if a disaster occurs.

XRC is a separate hardware function that tracks DB2 for z/OS logs. Likewise, RSR tracks IMS logs. RSR coordinates with XRC so that both IMS and DB2 for z/OS logs are tracked and can be synchronized. To request XRC tracking, in the DFSRSRxx IMS.PROCLIB member, you specify the session ID of the XRC session that is shadowing the DB2 for z/OS volumes. As a result, the RSR tracking subsystem processes and synchronizes with XRC processing.

XRC marks all updates with a time stamp on the volumes in the XRC session and uses the time stamps to control updates to XRC-managed secondary volumes. A time stamp called *XRC consistency time* is used to indicate what updates have been applied to the secondary volumes.

When engaged in XRC tracking, the IMS routing point is kept behind the XRC consistency time. The log router maintains the last known XRC time. When the log router encounters a log record that has a time stamp greater than the XRC time, the log router gets a new time from XRC. If the XRC time has not changed, routing will be suspended until a new XRC time has been obtained. Repeated attempts to obtain a new time will occur.

How to enable XRC tracking

To enable XRC tracking, specify the XRC session ID as a startup parameter for the IMS RSR tracking subsystem. You specify the XRC session ID in the DFSRSRxx IMS.PROCLIB member used by the RSR tracking subsystem.

This XRC session must shadow the volumes containing the active logs and BSDSs for the DB2 for z/OS systems that need to be synchronized with the IMS systems.

Related reference:

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

How to stop XRC tracking

To stop XRC tracking, issue the /STOP XRCTrack command.

The tracking system does not know when to stop XRC tracking. For example, the tracking system cannot determine if the DB2 for z/OS systems have shut down, and has no knowledge of what data is on the XRC volumes. The tracking system keeps the IMS routing point behind the last known XRC consistency time until it is explicitly told to stop XRC tracking.

You might encounter situations when it is desirable to keep IMS RSR tracking running even though the IMS tracking logs are out of sync with the DB2 for z/OS tracking logs. For example, if there is a problem with the XRC session, you might want to continue IMS tracking so that, when the XRC session is available, only the XRC data needs to be re-synchronized.

How to resume XRC tracking

To resume XRC tracking, issue the /START XRCTrack command. Restarting XRC tracking causes the resynchronization of IMS tracking with DB2 for z/OS tracking.

Note: /START XRCTrack cannot be used to enable XRC tracking. If XRC parameters were not specified in the DFSRSRxx member when the tracking subsystem was started, then XRC tracking is not available and a /START XRCTrack command will be rejected (message DFS4110I will indicate that the command is not valid).

Routing delays caused by XRC tracking

Message DFS4035A is issued to alert you that routing has been unable to proceed for a short period of time (approximately 2 minutes) because the XRC consistency time has not changed. When approximately 10 minutes have passed and routing is still suspended, the message is issued again and is repeated at regular 10 minute intervals until the consistency time has advanced.

Routing is also suspended when the tracking subsystem's request to XRC to get the consistency time fails with an error return code. Check the status of the XRC session and correct the problem. If the applications that were updating the XRC volumes have terminated, issue the /STOP XRCTrack command to allow routing to proceed without considering the XRC consistency time.

Gap processing

When XRC tracking is enabled on the tracking subsystem, log routing stops when a missing log gap is detected or when there is no conversation between the log router on the tracking subsystem and an isolated log send (ILS) component at the active site.

If there is a missing log gap, the log router acts on the presumption that the missing IMS log needs to be synchronized with the XRC volumes (and with other IMS logs that are being received), so routing is suspended.

If the log router is not in communication with ILS, routing is suspended because there could be missing log gaps that the log router is not able to detect.

The tracking system should always have a conversation with ILS so that gaps are filled quickly and the tracking system stays current with the active systems. If there is more than one Transport Manager Subsystem (TMS) at the active site, ILS should be started in each so that the log router can quickly establish a conversation with another ILS in the event that the ILS to which it is connected fails. Use the `/DISPLAY TRACKING STATUS` command on the tracking subsystem to determine if there is a conversation with ILS.

Shutdown processing

When the tracking subsystem processes a shutdown request and XRC tracking is not enabled, routing of the received log data is not immediately terminated. Rather, the log router routes the log streams as far as it can, until it reaches the end of the received log data or reaches a gap in the received log data.

If the tracking subsystem processes a shutdown request and XRC tracking is enabled, routing might not be able to proceed as far as it might have compared to a tracking subsystem that is not performing XRC tracking. After the log router has received the terminate request, routing will terminate as soon as an XRC query does **not** return a time stamp that allows routing to proceed. Otherwise, the tracking subsystem could be delayed indefinitely. For example, if activity on the XRC volumes is low or intermittent while RSR is producing log data at a steady rate, the XRC consistency time could remain constant for long periods of time and log routing would not be able to route all of the received log data.

When the log data is not completely routed, the remote site is not as current as it could have been. When the tracking subsystem is restarted, log routing will resume where it left off, provided that the XRC consistency time allows routing to proceed.

Restart considerations

When the tracking subsystem is restarted or cold started with the routing position restored from the position data set, it notes the current and previous status of XRC tracking and take action as appropriate.

If the previous tracking subsystem was performing XRC tracking but XRC parameters are no longer present in the DFSRSRxx member, message DFS2919A is issued to alert the user that IMS tracking is no longer being synchronized with XRC shadowing. XRC tracking is not initiated.

If XRC tracking was stopped by a `/STOP` command on the previous tracking subsystem and was not resumed, message DFS2920I is issued and XRC tracking remains stopped. A `/START XRCTACK` command must be entered to resume XRC tracking.

Planned remote takeover and XRC tracking

The processing of a planned remote takeover for the tracking subsystem is the same whether or not Extended Remote Copy (XRC) tracking is enabled. However, if routing encounters delays while waiting for a more current XRC consistency time, shutdown of the tracking subsystem delays.

In a normal scenario where the XRC session is active and the consistency time is advancing frequently and more or less steadily, there should be little or no delay in the shutdown.

Related concepts:

“Planned remote takeover” on page 181

Related tasks:

“Performing a planned remote takeover” on page 182

Unplanned remote takeover and XRC tracking

During an unplanned remote takeover, any logs that did not route to completion (because of log gaps or issues related to the Extended Remote Copy (XRC) consistency time) are truncated at their current routing point.

If this truncation occurs, the message DFS2933I is written to the system console (and also written to the MTO for a DB/TM or DCCTL tracking subsystem). This message includes a time stamp that represents the truncation point. You should capture this time stamp and specify it in the conditional restart records for the DB2 for z/OS systems.

Follow these procedures after a decision has been made to perform an unplanned remote takeover.

1. Issue either the XEND or XSUSPEND command. If the XRC session is still active when you want to initiate an unplanned takeover, end the session by issuing the XRC TSO command, XEND, or suspend it by issuing the XRC command, XSUSPEND.
2. Issue the XRECOVER command on the remote system to complete the updating of the secondary XRC volumes. XRECOVER also labels the secondary volumes with the volume serial numbers of the primary volumes.
3. Enter an /RTAKEOVER UNPLANNED command on the tracking subsystem. To allow the tracking subsystem to route the receive log data as far as possible, the /RTA UNPLANNED command must be issued after recovery (XRECOVER command) of the XRC volumes has completed.

If you enter the /RTAKEOVER command before recovery of the XRC session has occurred, the tracking subsystem, shutting down for unplanned takeover, might terminate routing prematurely.

Related concepts:

“Unplanned remote takeover” on page 183

Related tasks:

“Performing an unplanned remote takeover” on page 184

Related reference:

 /RTAKEOVER command (Commands)

IMS recovery using Extended Recovery Facility

To operate IMS in an Extended Recovery Facility (XRF) complex, you must understand what tasks to perform during each XRF phase, from initialization through termination.

- **Initialization:**

- How to bring up the active IMS subsystem
- How and when to bring up the alternate IMS subsystem

- **Synchronization:**

- Whether or not to enter the /CHECKPOINT SNAPQ command

- **Tracking:**

- How to use the /DISPLAY command and read the status line of the master terminal
- How to use online change
- How to change surveillance dynamically
- How to use other XRF-related commands
- How to use XRF in an IMS DBCTL environment
- How to use XRF in a data-sharing environment

- **Takeover:**

- How does automatic takeover differ from manual takeover
- How to work with the operator of the other subsystem and the operator of the network
- How to initiate a takeover
- How to ensure I/O prevention on the failed active subsystem
- When to end I/O toleration on the alternate subsystem
- How to switch DASD and class-3 terminals

- **Post-takeover:**

- How to ease the workload on the new active subsystem
- How and when to restore service to class-3 and class-2 terminals that failed to switch
- How to complete dynamic DL/I backouts
- How to perform diagnostics on the failed subsystem
- How to bring up a third subsystem as an alternate or to restore the original active as the alternate subsystem

- **Termination:**


- How to terminate an XRF complex
- How to stop surveillance

XRF requires the following resources:

- z/OS, IMS, and VTAM, running on each central processor complex (CPC).

- If USERVAR is used, Network Control Program (NCP) running on a 3725 Communication Controller for full terminal switching.
- IMS system logs— online data sets (OLDs) and write-ahead data sets (WADs)—and certain other system data sets must be on DASD shared by both IMS subsystems.

Related concepts:

 Overview of XRF (System Administration)

Planning for operating an XRF complex

When operating with XRF, it is important that the staff understand the “single-system image” and have an awareness of what physically belongs to active or alternate subsystems, what is shared, and generally where system and database data sets are located.

The operator preparation is mostly knowing the division of responsibility when processing moves to an a different processor as a result of a takeover.

Regardless of whether master terminal operations staff have master terminals close to one another, close communication is required.

There are several considerations when you plan changes to your operational control for XRF. These considerations are:

- Whether additional staff should be responsible for an alternate subsystem and assume full control when that subsystem becomes the active
- Whether a relocation of the IMS master terminal and z/OS alternate subsystem console is necessary
- How additional takeover activities and other procedure changes are to be incorporated into “run books”
- Whether a takeover should proceed automatically or by means of operator intervention when the alternate subsystem makes a takeover decision

Communication and console configuration among operators

To promote good communication among the operators of an XRF complex, place all consoles in a central location. The speed with which the takeover proceeds depends on operators confirming that a takeover is in progress and that I/O prevention is complete. This confirmation is easiest when operators are located in the same physical area.

With the operators in two separate physical locations, the chance for misunderstanding or lack of communication is great. The success of the takeover and the integrity of the databases lies with your operators and their awareness of what is going on in the entire complex.

The following figure shows a recommended configuration for the consoles that are responsible for managing XRF:

- A group of four logical consoles:
 - IMS master terminal for the active subsystem
 - IMS master terminal for the alternate subsystem
 - Multiple console support (MCS) console that receives messages from z/OS for the active subsystem
 - Multiple console support (MCS) console that receives messages from z/OS for the alternate subsystem

- IMS secondary master terminal with attached printer for the active subsystem
- IMS secondary master terminal with attached printer for the alternate subsystem
- Network terminal for VTAM and Tivoli® NetView® for z/OS
- Optionally, two additional consoles for JES3 in the two systems (not shown in the figure)

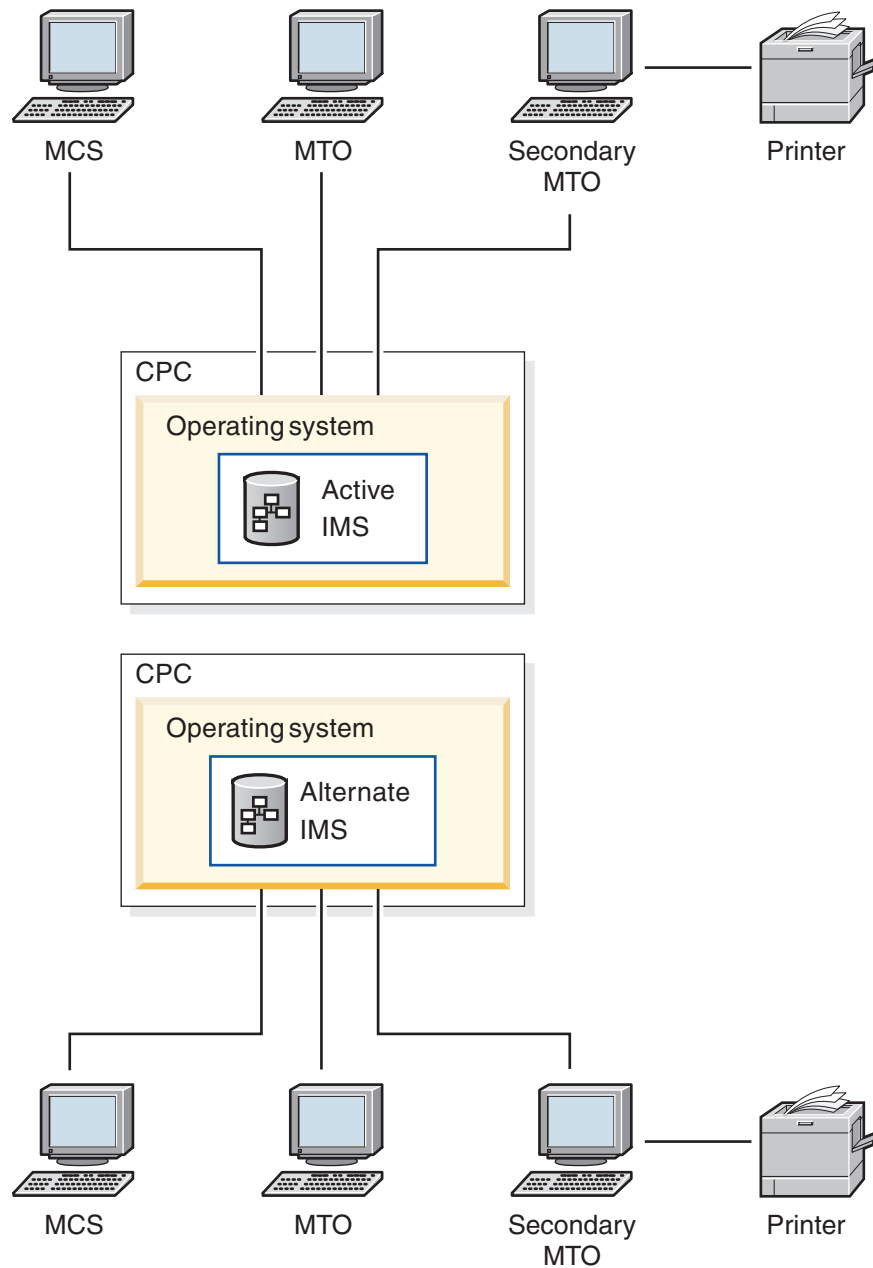


Figure 6. Recommended console configuration for XRF

It is important that the active and alternate IMS subsystems each have their own IMS master terminals. IMS limits the commands that the MTO can issue on the alternate subsystem.

On this set of terminals, your operators have a collective view of z/OS, IMS, and network events. Your operators cannot issue the z/OS SYSTEM RESET command on

the processor consoles, but they can receive all the availability manager (AVM) messages and can issue monitoring commands.

To understand how important it is to have the operators' consoles together, consider the implications of having separate centers for the active subsystem and the alternate subsystem operations. Operators at each center must have some understanding of the entire XRF complex. The state of the alternate subsystem is important to the operators for the active subsystem, and vice versa. For example:

- The operators for the active subsystem should know when the initialization of the alternate subsystem is complete.
- The operators for the active subsystem must know of a failure of the alternate subsystem.
- The operators for the alternate subsystem must know of the completion of I/O prevention in the failing active subsystem.

At changes of shifts, the operators on one shift must let the operators on the following shift know the current state of the complex.

Normal IMS operations in an XRF complex

During takeover, operators must be aware of response time, availability of terminals, data set access problems, and database access. You start the failed active subsystem as a new active subsystem in parallel with the alternate subsystem, providing interim end-user service.

In terms of command usage, the operator needs to know the following:

- When and how to initiate a takeover manually
- When and how to terminate an IMS alternate subsystem
- What terminal switching needs to be done
- What database control might be required
- How to interpret /DISPLAY output when operating in an XRF complex
- How to make changes to the takeover criteria and surveillance methods
- How to respond to a failure of the takeover process

Those terminals considered critical for continuous service should be defined as class-1 terminals. However, a certain number of these terminals might have to be switched manually. A convenient way of connecting these terminals to the new active IMS is to use an automated operator program. This program could issue a set of IMS /START or UPDATE commands. Another use of automated operator programs could be to modify surveillance intervals and takeover criteria.

Related concepts:

“Processing during takeover” on page 214

Online changes in an XRF complex

In addition to preparing the staging libraries, you need to prepare the inactive libraries.

Related concepts:

 Making local online changes in an XRF complex (System Administration)

XRF in a data sharing environment

If you are using IRLM in a sysplex data sharing environment and the alternate subsystem is available, there are no required tasks to maintain continued data availability.

XRF with a data propagation manager

To propagate changed data from IMS DL/I databases in an XRF complex to DB2 for z/OS tables, be sure IMS and DB2 for z/OS are available to both the active and alternate subsystems.

If DB2 for z/OS is not available on one of the systems, the data propagation from IMS to DB2 for z/OS fails until DB2 for z/OS is made available.

XRF initialization

During the initialization phase, in addition to standard IMS initialization tasks, the active subsystem and the alternate subsystems in an XRF complex perform tasks specifically related to XRF.

You need to know about the following:

- What IMS parameters are in effect for XRF.
- How to initialize MODSTAT, restart data set (RDS), and the message-queue data sets.
- Whether or not IMS joins the availability manager.
- How to successfully start the XRF complex.
- How to use VTAM USERVAR or VTAM Multinode Persistent Sessions (MNPS) to manage terminal sessions for XRF.

XRF parameter determination

The active subsystem uses the recoverable service element name (RSE name), not the z/OS SSID, when signing on to DBRC and identifying to IRLM.

Using the RSE name allows the alternate subsystem to inherit databases and their authorizations from the active subsystem during takeover.

Both subsystems read IMS.PROCLIB member DFSHSBxx for their RSEname, surveillance options, and takeover conditions.

Related reference:

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

Data set initialization

During initialization, the active subsystem can start but it is not XRF-capable if it does not find two MODSTAT data sets (unless global online change is enabled, in which case the MODSTAT data sets are not needed), two restart data sets (RDSs), and, if you used IMS Fast Path, four MSDBCPn data sets.

The alternate subsystem, however, cannot be started unless it finds these data sets on shared DASD.

Both the active subsystem and the alternate subsystem allocate the latest MODSTAT and RDS, but the alternate subsystem switches to the unused pair during startup. The active subsystem uses the last active MSDBCPn pair, and the alternate subsystem uses the other two.

In a shared-queues environment, the active IMS subsystem does not need the message-queue data sets, but the alternate subsystem IMS subsystem does use these data sets because the alternate subsystem does not participate in sharing the IMS message queues until it takes over the active IMS subsystem workload.

How to join IMS with the availability manager

The availability manager (AVM) has its own address space. You can start the AVM with the z/OS START command or let z/OS start it automatically when you issue the IMS /NRESTART or /ERESTART command. To start the availability manager, issue the z/OS START AVM command. When the AVM is initialized, you receive the AVM001I message.

Recommendation: Use the z/OS START AVM,SUB=MSTR command to start the AVM. If you omit the SUB=MSTR keywords, z/OS starts the AVM from the JES address space, and the resulting dependency on JES means you cannot stop JES without also stopping the AVM.

Both the active subsystem and the alternate subsystem notify the z/OS availability manager of their participation in an RSE during initialization. When IMS successfully joins the AVM and RSE, you receive messages DFS3871I and DFS3873I.

If IMS fails to connect to the AVM, you receive message DFS3872I. If the connection fails, the operator must initiate I/O prevention during XRF takeover.

You must have a procedure established in case the connection fails during initialization of either the active subsystem or the alternate subsystem. You might want to determine the cause of the failure and restart the initialization to ensure that the connection is successful.

How to start the subsystems in an XRF complex

Start the active subsystem first. Use the /ERE BACKUP command to start the alternate subsystem only after the active subsystem startup completes with the DFS994I message.

Recommendation: If the active subsystem begins processing before you start the alternate subsystem, issue the /CHECKPOINT command on the active subsystem to take a simple checkpoint. If you establish a checkpoint before starting the alternate subsystem, the alternate subsystem needs to read only the active subsystem's log from this checkpoint.

After an XRF takeover, restart the alternate subsystem first if it issued a simple checkpoint, even if you had not yet released it from I/O toleration (by issuing the /UNLOCK SYSTEM command or replying to an AVM message). Otherwise, restart the active subsystem first.

When you start an XRF alternate subsystem (using the /ERE BACKUP command), IMS establishes VTAM backup sessions (if USERVAR is used with XRF), an AVM node, and an ISC link to facilitate communications with the active subsystem. If the alternate subsystem fails and is restarted by the z/OS Automatic Resource Manager (ARM), the VTAM ACB and the VTAM master terminals will be opened towards the end of the alternate subsystem's initialization.

Startup procedures for subsystems:

You must have two distinct procedures to start an XRF complex because there are two separate IMS subsystems that perform different functions.

- One of these procedures must specify HSBID=1 in the IMS procedure and, in a non-shared-queues environment, contain DD statements for its set of unique message-queue data sets (LGMSG, SHMSG, QBLKS, LGMSG, SHMSG, and QBLKSL).

- The other procedure must specify HSBID=2 in the IMS procedure and contain DD statements for its set of message-queue data sets, which must be distinct from the data sets listed above.

In an XRF complex, the OLDS definition (OLDSDEF) for the active subsystem must match the OLDSDEF for the alternate subsystem. Also, all OLDSs used by the active subsystem must be accessible to the alternate subsystem, regardless of whether they are in the OLDS definition.

IMS identifier:

The IMS identifier (ID) identifies a particular IMS subsystem. When an IMS BMP job tries to link to an IMS subsystem, it attempts to link first to the subsystem specified by the **IMSID=** parameter, and next to any subsystem specified by the **ALTID=** parameter.

Specify the subsystem name in the IMSCTRL macro during system definition, and you can overwrite the subsystem name in JCL procedures.

If the active and alternate subsystems in an XRF complex are on the same central processor complex (CPC), their IMSIDs or ALTIDs must be different. If the subsystems are on different CPCs, you can specify either the same or different IMSIDs or ALTIDs.

If you specify different IMSIDs or ALTIDs, you must specify the correct IMSID or ALTID in all dependent region JCL procedures.

Related reference:

 Parameter descriptions for IMS procedures (System Definition)

Terminal session management for XRF

You can use VTAM USERVAR or VTAM Multinode Persistent Sessions (MNPS) to manage terminal sessions for XRF. The differences between each method are transparent to the end user.

In an XRF system with USERVAR, the alternate subsystem continuously maintains an open backup session for class-1 terminals and switches the terminals to that session during a takeover.

In an XRF system with MNPS, the alternate subsystem does not maintain backup sessions. Instead, the alternate subsystem opens a new instance of the MNPS ACB during a takeover. Sessions for class-1 terminals are then rerouted through the new MNPS ACB on the alternate IMS.

Recommendation: Because of performance limitations of MNPS, use the Communication Controller for Linux on System z® (CCL) to support XRF.

Related concepts:

 Using MNPS or USERVAR in XRF complexes (Communications and Connections)

Initializing the USERVAR to manage terminal sessions for XRF:

VTAM uses a variable called USERVAR to associate a user's logon message with the application name of the currently active IMS subsystem.

The two subsystems in an Extended Recovery Facility (XRF) complex have different application names.

During IMS initialization (and later, during XRF takeover), to allow terminals to log on to the correct subsystem, IMS or network operators must initialize the VTAM USERVAR tables.

If you are not using the VTAM USERVAR Management Enhancement in your network, the network operator must enter the `MODIFYprocname, USERVAR` command for the alternate subsystem VTAM if it owns terminals. You might also have to tell operators for other VTAMs that communicate with the XRF complex to enter the z/OS `MODIFY` command to change the application name in their USERVAR tables.

The active IMS subsystem issues a `MODIFY` command to tell its own VTAM to initialize the application name in its USERVAR table.

If you are using the VTAM USERVAR Management Enhancement in your network, the network operator should not initialize the VTAM USERVAR table nor use a Tivoli NetView for z/OS program (CLIST) to propagate the table to other VTAM nodes. It is also not necessary for the operator to enter the z/OS `MODIFYprocname, USERVAR` command to change the application name in the USERVAR tables. The VTAM USERVAR Management Enhancement does all of this for IMS.

Specifying MNPS to manage terminal sessions for XRF:

To have IMS use Multinode Persistent Sessions (MNPS) for Extended Recovery Facility (XRF) terminal switching, specify the `MNPS=name` optional keyword on the `DFSPBxxx` or `DFSHSBxx` `IMS.PROCLIB` members.

The *name* specifies the name of the MNPS ACB, which is used for terminal recovery. The MNPS ACB name must be the same for the active IMS and the alternate IMS, and only one MNPS ACB can exist at a time. If the `MNPS=` keyword is specified, the `USERVAR=` keyword is ignored.

Actions performed during synchronization

Actions that the operator may need to take during synchronization include issuing the `/CHECKPOINT SNAPQ` command, responding to inconsistent definitions, or responding to improperly allocated system data sets.

The alternate subsystem system automatically performs the following actions:

- Preopening backup sessions for class-1 terminals, if using USERVAR
- Pre-allocating and preopening DBDSs

Simple checkpoint request (/CHECKPOINT SNAPQ)

After you start the alternate subsystem with the `/ERESTART BACKUP` command, the active subsystem sends an “OKTOSYNC” message to the alternate subsystem over the ISC link when checkpoint processing is in progress. If you do not have an ISC link, the alternate subsystem monitors the restart data set (RDS) to determine when the checkpoint is taken.

After you start the alternate subsystem with the `/ERESTART BACKUP` command:

- If you do not have an intersystem communication (ISC) link, and you receive message DFS3801I on the alternate subsystem (saying that a checkpoint SNAPQ is required on the active subsystem):

1. Switch OLDSs to reduce contention on the OLDSs.
 2. Issue the /CHECKPOINT SNAPQ command on the active subsystem.
- If you do have an ISC link between the active subsystem and the alternate subsystems, and if the most recent checkpoint on the OLDS is not a SNAPQ, then the alternate subsystem sends the active subsystem a synchronization request to take a SNAPQ checkpoint.

If the alternate subsystem finds a SNAPQ (no matter how old it is) before finding any other checkpoint, it applies changes from the log beginning with that SNAPQ. Because this processing might stress your system, be careful when you start the alternate; you might not want to start it during peak processing hours.

If the most recent checkpoint is a SNAPQ, you can reduce processing during initialization by taking a simple checkpoint (using the /CHECKPOINT command) on the active subsystem before starting the alternate subsystem.

In a shared-queues environment, the /CHECKPOINT SNAPQ command synchronizes the active subsystem and alternate subsystems, but does not copy the message queues to the IMS log.

System definitions for XRF

The alternate subsystem uses the information on the OLDS from the most recent SNAPQ to ensure consistency of system definition for each IMS subsystem. If the definitions are inconsistent, the /ERESTART BACKUP command fails.

System definition for an Extended Recovery Facility (XRF) complex includes the following keywords, parameters, and statements:

- The following table shows the XRF definition in the system definition macros:

Table 26. Keywords, parameters, and statements for an XRF complex in the system definition macros

Required	Optional
HSB= (IMSCTRL macro)	BACKUP=(CTLUNIT, LINEGRP, MSLINK, MSPLINK, STATION, and TYPE macros)
APPLID= and PASSWORD= (COMM macro, if you use VTAM)	BACKUP= and ADDR= (LINE macro)
	BACKUP=, ADDR=, and NAME= (TERMINAL macro)
	BACKUP= (ETO logon descriptor)

- The following table shows the XRF definition in the IMS.PROCLIB member DFSHSBxx:

Table 27. Keywords, parameters, and statements for an XRF complex in the IMS.PROCLIB member

Required	Optional
RSENAME=. Specifies the RSE name, which must be the same for both the active and alternate subsystems.	ALARM=. Specifies whether the service processor alarm sounds for a takeover request.
	AUTO=. Specifies whether a takeover requires your approval.
	DEFERFIX=. Identifies the IMS.PROCLIB member DFSFIXnn.

Table 27. Keywords, parameters, and statements for an XRF complex in the IMS.PROCLIB member (continued)

Required	Optional
	USERVAR=. Specifies the USERVAR (if you use VTAM)
	SURV=. Specifies the surveillance mechanisms.
	LNK=. Specifies interval and timeout values for links.
	LOG=. Specifies interval and timeout values for the log.
	RDS=. Specifies interval and timeout values for the restart data set (RDS).
	SWITCH=. Controls whether absences of surveillance signals or failures of VTAM or of the IRLM cause the alternate subsystem to request or initiate a takeover.

The IMS.PROCLIB member DFSHSBxx controls the following aspects of the XRF complex:

- Surveillance mechanisms
- Takeover conditions
- Display of optional messages
- VTAM USERVAR name or MNPS ACB name (depending on whether USERVAR or MNPS is used)
- Name of the IMS.PROCLIB member that determines non-Fast Path control region page-fixing values for the alternate subsystem after a takeover
- Name by which the complex is known to z/OS, DBRC, and the IRLMs
- In the IMS.PROCLIB member DFSFIXnn:

This member specifies page-fixing options for the alternate subsystem after a takeover. You can move non-Fast Path related options of an existing DFSFIXnn member to this additional page-fix member to reduce the real storage requirements of the alternate subsystem during initialization and tracking. Use this member to fix buffer pool pages that you do not want fixed during initialization and tracking.
- In the IMS.PROCLIB member DFSVSMxx:


If the active and the alternate subsystems are on the same CPC, you must change the VSAMPLS keyword on the OPTIONS control statement to specify LOCL (local) instead of GLBL (global) for at least one of the subsystems. Only one control region on a processor can request a VSAM Shared Resource Pool. If you specify LS0=Y or LS0=S for the IMS control region, the default for VSAMPLS is LOCL. The OLDSDEF and WADSDEF control statements must be identical.
- In the EXEC statement for each control region:
 - HSBMBR=, which gives the suffix for the DFSHSBxx member of IMS.PROCLIB.
 - HSBID=, which identifies the message queues. Setting HSBID= to null deactivates XRF capability.

If you do not set HSBID= to null, you must specify either 1 or 2. The procedure for IMSA must specify a different value than you specify for IMSB.

If after a takeover, you bring up a third subsystem as an alternate subsystem, specify the same HSBID= as on the failing active subsystem.

- If the XRF subsystems participate in block-level data sharing, you need JCL statements for the IRLMs.
- DD statements for system data sets.

Related tasks:

 Guidelines for IMS system data set placement (System Administration)

System data sets for XRF

If the restart data set (RDS), MODSTAT , and MSDBCPn data sets are not properly allocated on the active and alternate subsystems, the alternate subsystem terminates. The alternate subsystem loads MSDBs from the SNAPQ records rather than from the MSDBCPn data set containing the last image.

Because IMS allocates data sets and pre-loads programs during initialization of dependent regions, you should exercise care if you make any changes in these areas after this time. Specifically:

- If programs running on the active subsystem modify data set allocation information, you might have to rewrite these programs to use dynamic allocation. Also note that IMS does not perform I/O prevention for any user data set.
- If you rebind preloaded programs on the active subsystem and subsequently activate them using the z/OS START command, you must also restart the associated regions on the alternate subsystem. The alternate subsystem does not automatically start or stop its dependent regions based on actions on the active subsystem.

Backup sessions for class-1 terminals

If USERVAR is used, the alternate subsystem establishes backup sessions for terminals that have the BACKUP=(, YES) option specified and that have open sessions (as recorded in the SNAPQ log records).

Tracking in the XRF complex

During normal operations, the operators do not have any tasks specifically for XRF, nor do they receive messages about XRF. During the tracking phase, your tasks are to monitor, maintain, and modify the XRF complex.

Because these tasks involve IMS commands that are also useful during other processing phases, these topics review all XRF-related IMS commands.

During the tracking phase, operators can issue the following commands to monitor the system:

- IMS /DISPLAY HSB command
- VTAM DISPLAY NET,ID and DISPLAY NET,USERVAR commands (if USERVAR is used)

These commands display the status of the XRF complex and monitor the status of sessions on class-1 terminals.

Operator commands on an XRF alternate subsystem

Because the XRF alternate subsystem does not process transactions, the operator for the alternate subsystem does not need some of the commands that the operator for the active subsystem needs.

IMS prevents the operator for the alternate subsystem from entering certain commands and displays one of the following messages if invalid commands are entered:

- DFS110I

Commands that are valid on an XRF alternate system, but which have keywords that are not valid when issued on that system, result in this message. For example, the command /START DATAGRP GROUP1 is not valid on an XRF alternate system. Although some forms of the /START command are allowed, the DATAGRP keyword is not allowed.

- DFS136I

Commands that are not allowed on an XRF alternate system result in this message.

- DFS3824I

Some commands contain a command verb and keywords that are each allowed on an XRF alternate system, but the syntax of the command is not allowed and results in this message. For example, the /START USER command is not allowed, even though the /START command verb is a valid command verb on an XRF alternate system, and the USER keyword is allowed for some commands on an XRF alternate system.

You can enter the following IMS commands on an alternate subsystem:

```
/ACTIVATE
/CANCEL
/CHANGE
/CLSDST
/COMPT
/DISPLAY
/END
/ERESTART
/FORMAT
/IDLE
/NRESTART
/OPNDST
/PSTOP
QUERY AREA
QUERY DB
QUERY TRAN
/RCLSDST
/RCOMPT
/RDISPLAY
/RMLIST
/RSTART
/SMCOPY
/TEST
/TRACE
UPDATE TRAN SET(MAXRGN)
UPDATE TRAN START|STOP(TRACE)
```

You can enter the following IMS commands on an alternate subsystem for VTAM master and secondary master terminals:

```
/RSTART NODE x  
/START NODE x  
/STOP NODE x
```

How to monitor an XRF complex

Monitoring an XRF complex requires that you know how to read the status line of the master terminal, and how to use the /DISPLAY command.

Reading the system status line

The top line of the screen of the IMS master terminal in an XRF complex gives the following information:

- The current date
- The current time
- The RSE name of the active and alternate subsystems
- The status of the subsystem:
 - ACTIVE for the active subsystem
 - BACKUP for the alternate subsystem
- The phase of the active and alternate subsystems
 - Before restart:
INITIALIZATION
 - After starting the active subsystem:
Null
 - After starting the alternate subsystem:
AWAITING SNAPQ
SYNCHRONIZATION
TRACKING IN PROGRESS
TAKEOVER REQUESTED
TAKEOVER IN PROGRESS
 - After XRF takeover:
AWAITING I/O PREVENTION (on the active subsystem)
I/O TOLERATION (on the alternate subsystem)
- The IMS subsystem identifier, such as IMSA

Displaying tracking and other XRF complex information

You can use the /DISPLAY command to display relevant tracking and other information for an XRF complex.

```
/DISPLAY STATUS DATABASE  
/DISPLAY DATABASE dbname BKERR
```

Use either one of these commands to display information about the following:

- Incomplete backouts that prevent use of a database
- EEQEs for a database

To complete backout, restart any database displayed. If the /START DB or UPDATE DB START(ACCESS) command to start the database fails because the

required OLDS have been archived and reused, use the date and time shown in the display to determine which SLDS to use and then perform batch backout.

/DISPLAY ACTIVE

Use this command on either the active or the alternate subsystem in an XRF complex to display information about active regions and, after the time stamp, the RSE name, the status of the subsystem (ACTIVE or BACKUP), and the phase of the XRF complex. The status can be one of the following:

AWAITING SNAPQ
SYNCHRONIZATION
TRACKING
TAKEOVER REQUESTED
TAKEOVER IN PROGRESS
I/O TOLERATION

/DISPLAY HSB

Use this command on either the active or the alternate subsystem to display:

- The RSE name
- The status of the subsystem (ACTIVE or BACKUP)
- The mode (whether the alternate subsystem is in I/O toleration after a takeover)
- The processing phase of the alternate subsystem (if entered on the alternate subsystem), which is one of the following:

AWAITING SNAPQ
SYNCHRONIZATION
TRACKING IN PROGRESS
TAKEOVER REQUESTED
TAKEOVER IN PROGRESS

- The IMSIDs of the active subsystem and the subsystem used to enter the command
- The VTAM USERVAR, if USERVAR is used
- The MNPS ACB name, if MNPS is used
- The time associated with the log record currently being read by the alternate system as it tracks the active subsystem
- The surveillance mechanisms in effect for the complex
- The takeover conditions specified on the SWITCH= keyword of the DFSHSBxx member of IMS.PROCLIB

Related concepts:

“Maintaining the XRF complex during tracking” on page 211

“Unplanned takeover” on page 214

Monitoring the status of the network

Some VTAM messages have a STATUS field that appears in response to the VTAM DISPLAY command. Examples in these topics describe the XRF-related keywords for the DISPLAY command and the STATUS fields that appear in the resulting messages.

The following graphic shows example system configuration for monitoring an XRF complex. The complex contains a primary and a backup session; IMSNAME1 is in

the primary session and IMSNAME2 is in the backup session.

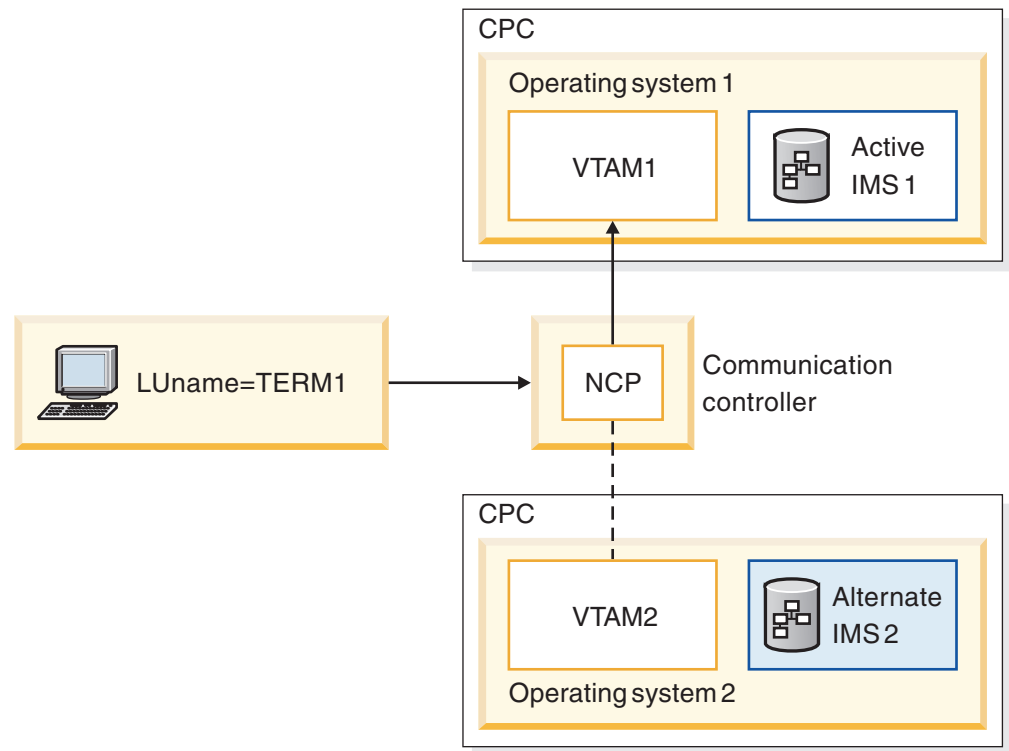


Figure 7. Monitoring status

Displaying the status of specific terminals:

The VTAM DISPLAY NET,ID command allows the network operator to see the status of the primary and backup sessions for a specific class-1 terminal.

For example, to display the status of both sessions for the class-1 terminal with logical unit name TERM1, the network operator issues the following command:

```
DISPLAY NET,ID=TERM1,SCOPE=ALL
```

The response to this command is the following message:

```
IST634I NAME      STATUS  SESSION ID SEND  RECV  VRN  TP  . . .
IST635I IMSNAME1  APRIM-PRI
IST635I IMSNAME2  ABKUP-PRI
IST835I END
```

Except for the STATUS field, the two IST635I messages should be familiar to your network operators. The NAME and STATUS fields in the two messages indicate the following:

- The following is the first IST635I message:
IST635I IMSNAME1 APRIM-PRI
In this message:
 - IMSNAME1 indicates that terminal TERM1 is in session with IMSNAME1.
 - A indicates that the session is active, that is, not pending.
 - PRIM indicates that the session with IMSNAME1 is the primary session.
 - PRI indicates that IMSNAME1 is a primary LU.
- The following is the second IST635I message:

```
IST635I  IMSNAME2  ABKUP-PRI
```

In this message:

- IMSNAME2 indicates that terminal TERM1 is in session with IMSNAME2.
- A indicates that the session is active, that is, not pending.
- BKUP indicates that the session with IMSNAME2 is the backup session.
- PRI indicates that IMSNAME1 is a primary LU.

Displaying the status of sessions connected to specific subsystems:

The VTAM DISPLAY NET,ID command allows the network operator to see the status of all sessions on class-1 terminals connected to a specific subsystem in the XRF complex.

For example, to see the status of all class-1 terminals that have sessions with the subsystem IMS1, the network operator issues the following command:

```
DISPLAY NET,ID=IMS1,SCOPE=ALL
```

The following message (IST635I) appears for every session that is communicating with IMS1.

```
IST635I  NAME  STATUS  SESSION  ID  SEND  RECV  VRN  TP  . . .  
IST635I  TERM1  APRIM-SEC
```

In this message:

TERM1 indicates that terminal TERM1 is in session with IMSNAME1.

A indicates that the session is active, that is, not pending.

PRIM indicates that the session is a primary XRF session. All sessions with IMSNAME1 will be primary XRF sessions

SEC indicates that TERM1 is a secondary LU.

Recommendation: If you have many class-1 terminals, do not display all sessions.

Displaying the number of primary and backup sessions in XRF with USERVAR:

The VTAM DISPLAY NET,ID command allows the network operator to see the number of primary and backup sessions on class-1 terminals communicating with XRF subsystems.

For example, to display the number of primary sessions in the XRF complex, the network operator issues the following command:

```
DISPLAY NET,ID=IMSNAME1,SCOPE=ALL
```

VTAM responds with this message:

```
IST809I  XRF SESSIONS - PRIMARY = nnnn  BACKUP = 0000.
```

To display the number of backup sessions in the XRF complex, the network operator issues the following command:

```
DISPLAY NET,ID=IMSNAME2,SCOPE=ALL
```

VTAM responds with this message:

```
IST809I  XRF SESSIONS - PRIMARY = 0000  BACKUP = mmmm.
```

Messages IST634I and IST635I follow with information about each session that is communicating with the IMS XRF complex.

An important use of this command is to determine if all the class-1 terminals have backup sessions established (if USERVAR is used). For example, before initiating a planned XRF takeover, an operator should make sure that all the users of class-1 terminals have backup sessions in place. In this case, the operator issues both DISPLAY commands discussed in these topics, compares the number of primary and backup sessions, and does not initiate a takeover unless the two numbers are identical.

Displaying the current application name in XRF with USERVAR:

The VTAM DISPLAY NET,USERVAR command allows the operator to see the application name that corresponds to a USERVAR in the USERVAR table.

For example, the operator issues the following command:

```
DISPLAY NET,USERVAR,ID=IMS
```

The response to this command is:

```
IST811I USERVAR(S):  
IST810I NAME = IMS, VALUE = IMSNAME1  
IST835I END
```

If your installation has defined more than one USERVAR, the network operator can issue the following command to display all the USERVARs and the corresponding application names:

```
DISPLAY NET,USERVAR
```

Maintaining the XRF complex during tracking

Maintaining an Extended Recovery Facility (XRF) complex requires that you know how to change surveillance dynamically, how to update system data sets, and how to suspend operations on the active or the alternate subsystem.

Although an end user cannot log on to the alternate subsystem during the tracking phase, the master terminal, secondary master terminal, optional ISC link between two subsystems, and the system console can communicate with the alternate.

You can enter the following commands on the alternate subsystem during tracking:

/DISPLAY

Use the /DISPLAY command to display relevant tracking and other information for an XRF complex.

/OPNDST NODE

/CLSDST NODE

Use these commands to open and close VTAM sessions for the master terminal and the ISC link. These commands can reestablish backup sessions for class-1 terminals that have active sessions, if USERVAR is used. Backup sessions are not maintained in XRF systems with MNPS.

If the active IMS has a session that is XRF capable, enter /OPNDST command on the alternate subsystem to establish a backup session, if it does not already have one.

In order to use the /OPNDST command successfully for a class-1 ETO terminal, the control blocks must exist and be coupled together, and an active session must exist on the active system. If not, IMS rejects the /OPNDST command, and issues message DFS3862.

/START LINE x PTERM y
/STOP LINE x PTERM y
/RSTART LINE x PTERM y

Use these commands to start, stop, and restart non-VTAM master and secondary master terminals.

/START NODE x
/STOP NODE x
/RSTART NODE x

Use these commands to start, stop, and restart VTAM master and secondary master terminals.

/START REGION
/STOP REGION

Use these commands to start and stop regions.

/START SURV
/STOP SURV
/CHANGE SURV

Use these commands to dynamically stop and start surveillance and to dynamically change surveillance values.

Related reference:

“How to monitor an XRF complex” on page 207

Controlling surveillance:

You can issue commands to start and stop surveillance for ISC link, the system log, and the restart data set. You can also change the mechanism, interval value, or timeout value that you established during initialization.

Use the following command to dynamically stop and start surveillance, and dynamically change surveillance values:

/START SURVEILLANCE LNK | LOG | RDS | ALL
/STOP SURVEILLANCE LNK | LOG | RDS | ALL

Use these commands to start and stop the operation of any (or all) of the three surveillance mechanisms:

- The ISC link
- The system log
- The restart data set

If you operate with all three surveillance mechanisms, IMS can detect a problem as soon as possible and is more likely to initiate a takeover only when it is appropriate. However, you can stop and start surveillance for a brief period to perform a system change, such as to restart a DASD control unit, that might trigger an unwanted takeover.

/CHANGE SURVEILLANCE [LNK | RDS | LOG | ALL] TIMEOUT | INTERVAL

Use this command to change the mechanism, interval value, or timeout value that you established during initialization. If you enter this command on the active subsystem, it affects both the active and the alternate subsystems; if you enter it on the alternate subsystem, it affects only the alternate subsystem.

Recommendation: Enter /CHANGE SURVEILLANCE on the active subsystem. However, if you enter it on the alternate subsystem, do not specify an interval or timeout value for the alternate subsystem that is less than the interval or timeout value on the active.

Updating system data sets during tracking:

Before you make a system change to an XRF-capable IMS subsystem, you must identically update the equivalent data set on the other subsystem in the XRF complex.

Update the following system data sets:

- IMS.ACBLIBA and IMS.ACBLIBB
- IMS.FORMATA and IMS.FORMATB
- IMS.MODBLKA and IMS.MODBLKB
- IMS.TFORMATA and IMS.TFORMATB

You can perform identical updates on both subsystems using an IMS restart or the Online Change Copy utility (DFSUOCU0). Use one set of staging libraries for both subsystems; a staging library is not a single point of failure because the online subsystem does not use it. After you have ensured that the inactive data sets on both subsystems are identical in content, use the /MODIFY PREPARE, /DISPLAY MODIFY, and /MODIFY COMMIT commands on the active subsystem. After these changes are committed on the active subsystem, the alternate subsystem automatically reprocesses the online change sequence.

Suspending operations on the active or alternate subsystem:

Do not suspend operations on or stop the z/OS system on which the active subsystem is running without first disabling surveillance, unless you want to initiate an XRF takeover.

Recommendation: Do not suspend operations on or stop the z/OS system on which the alternate subsystem is running without first disabling surveillance and stopping the active subsystem. If you do, the alternate subsystem could interfere with the performance of the active subsystem when you restart the alternate subsystem because it tries to reflect the changes that occurred on the active subsystem in the meantime.

IMS control region dump during tracking:

To prevent an unwanted takeover while processing an SDUMP, ensure that the XRF takeover interval is sufficient to allow the SDUMP to complete, and ensure that the system data set contains adequate space for capturing information for the dump.

If you specify QUIESCE=YES on the z/OS CHANGEDUMP command, any SVC dump, depending on the timeout values assigned to the LNK and RDS keywords in DFSHSBxx, might trigger a takeover request. If you specify QUIESCE=NO, only a dump of the IMS control region can trigger a takeover. If you are troubled by SVC dumps that cause unwanted takeovers, stop surveillance or specify QUIESCE=NO.

Recommendation: To prevent an unwanted takeover while processing an SDUMP:

How commands on the active subsystem affect the alternate subsystem:

Commands that make resources unavailable to the active subsystem also make these resources unavailable to the alternate subsystem after a takeover. This includes, for example, a /STOP PROGRAM command to stop the scheduling of transactions and any command to reassign terminals or stop physical terminals.

Processing during takeover

Most of the processing during takeover takes place in the alternate subsystem. The alternate subsystem detects the normal or abnormal termination of the active subsystem and takes control of its resources and workload.

Related concepts:

“Normal IMS operations in an XRF complex” on page 198

Planned takeover

You can apply hardware or software maintenance to either subsystem in an XRF complex as long as you do not change the log record format, make system definition changes (and bring up an alternate subsystem with those changes), or apply maintenance that requires an IMS cold start in a non-XRF environment.

To apply maintenance to the alternate subsystem: terminate it, leaving the active subsystem without an alternate subsystem for the time that you take to apply the maintenance. Then bring that subsystem back up as the alternate subsystem.

To apply maintenance to the active subsystem: initiate a planned takeover to shift its workload to the alternate subsystem. After you apply maintenance, bring up that subsystem as a new alternate subsystem. Initiate another takeover to return the subsystem to its role as the active subsystem.

Unplanned takeover

The alternate subsystem can take over the workload of the active subsystem without operator intervention.

In an XRF complex, the alternate subsystem uses log records and surveillance to detect events on the active subsystem that suggest the need for a takeover:

- It reads the log for evidence of an IMSabend.
- It reads the log for evidence of an IRLM failure.
- It detects an absence of signals across the ISC link, absence of new records on the OLDS, or absence of a new time stamp on the restart data set (RDS).

You can display the surveillance mechanisms and takeover conditions in effect by using the /DISPLAY HSB command.

The following conditions can trigger a takeover request:

- A surveillance-detectable failure of the IMS control region
- A total failure of z/OS
- A single central processor complex (CPC) failure
- An IRLM failure that causes IMS to call the IMS STATUS exit routine
- A VTAM failure that causes IMS to call the IMS TPEND exit routine

If you set the **SWITCH=** parameter of the DFSHSBxx member in IMS.PROCLIB to SWITCH=(LNK,LOG,RDS),(IRLM),(TPEND) (the default), the alternate subsystem requests or initiates a takeover for any one of the following events:

- Both the ISC link and the RDS fail to send signals, and new log records fail to appear after the timeout periods that you specify.
- The IRLM fails and calls the STATUS exit routine.
- VTAM fails and causes IMS to call the TPEND exit routine.

XRF does not take over for the following types of failures:

- A channel or link failure that causes a break in communication between the CPC and the communication controllers
- A failure in the telecommunication network, such as communication controllers, the network control program (NCP), telecommunication lines, or terminals
- An intersystem failure, such as those caused by JES3 or CTC links
- Loss of or damage to an IMS database
- A power failure that affects both CPCs in the XRF complex
- A failure of user catalogs that point to data sets, such as databases

You can establish your own procedures to detect more subtle causes of failures of the active to process transactions. If a failure situation occurs that the alternate subsystem does not detect, you can force a takeover by entering the /SWITCH SYSTEM command from either the active subsystem or the alternate subsystem.

Related reference:

“How to monitor an XRF complex” on page 207

The takeover flow

You can determine if you want a takeover to proceed automatically or for it to proceed only with your approval.

If you want a takeover to proceed automatically, specify AUTO=YES in the DFSHSBxx member of IMS.PROCLIB. When IMS detects a takeover condition, it does not give you the option of deciding whether a takeover is needed. You receive message DFS3890I saying that a takeover was requested, and the alternate subsystem takes over the workload from the active subsystem.

If you want a takeover to proceed only with your approval, specify AUTO=NO (the default) in the DFSHSBxx member of IMS.PROCLIB. When IMS detects a takeover condition, you receive message DFS3869 saying that the active subsystem might have failed. You must decide whether to proceed with a takeover. To initiate a takeover, issue the /SWITCH SYSTEM ACTIVE [FORCE] command from the active subsystem if you can. Otherwise, issue the /SWITCH SYSTEM FORCE command from the alternate subsystem.

Good communication between operators for the active and alternate subsystems is essential.

How to ensure I/O prevention:

After an XRF takeover, you must ensure that the failing active subsystem cannot update data sets used by the new active subsystem. If the new active subsystem can write to the databases before the failing active subsystem is prevented from changing this data, data integrity is compromised.

That is, you must ensure that the new active subsystem is in a state known as *I/O prevention*. I/O prevention can occur in one of the following ways, depending on the cause of the takeover:

- I/O prevention occurs automatically:
 - If IMS fails and z/OS continues running
 - If you enter the /SWITCH SYSTEM command on the active subsystem and IMS issues abend U0604
- If z/OS or the CPC fails, ensure I/O prevention by resetting the active CPC.
- For all other causes of takeover, terminate the active IMS subsystem with the z/OS MODIFY command.

Because there is no communication between the z/OS for the active subsystem and the z/OS for the alternate subsystem, the operator for the active subsystem must inform the operator for the alternate subsystem when I/O prevention is complete.

I/O prevention is completed in one of the following ways:

- By receiving the following message from the z/OS availability manager:


```
AVM006E  TELL OPERATOR AT BACKUP TO REPLY          "UNLOCK" TO MESSAGE AVM005A.
          I/O PREVENTION IS COMPLETE FOR
          SUBSYSTEM ssid,
          FAILING ACTIVE ELEMENT OF RSE rsename.
```
- By ensuring that the failed IMS control region and DL/I address space have terminated execution
- By ensuring that the CPC has failed
- By ensuring that the CPC has been reset
- By canceling IMS and waiting until all of the address spaces terminate
- By physically switching the DASD that contains the data sets that the active subsystem might access

If the availability manager's address space fails, IMS continues processing normally but takeover processing is degraded. You can choose to continue without the availability manager, however, you cannot reconnect IMS to the availability manager unless you terminate and restart IMS.

How to end I/O toleration:

Until the failed active subsystem has completed I/O prevention, the new active subsystem is in a state known as *I/O toleration*. You must be sure that I/O prevention is complete before you end I/O toleration. No message appears on the new active subsystem console to tell you that I/O prevention is complete.

The only way you can learn that I/O prevention is complete is from the operator of the failed active.

Recommendation: Because the new active subsystem is executing transactions while the failed active subsystem is in I/O prevention, it is better to wait until you are sure that the IMS control region and the DL/I address space on the failed active subsystem are fully terminated than to get out of I/O toleration prematurely. You can end I/O toleration in one of the following ways:

- By responding UNLOCK to the following message:


```
AVM005A  REPLY UNLOCK WHEN I/O PREVENTION
          COMPLETES FOR RSE rsename
```
- By entering the /UNLOCK SYSTEM command

Manual control at takeover

When IMS issues message DFS3890I to initiate the takeover request, the IMS operators must have specific procedures to follow. The IMS operator for the active subsystem first verifies that a takeover is necessary.

The message might indicate a condition that does not call for a takeover, such as a link failure or the active subsystem starting a dump. The message appears every 30 seconds as long as surveillance or the system log identifies the error condition to the alternate subsystem. If the active subsystem resumes normal processing, the message no longer appears.

You give an operator control during an extended recovery facility (XRF) takeover for one or more of the following reasons:

- To perform certain tasks
- To make the takeover smoother
- To make sure the takeover is necessary
- To try to correct the problem that caused the takeover request

If a takeover is necessary, before allowing the takeover to proceed, an IMS operator might:

- Switch telecommunications lines
- Suspend operations for certain workloads on the alternate subsystem, such as stopping a batch job
- Force TSO terminal users off the alternate subsystem
- Switch a 3274 terminal controller to allow recovery on local terminals
- Switch the databases connected to the active subsystem to the alternate subsystem

To attempt to fix the problem that caused the takeover request, the IMS operator for the failing active subsystem might:

- Restart the CPC
- Terminate the active IMS subsystem
- Respond to a missing interrupt handler (MIH) message on one of the z/OS consoles

To proceed with the takeover, the operator for the active subsystem issues the /SWITCH SYSTEM ACTIVE command from the IMS console. This command causes the active subsystem to terminate abnormally with a U0604 abend.

The operator for the alternate subsystem can request that the takeover proceed by issuing the /SWITCH SYSTEM FORCE command. The operator then terminates IMS in the active subsystem. When address spaces in the active subsystem are all terminated, the operator for the active subsystem can assume that I/O prevention is complete.

Recommendation: The operator for the alternate subsystem should initiate a takeover only when the active subsystem cannot receive commands.

At this point, procedures for the operators are the same as if the takeover proceeded automatically rather than manually.

Failure during takeover

If a failure of the alternate subsystem occurs during takeover (regardless of whether it is sharing data), use the /ERE to restart the subsystem that took the latest checkpoint valid for restart.

Network and user considerations

Because the purpose of an XRF complex is to diminish the disruption from failures and planned changes to end users, you must be aware of how an XRF takeover alters your network and terminal configuration.

XRF support of terminals:

XRF provides different service after takeover for the three different classes of terminals, depending on their characteristics.

The support that XRF provides to a particular terminal depends on several factors:

- Whether the terminal uses the SNA protocol
- Whether the terminal uses the logical unit type 6.2 (LU 6.2) protocol
- How you define the terminal on the UNITYPE keyword of the TYPE or LINEGRP macros during system definition or ETO logon description
- What priority you establish for the terminal on the BACKUP= keyword of a system definition macro or ETO logon descriptor
- Whether the terminal is controlled by a VTAM instance and an NCP that support XRF, if USERVAR is used
- Whether the terminal is connected to a 37x5 Communication Controller
- Whether the terminal is channel-attached to the active subsystem or the alternate subsystem
- Whether you specified BACKUP=(, YES) keyword for class-1 terminals or BACKUP=(, NO) keyword for class-2 terminals
- Whether the VTAM instance that owns the terminal is in the z/OS of the active subsystem, the z/OS of the alternate subsystem, or in a Communication Management Configuration (CMC)

Related concepts:

 Terminals in an XRF complex (System Administration)

XRF terminal class support:

During takeover, an XRF-capable IMS subsystem manages its three terminal classes in different ways.

Class-1 terminals

Class-1 terminals are managed differently depending on whether XRF with USERVAR or XRF with MNPS is used.

If XRF with USERVAR is used, IMS asks the Network Control Program (NCP) to switch VTAM sessions for class-1 terminals from their active sessions to their pre-initialized backup sessions and to terminate their active sessions.

If XRF with MNPS is used, VTAM holds the class-1 terminals' sessions open as persistent sessions when XRF takeover begins. The original MNPS ACB closes and the alternate IMS opens its own MNPS ACB. VTAM then reconnects the class-1 terminals' sessions to the new MNPS ACB on the new active IMS.

Class-2 and class-3 terminals

If XRF with USERVAR or XRF with MNPS is used, class-2 and class-3 terminals are managed as follows:

- IMS tries to reestablish sessions for class-2 terminals.
- IMS terminates sessions for class-3 terminals, as it would if XRF were not present. Logical unit type 6.2 (LU 6.2) devices are class 3.

Sessions for class-2 terminals terminate at the time of failure, but the alternate subsystem tries to reestablish sessions if:

- Paths between the alternate subsystem and the terminals exist.
Paths exist for channel-attached terminals connected to other hosts in the network. Paths do not exist for channel-attached terminals connected to the failing active.
- The alternate subsystem can find the host that owns the terminals.
It does this by specific definition in the Cross-Domain Resource Definitions or by using the VTAM facilities that try to locate the owning system services control program (SSCP) by referring to a default list of SSCPs.
- The owning SSCP is available.
The alternate subsystem cannot reestablish a session with a class-2 VTAM terminal if the failing active subsystem owns it.

XRF and TSO, batch, CICS, or another IMS:

During XRF takeover, z/OS might swap out swappable work, such as TSO and batch applications, until the new active subsystem is stable.

When z/OS swaps the work back in, it might run with degraded performance, depending on the demands of the IMS workload.

If you want to cancel these jobs on the new active subsystem, wait until takeover is complete, and cancel one job at a time. Processing for the CANCEL command takes priority over takeover processing, and could slow takeover.

Nonswappable work during takeover, such as CICS or another IMS, competes with takeover processing, and experiences some degree of degraded performance. z/OS tries to give IMS in the new active subsystem priority during takeover, but the nonswappable work might still interfere and prolong the takeover.

Nonswappable work after takeover competes with the IMS interactive workload. Response times for both sets of users might be affected, depending on how heavily system resources are used, and their relative priorities.

Change the USERVAR:

During a takeover, as during initialization, to allow terminals to log on to the currently active subsystem, IMS and network operators must initialize the VTAM USERVAR tables.

Changing the USERVAR follows this pattern:

- The alternate subsystem automatically issues a MODIFY USERVAR command to tell its own VTAM instance to initialize the application name in its USERVAR table.

- If you are using VTAM Release 3.1 or an earlier release, the network operator must enter a `MODIFY procname,USERVAR` command for any other VTAM instance that owns terminals that log on to the currently active subsystem in the XRF complex.

This might mean that you must tell operators for VTAM instances that communicate with the XRF complex to enter the `MODIFY procname,USERVAR` command to change the application name in their USERVAR tables.

If you use Tivoli NetView for z/OS, a command list (CLIST) is available that automatically propagates this change in application name to other VTAM nodes.

Dynamic DL/I backouts

During takeover, the alternate subsystem in an XRF complex backs out transactions that were in process when takeover was initiated. The alternate subsystem records information necessary for the backout and notifies DBRC. DBRC flags the DBDS record with “Backout required”. You can retry the backout as long as the OLDS containing the pertinent information is available.

For DEDBs on shared DASD, the alternate subsystem performs forward recovery. For DL/I databases on shared DASD, the alternate subsystem reads the database block to which the change log record refers, and then it backs out the change. If I/O toleration has completed, it writes the block back. If I/O toleration is still active, the alternate subsystem creates EEQEs for these blocks and places them in virtual buffers residing above the 16-megabyte line.

Transactions (running on the new active IMS) that access these blocks prior to the end of I/O toleration use these virtual buffers instead of the blocks on DASD. DBRC records the presence of the EEQEs and the IRLM propagates this information to all sharing subsystems. Thus, until I/O toleration ends, these databases can continue to be shared, but the presence of the EEQE prevents a sharing subsystem from accessing these particular blocks.

For DL/I databases on nonshared DASD, the alternate subsystem allows you to defer backout until you switch the DASD and restart the databases. The alternate subsystem schedules the appropriate PSBs for backout. It backs out all databases for each PSB that can access them and stops only the databases that it cannot access.

Post-takeover

After you receive message DFS994I (takeover is complete) on the alternate subsystem, you can finish other tasks related to the takeover.

Easing the workload after takeover

After you receive message DFS994I on the alternate subsystem after takeover, you might want to ease the workload on the new active subsystem.

You can perform the following tasks to help the new active subsystem adapt to its changed workload:

- Stop batch or TSO work that had been running on the alternate subsystem before the takeover.
- Ease the workload of the new active CPC to improve its performance for online users by stopping certain transactions or applications.
- Inform users of restricted operations (if any).

Recovering terminal sessions

In an XRF system with USERVAR, the alternate subsystem continuously maintains an open backup session for class-1 terminals and switches the terminals to that session during a takeover. If the alternate subsystem fails to reestablish a session with a class-2 terminal, use existing procedures to correct the cause of the failure and reestablish service. Use existing procedures to establish new sessions for class-3 terminals.

If class-1 terminals switch to the alternate subsystem (the new active subsystem), you receive the following messages for them on the master terminal or system console: DFS3859I and DFS3860I saying that all terminal sessions have been switched. If you do not receive these messages, identify the unswitched terminals and issue /CLSDST and /OPNDST commands for them.

If you need to switch many terminals manually, you can use an automated operator program. This program can invoke a set of /DISPLAY NODE, /CLSDST NODE, and /OPNDST NODE commands.

Related concepts:

Chapter 9, “Automated operations,” on page 329

Completing dynamic backouts and starting databases

Use the /DISPLAY DB BKERR command to determine whether dynamic DL/I backout has failed for any database on nonshared DASD. If it has, switch the DASD to make the database available to the alternate subsystem (using, for example, an IBM 3814 switching unit), and issue a /START DB or UPDATE DB START(ACCESS) command for the affected database to perform the incomplete backout. For shared databases, include the GLOBAL keyword on the /START or route the UPDATE command to all IMS systems in the IMSplex.

The /DISPLAY DB BKERR command does not show databases and VSAM KSDSs stopped because of extents by the failing active subsystem in progress during takeover. The following messages inform you of these: DFS2012I, DFS0772I, and DFS0773I. These messages list stopped databases, stopped shared databases, and KSDSs that were not backed out.

Notifying other VTAMs of the takeover in XRF with USERVAR

During takeover, if Extended Recovery Facility (XRF) with USERVAR is used, the alternate subsystem issues a MODIFY USERVAR command to change the VTAM application name in its USERVAR table. If you are using VTAM Release 3.1 or an earlier release, network operators in all VTAM domains except the one for the alternate subsystem, including the failed active, must change these entries. The network operators for each of these VTAM instances issue the same command.

When an entry in the USERVAR table changes, VTAM responds with this message:

```
IST813I USERVAR IMS CHANGED FROM name1 TO name2
```

Until the USERVAR is changed for these VTAM nodes, terminals owned by those VTAM instances cannot log on to the new active subsystem.

You might want your operator to delay issuing this command after a takeover until processing on the new active subsystem stabilizes. In fact, if the processing environment for the new active subsystem is inferior to that at the failed active subsystem, you might not want any more terminals to connect to the new active subsystem.

If your installation has installed Tivoli NetView for z/OS, you can use Tivoli NetView for z/OS to communicate the new application name.

Performing diagnostics after unplanned takeovers

If the takeover was unplanned, use existing procedures to determine the cause of the takeover on the failed active subsystem.

You might, for example:

- Request service aid information using a z/OS MODIFY IMS,DUMP command.
- Perform problem determination.
- Identify the component that caused the subsystem to fail.
- If the subsystem is dumping, let it continue.

Establishing a new alternate subsystem

Following a takeover, the new active subsystem runs without an alternate subsystem. You must establish a new alternate subsystem, and you can reestablish your original configuration.

Establish a new alternate subsystem by doing one of the following:

- Establish a new alternate subsystem in the CPC that had the failed active subsystem.
- Establish a new alternate subsystem in a different CPC without waiting for problem determination for the failed active subsystem to complete. Given a suitable configuration, this is the faster way to regain XRF capability.

Restriction: If the XRF complex is sharing databases with another system, do not establish an alternate on a different CPC.

If you want to reestablish your original configuration (active and alternate subsystems in their original CPCs), do so in stages:

1. Bring up the former active subsystem as the new alternate subsystem.
2. Perform a planned takeover.

Terminating subsystems

The termination phase separates the two subsystems in the XRF complex into independent IMS environments

To terminate the alternate subsystem and cancel backup sessions for class-1 terminals, issue the /STOP BACKUP command from the alternate subsystem.

To terminate both the active and the alternate subsystems, issue the enter the /CHECKPOINT FREEZE command from the active subsystem. When the active subsystem terminates, it writes a type X'06' record to the IMS log. The alternate subsystem reads this record and terminates

IMS DBCTL capabilities

In an IMS DBCTL environment, you can back up your system by using a second DBCTL subsystem as a standby alternate subsystem, or using XRF.

DBCTL standby

While one DBCTL subsystem is active, you can start another DBCTL subsystem as a standby alternate. This alternate subsystem does not track the active subsystem (as it would in a true XRF complex), but is a fully initialized IMS DBCTL

subsystem that is waiting for a restart command (for example, the /ERE command). This alternate subsystem is a *preinitialized DBCTL environment*.

The preinitialized DBCTL and active subsystem startup procedures must specify the DBRSE= keyword with the same string so that after restart the preinitialized DBCTL environment can assume the identity of the failed active. In addition, the active and preinitialized DBCTL subsystems must both specify unique IMSID= (or ALTID=) and CRC= parameters. Otherwise, the second region to start will fail during initialization.

Until you issue the restart command for the preinitialized IMS DBCTL subsystem, it cannot process any other commands. After restart is complete, everything is the same as if you had restarted the failed active system. Issuing a restart command for the preinitialized DBCTL subsystem before the active subsystem has shut down causes the restart command to fail and could cause a failure of the preinitialized DBCTL environment.

DBCTL and XRF

An IMS DB/DC XRF subsystem can provide DBCTL service. There are no changes to the XRF functions or to the DBCTL functions in this case.

For an XRF takeover to complete in this environment, the CCTL must be able to broadcast a /SWITCH command to the XRF active subsystem. The CRC must be unique to each IMS region that specifies it.

Recovery points

If you use either full-function or Fast Path DEDB databases, you can create recovery points in the database without having to take IMS offline. By creating recovery points you will improve the recovery of the database when an outage occurs.

When you quiesce the database you establish a point of consistency, which is when there are no updates pending for the database and the information stored on a direct access storage device (DASD) accurately reflects the current information stored in the database. When this point of consistency is reached, you create an image copy that you can use to provide a quick method of recovering the database to this point.

Important: This process must occur on every IMS in an IMSplex that is actively using the database.

You can create this point of consistency without making the database unavailable to applications. When the quiesce command is invoked, all updates that are in progress are committed and stored on DASD. However, the next DL/I call that attempts to access the database results in a wait until the quiesce on the database is released. Once the quiesce is released, these applications can access the database again.

There are two ways to quiesce the database using the UPDATE DB command with the START(QUIESCE) keyword:

START(QUIESCE)

This command causes the database to become quiesced and then releases the quiesce, thereby creating the new recovery point for the database. This is the default form of the quiesce command.

START(QUIESCE) OPTION(HOLD)

This command causes the database to become quiesced and keeps the database quiesced until an UPDATE DB STOP(QUIESCE) command is issued that releases the quiesce. While the database is quiesced, you can create image copies of the database data sets.

Creating recovery points

You use the IMS single point of control (SPOC) OM API to enter the UPDATE DB command with the START(QUIESCE) keyword to establish a recovery point. You can also enter an UPDATE DB START(QUIESCE) OPTION(HOLD) command to have a recovery point established, and keep the database quiesced.

The UPDATE START(QUIESCE) command stops all updates in progress and saves them to a direct access storage device (DASD). When this command is issued, all open ALLOC records on any sharing systems are closed. When the command completes, you have a recovery point, and applications can continue to make updates to the database.

To create recovery points for high-activity databases:

1. Enter one of the following commands from IMS SPOC:

Option	Description
UPDATE AREA NAME(<i>area_name</i>) START(QUIESCE)	To quiesce a Fast Path area and establish a new recovery point.
UPDATE DB NAME(<i>db_name</i>) START(QUIESCE)	To quiesce a full-function database and establish a new recovery point.
UPDATE DB NAME(<i>HALDBmaster_name</i>) START(QUIESCE)	To quiesce a full-function HALDB. Note: When the quiesce function is invoked on a full-function HALDB, it affects all of the partitions of the HALDB. The command does not set on the quiesce in progress flag on in the RECON data sets for the HALDB master. Instead it sets the quiesce in progress flag on in each HALDB partition.
UPDATE DB NAME(<i>DEDB_name</i>) START(QUIESCE)	To quiesce a Fast Path DEDB and establish a new recovery point.
UPDATE DATAGRP NAME(<i>dbgrp_name</i>) START(QUIESCE)	To quiesce a database group and establish a new recovery point.

IMS SPOC returns the following output:

DBName	MbrName	CC
DBXYZ	IM02	0
DBXYZ	IM01	0
DBXYZ	IM03	0

2. Verify that the command has processed successfully. Enter the following command from IMS SPOC: QRY DB NAME(DBXYZ) SHOW(STATUS)

IMS SPOC returns the following output:

DBName	MbrName	CC	TYPE	LclStat
DBXYZ	IM03	0	DLI	ALLOCS,OPEN
DBXYZ	IM01	0	DLI	ALLOCS,OPEN
DBXYZ	IM02	0	DLI	ALLOCS,OPEN

3. Create a concurrent image copy of the database. Verify that the image copy completes successfully.

Creating recovery points and keeping the database quiesced

The UPDATE DB START(QUIESCE) OPTION(HOLD) command causes the database to become quiesced and keeps the database quiesced until an UPDATE DB STOP(QUIESCE) command is issued that releases the quiesce. While the database is quiesced, you can create image copies of the database data sets.

To create recovery points for low-activity databases:

1. Enter one of the following commands from IMS SPOC:

Option	Description
UPDATE AREA NAME(<i>area_name</i>) START(QUIESCE) OPTION(HOLD)	To quiesce a Fast Path area and establish a new recovery point and keep the area quiesced.
UPDATE DB NAME(<i>db_name</i>) START(QUIESCE) OPTION(HOLD)	To quiesce a full-function database and establish a new recovery point and keep the database quiesced.
UPDATE DB NAME(<i>HALDBmaster_name</i>) START(QUIESCE) OPTION(HOLD)	To quiesce a HALDB and establish a new recovery point and keep the HALDB quiesced. Note: When the quiesce function is invoked on a full-function HALDB, it affects all of the partitions of the HALDB. The command does not set the quiesce in progress flag on in the RECON data sets for the HALDB master. Instead it sets the quiesce in progress flag on in each HALDB partition.
UPDATE DB NAME(<i>DEDB_name</i>) START(QUIESCE) OPTION(HOLD)	To quiesce a Fast Path DEDB and establish a new recovery point and keep the database quiesced.
UPDATE DATAGRP NAME(<i>dbgrp_name</i>) START(QUIESCE) OPTION(HOLD)	To quiesce a database group and establish a new recovery point and keep the database group quiesced.

The following IMS SPOC output is returned:

```
DBName  MbrName  CC
DBXYZ   IM02    0
DBXYZ   IM01    0
DBXYZ   IM03    0
```

2. Check that the command executed successfully. Enter the following command from IMS SPOC: QRY DB NAME(DBXYZ) SHOW(STATUS) IMS returns the following results:

```
IMS SPOC output:
DBName  MbrName  CC TYPE  LclStat
DBXYZ   IM03    0 DLI    ALLOCS,OPEN,QUIESCED
DBXYZ   IM01    0 DLI    ALLOCS,OPEN,QUIESCED
DBXYZ   IM02    0 DLI    ALLOCS,OPEN,QUIESCED
```

3. Create an image copy of the database. Verify that the image copy completes successfully. The database should be successfully backed up.
4. Enter one of the following commands to release the quiesce on the database, which is now successfully backed up.

Option	Description
UPDATE AREA NAME(<i>area_name</i>) STOP(QUIESCE)	To release the quiesce on a Fast Path area that was previously quiesced.

Option	Description
UPDATE DB NAME(<i>db_name</i>) STOP(QUIESCE)	To release the quiesce on a database that was previously quiesced.
UPDATE DB NAME(<i>HALDBmaster_name</i>) STOP(QUIESCE)	To release the quiesce on a HALDB that was previously quiesced.
UPDATE DB NAME(<i>DEDB_name</i>) STOP(QUIESCE)	To release the quiesce on a Fast Path DEDB that was previously quiesced.
UPDATE DATAGRP NAME(<i>dbgrp_name</i>) STOP(QUIESCE)	To release the quiesce on a database group that was previously quiesced.

IMS SPOC returns the following output:

DBName	MbrName	CC
DBXYZ	IM03	0
DBXYZ	IM02	0
DBXYZ	IM01	0

- Enter the following command to ensure that the quiesced database has been released successfully: QRY DB NAME(DBXYZ) SHOW(STATUS) IMS SPOC returns the following output:

DBName	MbrName	CC	TYPE	LclStat
DBXYZ	IM03	0	DLI	ALLOCS,OPEN
DBXYZ	IM01	0	DLI	ALLOCS,OPEN
DBXYZ	IM02	0	DLI	ALLOCS,OPEN

Executing recovery-related functions

While IMS is running, an IMS system programmer or operator might need to execute functions relating to the recoverability of the system. You can initiate recovery through the type-1 and type-2 IMS commands.

You must ensure that resource definitions have been exported to a resource definition data set (RDDS) if you use DRD and perform a cold start from an RDDS. Otherwise, any resource changes made dynamically are lost across a cold start.

Issuing DBRC commands

You can issue batch and online (/RMxxxxxx) commands and DBRC API requests to use DBRC functions. Authorization checking is supported for DBRC commands that are submitted as batch or online commands.

The /RMxxxxxx commands perform the following DBRC functions:

- Record recovery information in the RECON data set
- Generate JCL for various IMS utilities and generate user-defined output
- List general information in the RECON data set
- Gather specific information from the RECON data set

The IMS security administrator can use a security product like RACF (Resource Access Control Facility), a security exit routine, or both to control authorization of DBRC commands. Commands issued through the z/OS console are always authorized.

Recommendation: Allow operators or automation programs to issue the /RMLIST command (or DSPAPI FUNC=QUERY API request) and the /RMGENJCL command (or DSPAPI FUNC=COMMAND COMMAND=GENJCL

API request). Restrict the use of /RMCHANGE, /RMDELETE, and /RMNOTIFY commands (or equivalent API requests), because they update the RECON data sets.

Attention: When the /RMLIST command is issued from the OM API (such as TSO SPOC), the amount of output that is generated can be large. Because the RECON data set is reserved during command processing, if the generated output is large, do not issue the command during peak times. To prevent the command from being issued with the DBRC='RECON' option inadvertently, consider protecting the command with a security product.

Related concepts:

- ➡ DBRC and data sharing support (System Administration)
- ➡ Supporting data sharing (System Administration)
- ➡ DBRC API (System Programming APIs)

Dumping the message queues

You use the /CHECKPOINT SNAPQ command to save the message queues in a nonshared-queues environment. This command dumps the message queues to the log without terminating the online system.

Recommendation: Schedule the /CHECKPOINT SNAPQ regularly because it shortens the time required for emergency restart if a problem occurs on the message queue data sets. Consider the following intervals:

- Whenever the OLDS is switched
- Once each hour
- Once each shift
- Twice each day (midnight and noon)
- Once each day

For a shared-queues environment, use the /CQCHKPT SHAREDQ command to dump the shared queues.

Recovering the message queues

In a non-shared-queues environment, you can recover the message queues during an IMS restart if the previous shutdown included the DUMPQ or the SNAPQ keyword.

Specify the BUILDQ keyword on the /NRESTART or /ERESTART command to restore the messages to the message queue data sets from the IMS log. Specify the FORMAT keyword on the /NRE or /ERE command if you also want to reinitialize the message queue data sets.

In order to use the /NRE BUILDQ command, the system must be shut down using a /CHECKPOINT DUMPQ | PURGE command. To use the /ERE BUILDQ command, you need only a prior /CHECKPOINT SNAPQ command.

Restriction: If a /NRE BUILDQ or /ERE BUILDQ command fails and you cold start IMS or cold start DC, messages are lost and are not processed.

You can use the IBM IMS Queue Control Facility for z/OS to select messages from the OLDS (or SLDS) and reinsert them into the IMS message queues after an IMS cold start.

For a shared-queues environment, CQS automatically rebuilds the message queues if the coupling facility fails. You can also use the SETXCF START,REBUILD command to rebuild the queues manually.

Related reference:

 [IMS Queue Control Facility overview](#)

 [IBM IMS Queue Control Facility for z/OS interface \(Diagnosis\)](#)

Archiving the OLDS

You should archive the OLDS to an SLDS at regular intervals. If you are not using automatic archiving, the MTO or automation program should issue the DBRC GENJCL command (or DSPAPI FUNC=COMMAND COMMAND=GENJCL API request) at regular intervals to generate the JCL for the Log Archive utility, and should execute the utility.

Making databases recoverable or nonrecoverable

You can change recoverable full-function DBs to nonrecoverable (after deleting recovery-related records from the RECON data set) by using the CHANGE.DB NONRECOV command.

You can change to recoverable again by using the CHANGE.DB RECOVABL command.

Use the LIST.DB command or the DBRC QUERY TYPE DB API request to display whether a database is recoverable.

Running recovery-related utilities


Depending on your recovery strategy, the MTO might be responsible for running various recovery-related utilities at regular intervals.


These could include:

- Database Image Copy utility
- Database Image Copy 2 utility
- Online Database Image Copy utility
- Database Change Accumulation utility

The MTO should also run these utilities when a database changes from nonrecoverable to recoverable.

Related reference:

 [Database Image Copy utility \(DFSUDMP0\) \(Database Utilities\)](#)

 [Database Image Copy 2 utility \(DFSUDMT0\) \(Database Utilities\)](#)

 [Online Database Image Copy utility \(DFSUICP0\) \(Database Utilities\)](#)

 [Database Change Accumulation utility \(DFSUCUM0\) \(Database Utilities\)](#)

Chapter 6. Developing operating procedures

After you have made basic design decisions about operations and recovery, you are ready to implement these decisions. Primarily, implementation takes the form of developing procedures for operating IMS.

The importance of developing good procedures cannot be over emphasized. The installations with the most success are those with well-written, well-maintained, well-tested, and well-understood procedures.

Establishing procedures for operating IMS is not a simple task. Because each installation has unique requirements and resources, no single approach to operations is possible. These topics provide some basic guidelines for establishing procedures.

Operations personnel

It is important that you identify the operational requirements for IMS, and identify the people responsible for performing the various operations tasks. After assigning responsibilities, you are ready to set up the procedures necessary to coordinate the operations task.

Many people are involved, directly or indirectly, in the operation of a successful IMS system. To help you understand who needs operating procedures and what it is they are responsible for operating, these topics look briefly at the various functions performed by these people. This topic also gives you an understanding of the context in which MTOs operate the system, and users or remote terminal operators use it. A fundamental understanding of the MTO is important because the subsequent topics describe the procedures you develop for the MTO and user.

An IMS DBCTL environment includes no IMS terminals, so the DBCTL operator must use a z/OS console. Throughout these topics, this DBCTL operator is called the IMS MTO.

The following groups of people represent functions often defined at an installation:

- **z/OS and JES system operators.** This group is responsible for operating z/OS and JES.

If the IMS MTO does not have access to a z/OS console, you need to establish procedures so the MTO can communicate with the z/OS operator in case of a failure. This is particularly important when the MTO is physically distant from the z/OS console.

- **Network support group.** This group is responsible for operation of the network, including: VTAM, physical telecommunications lines, and modems.

You must establish procedures for communication between the MTO and network support personnel for initializing the online system and isolating and correcting errors related to the IMS network.

- **DASD and tape pool operators.** This group is responsible for managing tape and disk media.

The MTO must communicate with DASD and tape operators when mounting archive tapes, and when creating image copies, change accumulations, and so forth. Communication is also required when the MTO is analyzing and correcting I/O errors.

- **User liaison group.** This group is the primary interface for the end user.
The user liaison group needs procedures for answering user questions relating to both IMS and individual application programs. The user liaison group also needs procedures for problem determination and correction, and for communicating with the MTO when unable to solve a problem. The user liaison group might also need to communicate with application support groups when a user problem relates to a specific application program rather than to the IMS subsystem.
- **Application support groups.** This group develops and maintains IMS application programs.
You need to identify representatives for the support of each application running on the IMS subsystem. You should give the user liaison group and the IMS MTO procedures for contacting this group in case of application problems.
- **IMS system administrator, system programmer, or recovery specialist.** This person (or group) is responsible for installing IMS and designing the system configuration.
There are certain error conditions the MTO cannot correct. For such errors, the MTO needs instructions on how to contact the IMS system administrator, system programmer, or recovery specialist. This person (or group) will help resolve the more complex recovery problems.

The following figure provides an overview of the relationships between these functional groups.

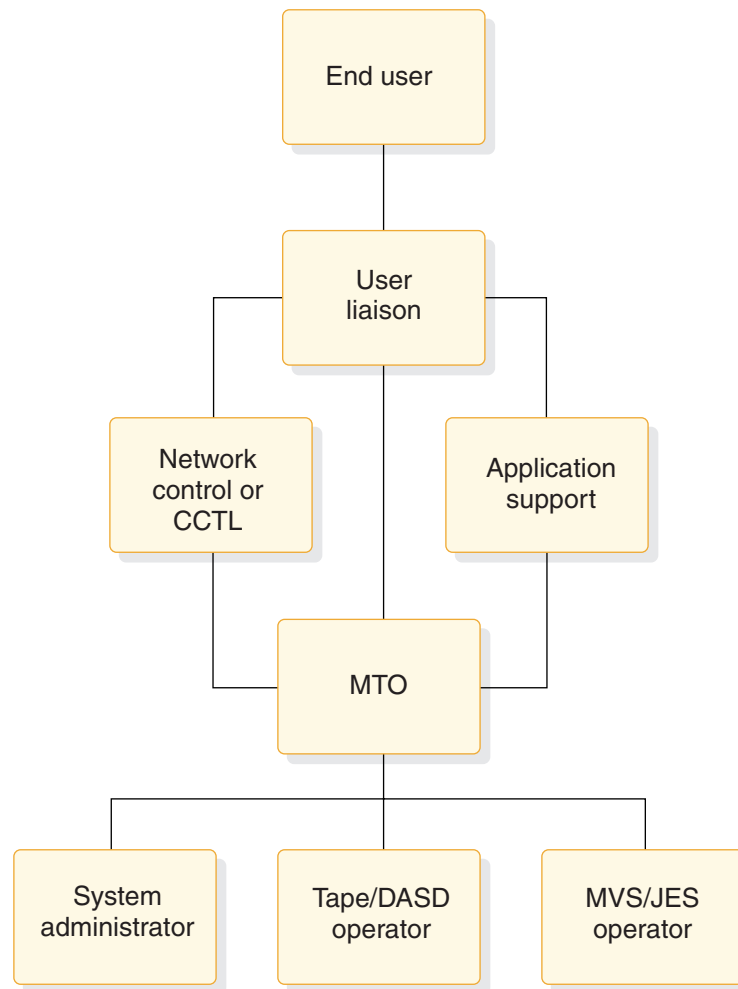


Figure 8. Relationships among operations groups

Depending on the size of your installation, some of these functions might be performed by the same person or group. In large installations, the same function might be performed by many different people. The IMS MTO, for example, might have access to a z/OS system console. This access would allow the MTO to enter z/OS START commands.

In this book, operations procedures are classified into two types:

- Those for operating IMS, which are generally performed primarily by the MTO and secondarily by the system administrator or recovery specialist
- Those for use by an end user

Developing procedures for an end user is not applicable in an IMS DBCTL environment, which has no DBCTL terminal end users. References to end users for a DBCTL environment actually refer to the CCTL end users.

Related concepts:

Chapter 7, "Developing user procedures," on page 285

Establishing operating procedure documents

You need to establish operating procedure documents as part of your overall strategy for controlling an IMS system. This will help you to ensure that operators will handle various situations in a consistent manner.

Establishing operations interactions

The online IMS environment should be transparent to general business procedures. It is probably only part of the data processing services. Your procedures, to some degree, reflect your installation's organization.

Definitions of support groups

Support groups provide specific functions for different groups of users, different environments, and even for different operations.

An *IMS operations group* must implement procedures and use tools, often specified by system programmers, that originate from requirements for database and data communication administration. The requirements for scheduling jobs, handling tapes and direct access devices, and producing printed output are developed in conjunction with general system operations procedures.

The *system operations group* provides z/OS support, and also often supports the hardware. The *software support group* provides support for the operating system and the language compilers. The *applications support group* provides support for the application programs, and is represented by user liaison and development maintenance groups.

The *user liaison group* is of special importance. Most of the problems that end users experience are application-program dependent. The user liaison group, acting as first point of contact, can answer such questions, but operations cannot. The user liaison group can also be the channel for problem resolution. They can initiate problem tracking, and their interpretation of error events can help all end users. The user liaison group analyzes problems to determine if they are application-program related, procedural, hardware, or system software. They can contact the master terminal control group regarding network or operations problems, or interact with application-program development support for programming errors. Very often, they are key personnel in any decisions affecting recovery, because they must interpret the actions taken on behalf of end users.

Operations contacts and responsibilities

Every organization needs to have layers of support and advice for control responsibilities for day-to-day operations. You need to plan for casual support and expert support.

Some suggestions for assigning control responsibilities are:

- Have a user liaison group answer day-to-day application-program questions, track problems, and monitor service.

Make the user liaison group the primary contact for the end user and for IMS operations. Encourage the user liaison group to be part of recovery decisions.

- Have a database-administration contact for recovery, design change evaluation, and database monitoring and analysis.

- Shield the MTO from most end-user inquiries so that the MTO can concentrate on availability, system status, and production statistics. Use other supervisory terminals to distribute the workload. Take advantage of automated operator programs to initiate MTO commands and reduce operator tedium.
- Establish system operations contacts for VTAM operation and hardware problems.
- Have a security contact for access problems. Other duties could be password and security maintenance, as well as violation monitoring.
- Have a monitoring and performance analysis coordinator who draws upon specialists for detailed studies.

These suggested contacts might only be suitable for prime shift or peak production hours. You should plan a contact list for other times, such as production cycle overruns and off-shift periods.

Plan for procedure requirements

The established operating procedures reflect many decisions and responsibilities. Development of the procedures takes place during the majority of the application implementation stages.

The content of the procedures draws upon several other administration tasks:

- The online IMS design as reflected in system definition
- The variety of application operation requirements
- The database administration procedures

The following table summarizes the tasks involved in developing a procedural document, relates them to subtasks, and briefly indicates the other groups of personnel involved. The tasks toward the bottom of the figure (starting with "Testing") are concerned with refinement of the procedures.

Table 28. Elements for establishing operation procedures

Development task	Operations procedure subtasks	Related support group
Initial structure	Plan terminal operator information Verify IMS message use Plan for automated operator Plan organization and contacts	User liaison (System support) Development
Application requirements	Plan scheduling Plan database availability Assess operator intervention Plan output handling	User liaison and DBA DBA User liaison and DBA Operations
Network control	Plan startup terminals Coordinate operator responsibilities Plan phased connections Plan alternate connections	User liaison and DBA Operations Operations Operations
System definition	Assess restrictions for MTO	Operations
Security	Design procedures Assess operator requirements	User liaison Security control

Table 28. Elements for establishing operation procedures (continued)

Development task	Operations procedure subtasks	Related support group
Administration	Establish system log control Plan error recording Plan network control desk	Operations Operations Operations
Problem procedures	Plan resource recovery	DBA and Development
Monitoring procedures	Plan IMS Monitor use Plan statistics gathering	Performance support User liaison
Recovery procedures	Decide on operator control Decide on problem analysis	Development Support groups and DBA
Testing	Respond to feedback	Development
Master terminal operator training	Respond to feedback	User education
Production control	Audit activities Respond to feedback	Operations User liaison
Modifications to applications	Respond to application-program changes Make tuning changes	Development Performance support

Planning the content of procedures

Your procedures for the master terminal operator (MTO) will vary depending upon their knowledge of IMS operations and the application-program logic

The operating procedure includes four types of instructions:

- The sequence of events for normal operation
- The information to be recorded
- The responsibilities and degree of autonomous action for operators
- The steps to be taken for error incidents or recovery action

Establish the pattern for normal processing by accumulating and combining the different requirements of both end users and databases.

The information the operators should record primarily concerns recovery actions; however, they should also record events that occur during the production cycle (in a form that can be audited). Many installations use a run book or system operation log for recording events during normal operations.

Whom you choose to have operational control strongly influences your procedure documentation. The instructions can vary a great deal, depending on their knowledge of IMS operation and the application-program logic. Based on your needs, you can train your MTO in one or more of the following areas:

- General operations
- Control of an application on a dedicated system
- Service for multiple applications

The arrangements you make to handle error or recovery events are also important in your control of the online IMS system. A knowledgeable operator can restart IMS and fall back to a valid checkpoint. An operator who suspects database damage can involve personnel from database administration or can schedule the required backout or change accumulation jobs.

A suitable document to guide the actions of MTOs should have several sections:

- An introduction

This section gives an overview of IMS operation (both generally and specifically for your installation), how the network is controlled, and how the master terminal is used. The intended audience is new or substitute operators.

- Operating procedures

This section shows the normal progression of events, but also provides supplementary instructions for error conditions. You can organize this section as a series of flowcharts.

- Commands reference section

This section shows syntax and recommendations for using IMS commands and other commands as appropriate. Include only those command keywords and parameters that you consider necessary for situations the MTO should handle without assistance from other groups. You might also want to include customized explanations of selected IMS messages to guide the MTO's actions.

- Configuration tables

This section contains descriptions of the IMS and system configuration. These descriptions help the operators to visualize the resources under their control. For example, you can include a chart of terminal connections related to application-program use, or representations of the databases that are online (together with the logical dependencies). You could also display larger copies in the operations area.

- Supplementary procedures

This section contains procedures not covered in earlier sections. With any large system, you can make alternative configurations available. This section can also contain recovery procedures indexed by symptom.

- A chronological record of events

This section lists the various events an MTO is likely to encounter during normal operations. This section could also list the types of events an MTO should record in an operations log or run book. Because an online IMS operation is likely to span several operator shifts, some type of record of the MTO's actions is desirable.

- A problem-reporting procedure

This section describes how to report problems or unusual events. An incident report is a useful tool to help you control the integrity of the system and service for the end user. Any of the following personnel or groups could fill in this report form: the end user or user liaison group, the IMS MTO or network administrator, the database administrator, or the system programmer.

Recommendation: Regardless of who fills in the form, you should require that it be completed, even though the person completing the form contacts other operating personnel. This form can become a controlling document in the case of a recovery action, but can also be used for minor complaints. To foster cooperation, you should return a copy of the report, showing actions taken, to the originator. You should also periodically review these reports to discover recurring problems and make valuable procedural changes.

Related concepts:

“Procedures for user terminal operators” on page 285

Related tasks:

“Planning availability of IMS service” on page 243

How to set the level of operator control

IMS is designed so processing is controlled by message traffic, which places control of the operation primarily in the hands of the master terminal operator. From an administration point of view, control of operations is far more than a structured use of IMS commands.

For a complex system such as IMS, you need to carefully match operational control to operator capability. You must decide whether an operator or an application administrator should have primary control.

The MTO can have a wide range of responsibilities. At one end of the spectrum, you have an operator using a standard set of instructions, designed using a cook-book approach: “if such-and-such happens, perform procedure X”. At the other end of the spectrum, you have a system programmer who understands the flexibility of IMS control as well as the application program and database requirements. The first approach requires good symptom-based instructions; the second requires clear guidelines and periodic auditing.

If an operator handles your operations, the two areas for which they need instructions, detailed to a level appropriate for their knowledge and responsibilities, are:

- Normal operations

You need to document instructions for making specific resources available for committed end-user service. The instructions must be especially clear about when and how to shut down and restart IMS. If alternative network connections are part of the system design, the instructions should have a set of additional procedures for them.

Operators who understand the end-user priorities should be allowed to change the scheduling algorithm by making reassignment of message priorities or even of message region classes.

- Error and recovery operations (triggered by symptoms or user feedback)

These events range from temporary transaction processing problems to problems that have serious database implications. Your recovery procedures should include carefully documented and tested actions for your operators to take.

Recommendation: Design your procedures so the operator initiates recovery actions in cooperation with your user-liaison group and your system-support personnel. Structure your operations to take advantage of as much consultation as possible. Insist that your operators record all recovery actions in a run log, including justification for decisions made.

An important part of the operations strategy is the recovery procedures used to restart the online IMS subsystem after a failure of the operating system or the IMS control region.

Record-keeping procedures for the MTO

The master terminal operator (MTO) needs to keep various types of information that document system activity. Consider designing standard forms or online panels for recording this information.

The MTO needs to record the following general types of information:

- IMS control region activity
- Utility executions
- Application program activity
- Problems or unusual events
- Startup, shutdown, and system log activity

Some of this information can be recorded automatically. DBRC records all OLDS and archive activity, and can also record the execution of batch jobs and recovery-related utilities.

By default, the master secondary terminal provides a hardcopy log of IMS system activity, including checkpoint information and IMS messages. The logging of IMS system messages to the secondary master terminal can be disabled by using the /SMCOPY MSG OFF command.

You can also have the secondary master terminal provide a hardcopy logging option for certain IMS commands and command responses by specifying the COPYLOG keyword in the COMM system definition macro or by issuing the /SMCOPY command with the appropriate keyword.

The information logged to the secondary master terminal can be helpful when tracking a series of events affecting the operation of IMS, or when verifying that an operations task was performed. Information useful for analyzing an IMS error situation is often in the z/OS hardcopy log.

You can select whether IMS logs commands issued from the master terminal or commands entered from other terminals, or both. In either case, IMS only logs certain commands. When you issue a command from a terminal other than the master terminal, IMS logs output and identifies its node, or line and physical terminal number.

The /SMCOPY command starts and stops the hardcopy logging facility for commands, system messages, or both.

IMS logs the following commands to the secondary master terminal:

```
/ACTIVATE  
/ALLOCATE  
/ASSIGN  
/CHECKPOINT  
/CLSDST  
/COMPT  
/DBDUMP  
/DBRECOVERY  
/DELETE  
/DEQUEUE
```

/DISPLAY
/IDLE
/MODIFY
/MONITOR
/MSASSIGN
/OPNDST
/PSTOP
/PURGE
/QUIESCE
/RCLSDST
/RCOMPT
/RMCHANGE
/RMDELETE
/RMGENJCL
/RMINIT
/RMLIST
/RMNOTIFY
/RSTART
/START
/STOP
/SWITCH
/TRACE
/UNLOCK
SYSTEM

Recording IMS control region activity

You should design a form on which the master terminal operator (MTO) can document executions of the control region. Details to be recorded on the form include: the time when the system is started or shut down, the type of restart, and the reason for shutdown. You should make provisions on the form for recording error situations, such as DASD I/O errors, and actions taken to recover from such errors.

The operator can file the output of the DBRC LIST.LOG command (or DSPAPI FUNC=QUERY API request) with this form to provide a record of log archive activity.

Utility recording forms for IMS utilities

How much manual record keeping is necessary for IMS utilities depends on your use of DBRC. If you have registered all your databases with DBRC, DBRC records the execution of the recovery-related utilities in the RECON data set. By listing the RECON data set (LIST.RECON command or DSPAPI FUNC=QUERY API request) at regular intervals, you can easily maintain records of utility executions.

If you have not registered all your databases with DBRC, you should design forms for recording execution of the following utilities:

- Database Image Copy
- Database Image Copy 2
- Online Database Image Copy

- Database Change Accumulation
- Database Recovery
- Those utilities involved in database reorganizations

You also need to record the execution of batch database reorganizations. A database cannot be backed out to a point earlier than its latest batch reorganization. Nor can a forward recovery be done from an image copy preceding the latest batch reorganization. Thus, a knowledge of recent reorganizations is important when recovering a database.

The following graphics are sample forms for different utilities. These forms assume that changes for all databases are accumulated together. The first graphic shows the input log data sets, and the output change accumulation data sets.

Run		Data set name	Volume serial number	Purge		Input data sets*	
Date	Time			Date	Time	Name	Volume serial number
91005	1520503	IMS.CUM	CUM03	91004	1200000	IMS.CUM	CUM02 SLDS01 RLDS01
91006	1510031	IMS.CUM	CUM04	91005			

*Includes old change accumulation data set, and SLDS or RLDS

Figure 9. Sample form: running of the Database Change Accumulation utility

Image copy record for database _____ data set _____			
Creation		Data set name	Volume serial number
Date	Time		
91004	1100100	DB1.IC1	DB1IC1
91005	1100051	DB1.IC2	DB1IC2

Figure 10. Sample form: running of the Image Copy utility

These forms are only samples. Use them as templates to design forms that suit your installation's procedures.

Application program recording forms

You might want to design forms to record the execution of BMPs and batch application programs. The latter is particularly important if you are not using DBRC to record logs produced by batch programs. Other information you might want to record includes checkpoint IDs and any error messages produced by the programs.

Forms for recording problems or unusual events

You might want to design an incident report form for recording any system problems or unusual events. Incident reports help you control the integrity of the system and services for end-users. This form might be filled in by the user liaison group, the MTO, or the recovery specialist.

Recommendation: Regardless of who fills in the form, you should require that it be completed, even though the person completing the form contacts other operating personnel. This form can become a controlling document in the case of a recovery action, but can also be used for minor complaints. To foster cooperation, you should return a copy of the report, showing actions taken, to the originator. You should also periodically review these reports to discover recurring problems and make valuable procedural changes.

	Hardware	Software	Sequence number	
Component in error & symptom			Date	
			Time	
			Incident reported to	

Type of dump available	(x)	Time DUMP was reported
SYSUDUMP		
MVS DUMP		
Stand-alone DUMP		

Full description of problem	Action taken

Figure 11. Sample form: incident report

Forms for recording startup, shutdown, and system log activity

You might want to design an online run sheet for recording startup, shutdown, and their relationship to system log activity.

/NRE	BUILDQ	Date	
/ERE	FORMAT ALL	Start time	
/CHKPT		Stop time	
		MTO name	
/CHE	DUMPQ FREEZE		
Other (Specify)			
Last CHKPT ID			

Time	Comments/Incidents

Figure 12. Sample form: online run sheet

The MTO fills in the online run sheet for each execution of the online system. When starting IMS, the MTO records (circles) the parameters used in the start command, along with the checkpoint ID used. When shutting down IMS normally, the MTO records (circles) the parameters used. When shutting down IMS abnormally, the MTO records the reason for the failure and fills in an incident report form. If the MTO shuts down IMS using a /CHECKPOINT DUMPQ command, that person records the closing checkpoint ID from the DFS994I message.

The MTO should use the bottom half of the form (comments/incidents) to record any unusual incidents. The MTO should note such things as: hardware errors on telecommunication lines, details about reassignment of LTERM names, and abending transactions. Finally, the MTO should add a cross reference to the incident report, where necessary.

DBRC records information about use of OLDSs in the RECON data set. After IMS archives an OLDS for an execution of the online system, the MTO can issue a LIST.LOG command to get information about the use of the OLDS. The MTO should file the hardcopy listing of this information with the online run sheet. This information provides an important means of communication between one shift of MTOs and the next. This information can also be used to assist in reconstruction of the RECON data set, if it becomes damaged.

The MTO should also file the complete hardcopy listing from the master terminal printer with the online run sheet.

OM audit trail

The IMS Operations Manager (OM) uses the z/OS system logger to provide an audit trail. IMS operator commands issued through OM, as well as their command response, are saved to the system logger. In addition, some unsolicited messages from OM, SCI, and IMS are also logged in the audit trail.

All input is logged before OM calls the OM Input user exit. Input is logged upon return from the OM Input user exit only if the exit modified the input. All output is logged before OM calls the OM Output user exit. Output is logged upon return from the OM output user exit only if the exit routine modified the output.

System logger

The system logger is a z/OS component that allows an application to log data from a sysplex environment. Data can be logged from one z/OS system or from multiple systems across the sysplex.

The z/OS system logger is a set of services that allows an application to write, browse, and delete log data. You can merge the log data from applications in a sysplex into a log stream, which is simply a collection of data in log blocks residing in the coupling facility; as the coupling facility fills, older data might be off-loaded to archival log stream data sets on DASD. Log blocks in the coupling facility can be backed up either in storage in each system or on DASD in staging data sets. When the log blocks in the coupling facility reach an installation-defined threshold value, they are off-loaded to DASD log data sets. Therefore, at any point in time, the log stream consists of records on the DASD log data sets (where data is hardened for longer term access) and the log blocks that currently reside in the coupling facility.

A system logger configuration includes the system logger address space in each system, the LOGR couple data set, a log stream structure in a coupling facility, DASD log data sets for off-loaded data from the coupling facility log stream, and optionally staging data sets for a backup copy of the log blocks resident in the log stream structure.

The system logger address space provides the application with services and connections to the coupling facility. An installation must:

1. Plan and predefine the coupling facility structures
2. Format and define a policy for the LOGR couple data set.
3. Define the DASD log data sets.
4. Optionally create staging data sets.

Related concepts:

 Using system logger services

 System Logger

Planning availability of IMS service

In order to provide a framework for the operator control of the online IMS subsystem, you need to obtain overall requirements. You can then proceed gradually to fill in the details of specific requirements of various user groups.

Related concepts:

“Planning the content of procedures” on page 234

Checklist for defining the production cycle

When defining the production cycle, obtain the overall production cycle requirement and the support requirements for databases and batch-program scheduling. Collecting the initial availability requirements ensures that you can plan your follow-up activities accordingly.

Take the following suggested steps to define the production cycle:

1. Obtain the overall production cycle requirement.

For example, suppose the execution of an application program requires the online subsystem to be available from 0400 to 2000 hours, with no weekend requirements. Very often these requirements are made to fit the data processing organization's constraints. In this example, the operations group must be agreeable to the early-morning scheduling (which is probably required for transaction input from locations in different time zones).

2. Obtain support requirements for databases and batch-program scheduling.

For example, a daily batch-report program might cause certain databases to be unavailable until 0800. A /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command stops additional transaction or CCTL access and allows offline processing for the database named in the command. Another example is a priority transaction, scheduled for database integrity purposes, that preempts other transaction processing.

The following table shows a checklist of initial availability requirements you need to collect.

Table 29. Availability requirements checklist

Availability requirement	Follow-up activity
What are the available hours: Daily? Weekly?	Check if a regular weekly time. Check daily start and stop times. Allow for weekends, holidays, and preventive maintenance.
Which are services critical to the user?	Summarize the transactions, programs, and database needs. Summarize the CCTL's dependencies for programs (PSBs) and databases.
What is the target maximum down time?	Establish guidelines for duration.
Are there any planned interruptions?	Check requirements for normal restart. Determine how queues are to be handled.
Are there variations of nucleus?	Match nucleus suffix to system definition schedule.
Is there any special message region scheduling?	Develop IMSMSG procedure variation. Find conditions that prompt scheduling.
What are the BMP scheduling requirements?	Tailor IMSBATCH procedure for BMP. Find z/OS file requirements.
What databases are not always online?	List the databases, their conditions, and the reasons for their being offline.
What batch programs are to run during online schedule?	Find any database conflicts. Find the required database levels and restart instructions.
What is the timing of database reorganization and related jobs?	Find jobs that run during online execution. Obtain master schedule from DBA.

Table 29. Availability requirements checklist (continued)

Availability requirement	Follow-up activity
Do any application programs require exclusive processing?	Find the effect on other online processing.
Are there configuration variations or multiple IMS subsystems?	Develop procedure support: plan for MSC, data sharing, DL/I separate address space, DBCTL, DCCTL, APPC/IMS, OTMA, or ETO.
Are you using a data propagation manager?	Make sure both IMS and DB2 for z/OS are available on the same z/OS system. Make sure the program libraries that contain the data propagation user exit routines are available to all dependent IMS regions.

Devices available to the network

Document which devices are to be initially available, those that should always be active, and those that have restricted usage. The important attribute is whether the MTO needs to activate the connection or must make any alternative connections in case of terminal or line failures.

A second aspect of availability is to ensure proper control of the network. You need to document what terminals are connected and when.

The following table identifies some connection data you need to collect.

Table 30. Checklist for network control connection data

Connection data	Follow-up activity
Which terminals are associated with services critical to the user?	List the connection requirements. List the LTERMs.
Are VTAM terminals to be continuously connected?	List those nodes and their LTERMs.
Are VTAM terminals to be connected on demand?	Identify how each is controlled (MTO, VTAM operator, or remote terminal operator).
Are there planned reconfigurations?	Develop sets of terminals and identify how they should be controlled.
Are there alternate connection plans?	Identify if there are alternate pairings, or occasional usage.
Are intelligent stations as connected subsystems used?	Identify operations for connected subsystems or terminals, and monitoring requirements.
Are switched lines used?	Identify the users and give the LTERM names in the group.
Is there batched data entry or high-volume input?	Identify the pattern of batched devices.
Are there special output storage devices?	Identify use of printers or data devices.

Operations responsibilities for online control services

It is important to understand why an application program places restrictions on what commands you can issue. For example, online image copy is required at end-of-shift, but, because of the requirements of an IMS restart, no conflict in database level is permitted. This kind of restriction might require you to prevent certain programs from executing.

A third checklist identifies the use of support services and expected MTO control actions. The following table presents this information; your emphasis should be on the operational requirements.

Table 31. Checklist for online control services

Check requirement for:	Follow-up activity
Dynamic allocation of databases	Requires coding of DFSMDA macro and prior cataloging of databases.
Online database image copy	List databases and schedule timing. Is image copy restart allowed?
Database surveyor	When to be scheduled. Who is the contact for range specification?
Database recovery control	Arrange for DASD dump entries. How are RECON data sets to be controlled? How are JCL generation facilities to be used?
Simple checkpoint	What conditions require the checkpoint? Is SNAPQ required?
/CHECKPOINT PURGE /CHECKPOINT DUMPQ /CHECKPOINT FREEZE /NRESTART	Must all queues be emptied? Can queues be held? What conditions warrant use of this command? Is queue rebuilding required?
Online system log management	Determine archive procedures and whether WADS and OLDS are to be dynamically allocated.
Online change management	Determine the scope and timing of online changes and how to control processing.
Offline dump formatting management	Establish transfer, archive, and formatting procedures.
IBM 3850 Mass Storage System (MSS)	Obtain scenario of IBM 3850 usage.
Monitoring	How will IMSMON data set be handled? What requirements exist for traces or statistics?
Security	Any override conditions?
Signon	Any override conditions?
Transaction	What is update frequency and control procedure?

MTO abilities and operator responsibilities

The MTO's primary duty is to monitor the condition of the total IMS online environment, and to maintain normal service. Any alterations in committed service or recovery decisions are really part of a wider set of business responsibilities, not necessarily those of the MTO.

When you are defining the duties of the master terminal operator, you must decide what actions are part of the normal duties, and what decisions require consultation with other groups. From the end-user viewpoint, the MTO could represent the data-processing staff, as well as the application-program contact. You need to arrange ahead of time what kind of support can be provided by the MTO, and if there is any way of shielding the MTO from unnecessary interruption.

The master terminal operator has the following abilities and responsibilities:

- Responsibility for running IMS

The MTO starts and shuts down the IMS subsystem, starts regions, and manages the system log.

- Knowledge of the ongoing status of the IMS subsystem

The MTO continuously monitors processing and detects any error situations.

- Control over contents of the system and the network

The MTO can control the network, connect other IMS systems, and perform other prearranged tasks.

- Privileged commands

In addition to the routine work, the MTO responds to error conditions, changes the scheduling algorithm, alters passwords, and reconfigures the system as necessary.

Normal operator actions (DB/DC or DCCTL)

The z/OS system operator needs to initialize the IMS control region, but you should plan to reduce the need for further involvement by the z/OS operator in IMS operations. In a busy system, responses to messages or instructions sent by an application program to the z/OS operator can be erratic. Such requirements are better handled by output messages to the IMS master terminal.

You should ask the z/OS system operator to relay any IMS messages (ones beginning with any of the following prefixes: DFS, DSP, DXR, ELX, BPE, or CQS) to the IMS master terminal. You might also give the z/OS operator a list of the IMS job names. Because often the z/OS system console is physically separated from the IMS master terminal, you need a local telephone or a paging line to maintain communication between them.

You need to decide whether the z/OS operator or the IMS MTO should start IMS regions. If you want the z/OS operator to start them, one of the following IMS procedures can be used:

- IMSBATCH procedure to start batch message processing regions
- IMSRDR procedure to start message processing regions

The advantage of allowing the z/OS operator to start IMS regions is that the operator can monitor any abnormal termination patterns and I/O requests.

The following table shows the actions usually performed by the MTO and the commands usually reserved for the IMS master terminal.

Table 32. Master terminal operator actions usually performed

Activity	IMS command
Activate IMS (cold start)	/ERESTART COLDSYS
Start a message region	/START REGION IMSMSG1
Start communication lines	/START LINE ALL
Display message queues	/DISPLAY
Start another message region	/START REGION IMSMSG3
Prepare for VTAM communication	/START DC
Initiate static VTAM sessions	/OPNDST NODE ALL
Initiate dynamic VTAM sessions	/OPNDST NODE AOI
Send message to terminals	/BROADCAST
Shut down VTAM terminals and IMS	/CHECKPOINT FREEZE QUIESCE

Table 32. Master terminal operator actions usually performed (continued)

Activity	IMS command
Restart IMS (warm start)	/NRESTART

The IMS master terminal also works well as a display device. On the bottom of the screen are two input lines, with room for a warning message. The remainder of the screen is divided into two areas, each of 10 lines. The top portion shows the 10 most recent IMS messages; new messages overlay old ones, so the display cycles through the 10 lines. The lower portion is used to page through the output from the /DISPLAY command.

If you set the master terminal to use a 3270 format, some of the program function keys are predefined as follows:

PF Key	Command or function
1	/DISPLAY
2	/DISPLAY ACTIVE
3	/DISPLAY STATUS
4	/START LINE
5	/STOP LINE
6	/DISPLAY POOL
7	/BROADCAST LTERM ALL
11	Next message page (NEXTMSGP)
12	Copy function

Tip: Create a PF key panel that shows these settings and place it with the keyboard as a reminder for the operator.

Normal operator actions (DBCTL)

The DBCTL MTO can control the DBCTL environment from any z/OS system console. If your MTO operates from a different console (secondary system console) from the primary one where the z/OS operator sits, you must ensure a means of communication between the two operators.

DBCTL requests for tape or DASD mounts must be given to the system console operator. Some IMS messages (DFS prefix) originate from the CCTL region and thus appear on the primary system console; these messages should be relayed to the MTO in case some action needs to be taken. For example, when the CCTL tries to connect to a DBCTL environment and the DBCTL subsystem is not yet running, the IMS database resource adapter (DRA), which is part of the CCTL region, issues message DFS690A. In general, communication between the CCTL operator and the MTO is important.

As in the DB/DC or DCCTL environments, the z/OS operator can start IMS BMPs by using either the IMSRDR or the IMSBATCH procedure. Any activity of IMS dependent regions or utilities the z/OS operator controls should be recorded or communicated so that the MTO is aware of all IMS-related activity.

The following table shows the actions usually performed by the MTO and the commands usually reserved for the IMS DBCTL master terminal.

Table 33. Master terminal operator actions usually performed

Activity	IMS command
Activate IMS (cold start)	/ERESTART COLDSYS
Start a BMP	/START REGION IMSBMP
Display dependent regions	/DISPLAY ACTIVE
Shut down IMS	/CHECKPOINT FREEZE
Restart IMS (warm start)	/NRESTART

Resources controlled by the MTO

Resources controlled by the master terminal operator (MTO) are listed, including communications, scheduling, and integrity.

The following table summarizes the resources controlled by the MTO.

Table 34. Resources controlled by the MTO

Communications	Scheduling	Integrity
Telecommunication lines	Regions	Checkpoint
Physical terminals	Databases	Restart
Logical terminals	Transactions	Database copy
Spooled Output	Programs	Database availability
Users		System status

Connections for network operations

The major part of any detailed instructions for the MTO is the required connections or any alternate connections (in the case of terminal or line failures). A useful aid for the MTO is a chart showing the following:

- The network
- The device type
- PTERM and LTERM names
- Availability requirements
- A summary of the application usage

The chart should specify the contacts for problem resolution, whether they are user liaison, VTAM operator, or some other form of network control.

Session initiation and termination options

Because of their functional capability, many VTAM-supported terminal devices are shared by IMS and other users. You have several choices for handling connect and disconnect. You can have the remote operator or master terminal operator be responsible, or you can call upon intervention by a VTAM operator. The following table summarizes your choices.

Table 35. Session initiation and termination options

Control by	Initiate by	Terminate by
MTO	/OPNDST	/CLSDST
Remote operator	LOGON IMS	/RCLSDST
(VTAM,Defn)	Initialization list	-
VTAM operator	VARY command	VARY command

Secondary master terminal and auditing operational control

Some IMS commands are automatically copied to the terminal designated as SECONDARY in the system definition NAME macro. If you specify COPYLOG=ALL in the COMM macro during system definition, then when other terminals besides the master terminal enter any of these commands, copies of an entered command and its response are sent to the secondary master terminal.

To understand the actions taken by the MTO, you should obtain a permanent printed copy of the command sequences issued by the MTO. This function is termed *hardcopy logging*.

The MTO can use the /SMCOPY MASTER OFF command to override the logging of commands to the secondary master terminal. The MTO can also use this command to turn off logging of command entry and responses from other terminals. To resume logging to the secondary master terminal, the MTO can use the /SMCOPY MASTER ON [TERMINAL ON] command.

This more permanent record is valuable for several reason, but the principal one is that of accountability. You can also use the output to audit whether the operations command sequence and timing are providing committed service. The checklist in the following table identifies the types of activities and related commands, and shows some of the ways you can use the secondary master terminal output.

Table 36. Using secondary master terminal output to audit performance.

Type of activity	Command	Output usage
Network control	/OPNDST	Check VTAM connection procedure.
	/CLSDST	Check VTAM disconnect procedure.
	/RCLSDST	Check expected disconnect.
	/START	Check network availability.
	/RSTART	Coordinate with terminal usage.
	/STOP	Investigate stopped resource.
	/IDLE	Investigate if there is lost or waiting output.
	/MSASSIGN	Check if planned MSC connect.
Device control	/COMPT	Check planned VTAM component switch.
	/RCOMPT	Check remote VTAM component allowed.

Table 36. Using secondary master terminal output to audit performance (continued).

Type of activity	Command	Output usage
Resource control	/ASSIGN	Check that all uses fit operations plan.
	/START	Check region availability and queue control.
	/STOP	Investigate all events.
	/PSTOP	Investigate why input stopped.
	/PURGE	Investigate why service was restricted.
	/MODIFY	Audit the extent and timing of online changes.
	/CHANGE	Check size of MFS dynamic directory.
Security	/MODIFY	Set timing of authorization and password changes.
Recovery	/DBRECOVERY	Investigate database errors.
	/DBDUMP	Investigate database errors and subsequent processing.
	/DEQUEUE	Investigate lost message or output.
	/CHECKPOINT	Coordinate with reported errors.
	/SWITCH	Investigate reasons for unplanned takeover activity.
	/UNLOCK SYSTEM	Investigate takeover activity.
Monitoring	/DISPLAY	Incorporate in monitoring feedback. Check use for scheduling.
	/TRACE	Coordinate with trace output.
HALDB OLR	/INITIATE	Initiate or resume HALDB Online Reorganization (OLR).
	/TERMINATE	Stop HALDB OLR.
	/UPDATE	Change the rate or the disposition flag of the DEL or NODEL data set of a HALDB OLR, or the disposition flag of REL or NOREL to release ownership of OLR at normal or abnormal shutdown of IMS.

Operator control of conversational transactions

The master terminal operator should be aware of the potential use of conversational transactions because their queue characteristics are somewhat different from those of nonconversational transactions. The number of conversational transactions that can be active at one time is 65535 per user or terminal. Also, the same individual transaction can be scheduled many times, but the processing patterns can be quite different among the set.

For example, one conversational transaction could involve an extended search of a database, while another, using multiple-choice menus, might retrieve a single item of data.

Another consideration for conversational transactions is whether the conversations are to be active until they are complete or whether they can be held while the end user is engaged in an offline activity.

Controlling conversational transactions

The following table summarizes the type of information that influences the content of your operational procedures when the online IMS subsystem includes conversational transactions.

Table 37. Operator information for conversational transactions

Transaction attribute	Information required	Procedure impact
Which transactions are conversational?	Summary of purpose and processing characteristics.	Multiple messages might need the message format buffer pool and message queues to be monitored.
Any dedicated terminals? What is the maximum number allowed to run?	Name LTERM and PTERM. Find maximum storage using DASD and expected loads.	Know impact of terminal problem. Advise remote terminal operator of restrictions and status.
Can conversations be held or released?	If so, how long? Is operator allowed to terminate?	MTO awareness of status. Use of /PSTOP and /EXIT commands. Use of /MODIFY warnings.
What recovery actions are to be taken?	At restart. Are abnormal termination exit routines used?	Notify remote terminal operator and other recovery personnel of output loss.

Remote terminal operator control for conversational transactions

A remote terminal operator can use a /FORMAT command to begin a conversation. This command formats an MFS-supported terminal screen. Or, the terminal operator can enter the initial transaction. To terminate the conversation before the normal program end, the terminal operator can use the /EXIT command.

The operator can suspend the conversation (if the operator needs to leave the terminal area) by using the /HOLD command. In response, IMS issues message DFS999I, which displays the identification number needed to resume the conversation with the transaction program. The terminal is then available for other transaction activity.

On return to the terminal, the terminal operator can issue the /RELEASE CONVERSATION *nnnn* command to resume the conversation, where *nnnn* is the conversation ID from the DFS999I message. As a convenience, IMS re-sends the last output message to the terminal. Or the operator can issue the /EXIT CONVERSATION *nnnn* command to terminate the conversation.

If the MFS formatting or the output content of the screen is lost during a conversation, enter the /HOLD command followed by the /RELEASE command. These commands resend the first physical page of the current message and restore the format.

MTO control for conversational transactions

The master terminal operator can obtain a status of all conversations, held or active. The following output from the /DISPLAY CONVERSATION command displays the terminal address, the conversation identification number, and status and type of all conversations:

```

TERMINAL USER    ID STATUS
11-   2          0001 HELD
  4-   2          0002 ACTIVE,SCHEDULED
  4-   1          0011 HELD
VT03              0012 ACTIVE

```

The MTO can terminate a conversation. For example, the /EXIT CONVERSATION 0001 LINE 11 PTERM 2 command terminates the first conversation shown in the command output above. A /START LINE, /START NODE, or /START USER

command also terminates the conversation, unless a program is already scheduled to process a portion of the conversation exchange.

Recommendation: Before terminating a conversation, the MTO should broadcast a warning message to that terminal.

Exit routines for abnormal termination of conversations

When a conversation terminates abnormally, a termination exit routine named DFSCONE0 can gain control. Its use might be to initiate cleanup activity, or simply to notify the input terminal of the canceled conversation. If the exit routine determines that more elaborate correction is required, it can schedule a transaction. This transaction can include the latest SPA, or a modified version of it, as a message segment. IMS provides a default exit routine (DFSCONE0) which schedules a transaction that includes the old SPA.

You can use the conversational termination exit routine to determine the effect of operator interruption.

Related reference:

 Conversational Abnormal Termination exit routine (DFSCONE0) (Exit Routines)

MTO and tracing operations

The MTO or an authorized remote terminal operator can initiate several kinds of trace activity with the /TRACE command. Even though you might use the IMS Monitor fairly regularly, you should regard all tracing operations as “special requests”.

Because trace or monitoring actions affect the performance of the online IMS system, they are not part of normal operations. The validity of the trace depends upon the processing taking place. Also, the operator can control the trace interval with the /TRACE command, so the operator must be aware of the trace timing.

Some types of trace requests do not require the person requesting the trace to be present during normal operations. The request should state the approximate duration of the trace and the type of monitoring. The main section should specify the initial start and termination conditions, as well as any monitoring constraints to be applied. Other sections should prompt for other details, such as queue levels, transactions to be active, or the time of day.

The types of tracing that can be requested are summarized in the following table.

Table 38. Trace and monitor options.

Type of trace	/TRACE parameters	Operational effect
IMS Monitor	SET MON opt	Initiates IMS Monitor trace.
	SET OFF MON	Terminates the trace.
	opt=ALL	Traces all events.
	opt=APDB	Traces activity between application programs and databases.
	opt=APMQ	Traces activity between application programs and message queues.
	opt=LA	Traces line and logical link activity.
	opt=SCHD	Traces scheduling and termination events.

Table 38. Trace and monitor options (continued).

Type of trace	/TRACE parameters	Operational effect
DL/I Image Capture	SET PSB name opt	Traces all DL/I calls for the named PSB.
	SET OFF PSB name	Terminates the trace.
	opt=COMP	Generates PCB and data-compare statement images.
	opt=NOCOMP	Generates no compare statement images. Default.
	SET PI OPTION opt	Traces program isolation activity, including ENQ, DEQ, and DL/I calls. Trace data is written to system log.
	SET PI	Traces PI activity. Trace data is kept in storage, and not written to system log.
	SET OFF PI	Terminates the trace.
	opt=LOG	Writes trace data to system log. Default.
	SET OFF PI OPTION LOG	Continues trace, but does not write buffers to system log.
	opt=TIME	Adds elapsed time for each ENQ and DEQ call when IMS has to wait.
	SET OFF PI OPTION TIME	Continues trace, but only records time of day.
	opt=ALL	Equivalent to opt=LOG TIME.
Line, node, and PTERM tracing	SET LINE nn opt	Traces line activity.
	SET OFF LINE nn	Terminates the line trace.
	SET NODE nn opt	Traces node activity.
	SET OFF NODE nn	Terminates the node trace.
	opt=LEVEL n	Sets detail level (1-4) for trace.
	opt=MODULE mod	Traces modules for IMS communications manager.
	mod=DDM	Traces device-dependent modules.
	mod=MFS	Traces MFS modules.
	mod=ALL	Traces both DDM and MFS modules.

Table 38. Trace and monitor options (continued).

Type of trace	/TRACE parameters	Operational effect
Internal Trace Tables	SET TABLE opt OPTION LOG	Initiates tracing into specified trace tables. IMS writes trace tables to log before reusing them.
	SET TABLE opt OPTION NOLOG	Initiates tracing into specified trace tables. IMS does not write trace tables to log. Default.
	SET OFF TABLE Opt	Terminates trace.
	opt=ALL	Specifies that IMS enable or disable all traces. Default.
	opt=DISP	Traces dispatching events.
	opt=DL/I	Traces DL/I calls.
	opt=DLOG	Traces log activity.
	opt=FAST	Traces Fast Path activity.
	opt=IDC0	Traces errors in modules DFSCNXA0 and DFSIDC00.
	opt=LATC	Traces latching activity.
	opt=LOCK	Traces locking and PI activity.
	opt=LRTT	Traces log router activity.
	opt=LUMI	Traces LU 6.2 and APPC activity.
	opt=OTMT	Traces OTMA activity.
	opt=RETR	Traces DL/I retrieve calls.
	opt=SCHD	Traces scheduling events.
	opt=STRG	Traces storage manager activity.
	opt=SUBS	Traces external subsystem events.
Link Tracing for MSC	SET LINK nn opt	Traces logical link.
	SET OFF LINK nn	Terminates link trace.
	opt=LEVEL n	Sets detail level (1-4) for trace.
	opt=MODULE mod	Traces modules for IMS communications manager.
	mod=DDM	Traces device-dependent modules.
	mod=MFS	Traces MFS modules.
	mod=ALL	Traces both DDM and MFS modules.

IMS can trace several different types of activity concurrently. The IMS Monitor writes its output to a data set specified on the IMSMON DD statement or dynamically allocated. IMS writes other traces to type X'67' log records.

Related reference:

 /TRACE commands (Commands)

Offline dump formatting operations

IMS operators need to be aware that their intervention might be required if a dump is requested at a time when the dump data sets are unavailable or when some other task in the z/OS system is currently taking an SDUMP.

In these situations, IMS issues message DFS3906 or DFS3907. In addition, CCTL operators must be aware that the CCTL can issue these messages when the IMS database resource adapter (DRA) is trying to generate an SDUMP.

Designing operating procedures

You can document operating procedures using flowcharts or a narrative style.

Operating procedure design using flowcharts and other graphic techniques

Flowcharts have several obvious advantages over narrative presentations of procedures. For example, with a flowchart it is easy to get an overall picture of what needs to be done. Another advantage is that when text accompanies each flowchart, you can skip reading any text associated with a step that you already understand.

If the flowchart, for example, says to start IMS and this step is understood, the reader does not need to read that “the system operator enters the /NRE command from the system console...”.

When you use flowcharts or other graphics to document operating procedures, you generally accompany them with an expanded textual description of the various steps in the flowchart. The main disadvantage to using flowcharts is that they are often difficult to maintain. Small changes to an operating procedure could require you to redesign an entire flowchart.

You might not need to use flowcharts for the IMS MTO. However, it is probably the preferred technique when designing procedures for a recovery specialist or others who need an overview of how something works or what actions are to be performed.

Operating procedure design using narrative

The advantage to using a narrative style is that it is relatively easy to maintain. There are no graphics to redraw, and changes to the procedures seldom require much redesign.

The disadvantage is that the overall operation being performed is not always obvious. In addition, if the users of the procedure suspect that they have taken a wrong branch in the procedure, it is difficult to backtrack through the steps.

There are many ways to document procedures using a narrative style, rather than graphics. For example, you could use a numbered list, where the reader starts at number 1, compares the situation with the information in number 1, then continues to number 2 or maybe jumps to number 15 depending on the decision made in number 1. This technique is essentially a narrative flowchart. As another example, you could write scenarios, wherein the reader picks a scenario that matches their situation, and follows all the steps in the scenario.

IMS-to-IMS TCP/IP connection operations

The communication path of an IMS-to-IMS TCP/IP connection passes through multiple IMS components: the IMS control region, the Structured Call Interface (SCI) of the Common Service Layer (CSL), IMS Connect, and either the IMS Multiple Systems Coupling (MSC) component or the IMS Open Transaction Manager Access (OTMA) component.

Operational tasks for an IMS-to-IMS TCP/IP connection in one segment of an IMS-to-IMS TCP/IP connection impacts the connection in the other segments. Also, general operational tasks in the IMS system are also likely to impact any IMS-to-IMS TCP/IP connections.

Related concepts:

“MSC operations” on page 267

Viewing configuration and status information for IMS-to-IMS TCP/IP connections in IMS Connect

You can view configuration and status information for IMS-to-IMS TCP/IP connections by issuing IMS Connect commands.

The connection information you can view is generally the same regardless of whether the connection is used for Multiple Systems Coupling (MSC) or Open Transaction Manager (OTMA). However, if you use MSC, you can also display information about the MSC link.

Viewing MSC connection information for IMS-to-IMS TCP/IP communications

You can view configuration and status information for IMS-to-IMS TCP/IP connections that are used for MSC by issuing IMS Connect commands.

Each MSC link requires two socket connections: a send socket connection and a receive socket connection. You display information about the send socket connections by querying RMTIMSCON connection information. You display information about the receive socket connection by querying the port information.

You can also display information about the definition and status of an MSC logical link by querying the MSC configuration information in IMS Connect.

- To display connection information on the sending side of a connection that is used for MSC, issue any one of the following IMS Connect commands that display the RMTIMSCON connection information:
 - In the IMS type-2 format, QUERY IMSCON TYPE(RMTIMSCON)
 - In the WTOR format, VIEWRMT
 - In the z/OS MODIFY format, QUERY RMTIMSCON
- To display connection information on the receiving side of a connection that is used for MSC, issue any one of the following IMS Connect commands that display port information:
 - In the IMS type-2 format, QUERY IMSCON TYPE(PORT)
 - In the WTOR format, VIEWPORT
 - In the z/OS MODIFY format, QUERY PORT
- To display link information on either side of a connection that is used for MSC, issue any one of the following IMS Connect commands that display link information:

- In the IMS type-2 format, QUERY IMSCON TYPE(MSC)
- In the WTOR format, VIEWMSC
- In the z/OS MODIFY format, QUERY MSC

Example: Viewing an MSC TCP/IP connection at the local and remote IMS Connect instances

In the following example, the VIEWRMT command is issued at a local IMS Connect instance for a RMTIMSCON connection, ICON2, that is used by MSC to send messages to the remote IMS installation.

```
nnVIEWRMT ICON2
```

The example output shows that two send socket connections, MSCBB435 and MSC84CF7, are in CONN status on the RMTIMSCON connection ICON2. The SENDCLNT IDs are auto-generated by the local IMS Connect instance when establishing the send socket connections to remote IMS Connect instance. In the example, PORT=5555 shows that the local IMS Connect expects the remote IMS Connect to be listening on port 5555. The presence of the LCLPLKID values indicate that these are MSC connections and identify the local MSC physical link that they are using.

```

HWSC0001I      RMTIMSCON=ICON2      STATUS=ACTIVE
HWSC0001I      IP-ADDRESS=009.030.221.055  PORT=5555
HWSC0001I      HOSTNAME=ECSER14.VMEC.SVL.IBM.COM
HWSC0001I      AUTOCONN=N  PERSISTENT=Y
HWSC0001I      IDLETO=0
HWSC0001I      RESVSOC=2  NUMSOC=2
HWSC0001I      SENDCLNT LCLPLKID STATUS          SECOND SENDPORT
HWSC0001I      MSCBB435 MSC12  CONN              6 1028
HWSC0001I      MSC84CF7 MSC12  CONN              56 1027
HWSC0001I      TOTAL SENDCLNTS=2  RECV=0  CONN=2  XMIT=0  OTHER=0

```

In the following example, the VIEWPORT command is issued at the local IMS Connect instance for port 9999.

```
nnVIEWPORT 9999
```

The example output shows the receive socket connections, MSCC73E0 and MSC0EBB0, that are in RECV state on port 9999 at the local IMS Connect instance. Again, the presence of the LCLPLKID values indicate that these connections are being used for MSC, as well as the local MSC physical link that they are using.

```

HWSC0001I      PORT=9999      STATUS=ACTIVE      KEEPAV=0  NUMSOC=3      EDIT=      TIMEOUT=0
HWSC0001I      CLIENTID LCLPLKID STATUS          SECOND CLNTPORT IP-ADDRESS
HWSC0001I      MSCC73E0 MSC12  RECV              6 1026      0:0:0:0:FFFF:91E:DD37
HWSC0001I      MSC0EBB0 MSC12  RECV              56 1025      0:0:0:0:FFFF:91E:DD37
HWSC0001I      TOTAL CLIENTS=2  RECV=2  READ=0  CONN=0  XMIT=0  OTHER=0

```

The following examples for a remote IMS Connect instance correspond to the preceding examples for a local IMS Connect instance.

In the following example, the VIEWPORT command is issued for port 5555 at the remote IMS Connect instance.

```
nnVIEWPORT 5555
```

The output for the VIEWPORT command shows two receive socket connections, MSCBB435 and MSC84CF7, on port 5555. These CLIENTID values match the SENDCLNT values shown in the preceding example of the VIEWRMT command at the local IMS Connect instance.

```

| HWSC0001I  PORT=5555      STATUS=ACTIVE  KEEPAV=0 NUMSOC=3      EDIT=      TIMEOUT=0
| HWSC0001I  CLIENTID LCLPLKID STATUS      SECOND CLNTPORT IP-ADDRESS
| HWSC0001I  MSCBB435 MSC21      RECV        18 1028      0:0:0:0:FFFF:91E:7326
| HWSC0001I  MSC84CF7 MSC21      RECV        69 1027      0:0:0:0:FFFF:91E:7326
| HWSC0001I  TOTAL CLIENTS=2 RECV=2 READ=0 CONN=0 XMIT=0 OTHER=0

```

In the following example, the VIEWRMT command is issued at the remote IMS Connect instance for RMTIMSCON connection, ICON1.

```
nnVIEWRMT ICON1
```

The example output shows that two send socket connections, MSCC73E0 and MSC0EBB0 under SENDCLNT, are in CONN status on the RMTIMSCON connection ICON1. These SENDCLNT values match the CLIENTID values shown in the preceding example of the VIEWPORT command at the local IMS Connect instance. In the example, PORT=9999 shows that the remote IMS Connect expects the local IMS Connect to be listening on port 9999. The presence of the LCLPLKID values indicate that these are MSC connections and identify the local MSC physical link that they are using.

```

| HWSC0001I  RMTIMSCON=ICON1 STATUS=ACTIVE
| HWSC0001I  IP-ADDRESS=009.030.115.038 PORT=9999
| HWSC0001I  HOSTNAME=ECSER13.VMEC.SVL.IBM.COM
| HWSC0001I  AUTOCONN=N PERSISTENT=Y
| HWSC0001I  IDLET0=0
| HWSC0001I  RESVSOC=2 NUMSOC=2
| HWSC0001I  SENDCLNT LCLPLKID STATUS      SECOND SENDPORT
| HWSC0001I  MSCC73E0 MSC21      CONN        18 1026
| HWSC0001I  MSC0EBB0 MSC21      CONN        69 1025
| HWSC0001I  TOTAL SENDCLNTS=2 RECV=0 CONN=2 XMIT=0 OTHER=0

```








Related concepts:

“Displaying information about an MSC network” on page 276

Related tasks:

“Viewing MSC link information in IMS Connect”

Related reference:

-  QUERY IMSCON commands (Commands)
-  VIEWPORT command (Commands)
-  VIEWRMT command (Commands)
-  VIEWMSC command (Commands)
-  IMS Connect QUERY MSC command (Commands)
-  IMS Connect QUERY PORT command (Commands)
-  IMS Connect QUERY RMTIMSCON command (Commands)

Viewing MSC link information in IMS Connect

When an IMS Connect to IMS Connect connection is used for MSC, you can view configuration and status information about the MSC links by issuing an IMS Connect command.

You can display configuration and status information for MSC links by the MSC physical link that are defined to IMS Connect or by logical link name.

- To display connection information about the configuration and status of MSC physical links, issue any one of the following IMS Connect commands:
 - In the IMS type-2 format, QUERY IMSCON TYPE(MSC). This command also displays some information about the MSC links.

- In the WTOR format, VIEWMSC
- In the z/OS MODIFY format, QUERY MSC
- To display status information for MSC logical links, issue the following command:
 - In the IMS type-2 format, QUERY IMSCON TYPE(LINK)

Related concepts:

“Displaying information about an MSC network” on page 276

Related tasks:

“Viewing MSC connection information for IMS-to-IMS TCP/IP communications” on page 257

Related reference:

 [QUERY IMSCON commands \(Commands\)](#)

 [VIEWMSC command \(Commands\)](#)

 [IMS Connect QUERY MSC command \(Commands\)](#)

Viewing OTMA connection information for IMS-to-IMS TCP/IP communications

You can view configuration and status information for IMS-to-IMS TCP/IP connections that are used for OTMA by issuing IMS Connect commands.

The information that you can view for an IMS-to-IMS TCP/IP connection that is used for OTMA is different depending on whether you are viewing the information on the sending side of a connection or the receiving side of a connection. For this reason, the commands you use to display the connection information are also different on the sending and receiving sides of the connection.

- To display connection information on the sending side of a connection that is used for OTMA, issue any one of the following IMS Connect commands that display the RMTIMSCON connection information:
 - In the IMS type-2 format, QUERY IMSCON TYPE(RMTIMSCON)
 - In the WTOR format, VIEWRMT
 - In the z/OS MODIFY format, QUERY RMTIMSCON
- To display connection information on the receiving side of a connection that is used for OTMA, issue any one of the following IMS Connect commands that display port information:
 - In the IMS type-2 format, QUERY IMSCON TYPE(PORT)
 - In the WTOR format, VIEWPORT
 - In the z/OS MODIFY format, QUERY PORT

Example: IMS Connect to IMS Connect connections for OTMA

In the following example, the VIEWRMT command is issued at a sending IMS Connect instance for a RMTIMSCON connection, ICON2B.

```
nnVIEWRMT ICON2B
```

The example output shows that there is one active send client socket connection, OTM924FA, on the RMTIMSCON connection ICON2B. The SENDCLNT ID is auto-generated by the sending IMS Connect instance when establishing the connection with the receiving IMS Connect instance. In the example, PORT=5555 shows that the sending IMS Connect expects the receiving IMS Connect to be listening on port 5555.


```

| HWSC0001I  RMTIMSCON=ICON2B    STATUS=ACTIVE
| HWSC0001I  IP-ADDRESS=009.030.221.055  PORT=5555
| HWSC0001I  HOSTNAME=ECSER14.VMEC.SVL.IBM.COM
| HWSC0001I  AUTOCONN=N  PERSISTENT=Y
| HWSC0001I  IDLETO=3000
| HWSC0001I  RESVSOC=4    NUMSOC=1
| HWSC0001I  SENDCLNT USERID  STATUS      SECOND SENDPORT
| HWSC0001I  OTM924FA APOL1    CONN        5941 1026
| HWSC0001I  TOTAL SENDCLNTS=1 RECV=0 CONN=1 XMIT=0 OTHER=0

```

In the following example, which corresponds to the preceding example, the VIEWPORT command is issued at a receiving IMS Connect instance for port 5555.

```
nnVIEWPORT 5555
```

The example output shows that there is one receive socket connection, OTM924FA, in RECV state on port 5555. The CLIENTID OTM924FA of the socket connection matches the SENDCLNT ID shown in the example command output for the sending IMS Connect instance.

```

| HWSC0001I  PORT=5555    STATUS=ACTIVE    KEEPAV=0 NUMSOC=2    EDIT=      TIMEOUT=0
| HWSC0001I  CLIENTID USERID TRANCODE DATASTORE STATUS  SECOND CLNTPORT IP-ADDRESS      APSB-TOKEN
| HWSC0001I  OTM924FA APOL1  APOL11  IMS2    RECV      11 1026    0:0:0:0:0:FFFF:91E:7326
| HWSC0001I  TOTAL CLIENTS=1 RECV=1 READ=0 CONN=0 XMIT=0 OTHER=0

```

Related reference:

- [➞ QUERY IMSCON commands \(Commands\)](#)
- [➞ VIEWPORT command \(Commands\)](#)
- [➞ VIEWRMT command \(Commands\)](#)
- [➞ IMS Connect QUERY PORT command \(Commands\)](#)
- [➞ IMS Connect QUERY RMTIMSCON command \(Commands\)](#)

Stopping a connection to a remote IMS Connect instance for IMS-to-IMS TCP/IP communications

You can stop an IMS-to-IMS TCP/IP connection to a remote IMS Connect instance by issuing an IMS Connect command from the local IMS Connect.

Depending on whether the connection is used for Multiple Systems Coupling (MSC) or Open Transaction Manager (OTMA), the implications of stopping an IMS-to-IMS TCP/IP connection to a remote IMS Connect instance are different.

To stop an IMS-to-IMS TCP/IP connection to a remote IMS Connect instance:

Issue any one of the following IMS Connect commands to the local IMS Connect instance:

- In the IMS type-2 format, UPDATE IMSCON TYPE(RMTIMSCON) NAME(*rmtimscon_name*) STOP(COMM)
- In the WTOR format, STOPRMT
- In the z/OS MODIFY format, UPDATE RMTIMSCON STOP(COMM)

IMS Connect issues message HWST3505I after the connection is terminated.

If you are stopping a connection to a remote IMS Connect instance that is used for MSC, the following actions occur when the connection is stopped:

- The local IMS Connect issues an informational message to the console and notifies MSC that the connection was stopped.

- IMS stops the MSC logical links that use the IMS-to-IMS TCP/IP connection.
- The local IMS Connect disconnects the TCP/IP sockets associated with the IMS-to-IMS TCP/IP connection.
- If the local IMS Connect receives any messages from MSC for the stopped connection, IMS Connect rejects the message and returns an error message to MSC.

If you are stopping a connection to a remote IMS Connect instance that is used for OTMA, the following actions occur when the connection is stopped:

- The local IMS Connect disconnects the TCP/IP sockets associated with the IMS-to-IMS TCP/IP connection.
- If an acknowledgement (ACK) response is pending on any socket that uses the connection, the local IMS Connect returns a negative acknowledgement (NAK) to OTMA in the local IMS system. IMS Connect issues message HWST3570E.
- If the local IMS Connect receives any OTMA messages from the local IMS system for the stopped connection, IMS Connect returns a NAK to OTMA with sense code 002A/0008 and issues message HWST3575W. OTMA leaves the message at the front of the output queue.

Related reference:

 [IMS Connect commands \(Commands\)](#)

 [UPDATE IMSCON commands \(Commands\)](#)

Restarting a connection to a remote IMS Connect instance for IMS-to-IMS TCP/IP communications

You can restart a stopped IMS-to-IMS TCP/IP connection in a local IMS Connect instance by issuing an online IMS Connect command.

Depending on whether the connection is used for Multiple Systems Coupling (MSC) or Open Transaction Manager (OTMA), the implications of restarting an IMS-to-IMS TCP/IP connection from IMS Connect are different.

If you are restarting a connection that is used for MSC, MSC communication does not resume until the MSC physical link is also started in the local IMS Connect instance.

If you are restarting a connection that is used for OTMA, the following actions occur when the connection is started:

- If the IMS-to-IMS TCP/IP connection is defined locally with AUTOCONN=Y, the local IMS Connect establishes socket connections with the remote IMS Connect instance. The number of socket connections opened is determined by the RESVSOC parameter in the RMTIMSCON configuration statement.
- The local IMS Connect notifies the local OTMA that the connection has been restarted and OTMA resumes sending messages on the connection.

To restart an IMS-to-IMS TCP/IP connection from the local IMS Connect instance:

1. Issue any one of the following IMS Connect commands:
 - In the IMS type-2 format, UPDATE IMSCON TYPE(RMTIMSCON) NAME(*rmtimscon_name*) START(COMM)
 - In the WTOR format, STARTRMT
 - In the z/OS MODIFY format, UPDATE RMTIMSCON START(COMM)

2. If the connection is used for MSC, start the MSC physical link in the local IMS Connect by issuing any one of the following IMS Connect commands:
 - In the IMS type-2 format, UPDATE IMSCON TYPE(MSC) NAME(*lclPlkid*) START(COMM)
 - In the WTOR format, STARTMSC *lclplk_id*
 - In the z/OS MODIFY format, UPDATE MSC NAME(*lclPlkid*) START(COMM)

IMS Connect issues message HWS3500I after the connection is restarted.

Related reference:

 IMS Connect commands (Commands)

 UPDATE IMSCON commands (Commands)

Stopping IMS Connect send client sockets on IMS-to-IMS TCP/IP connections

You can stop IMS Connect send client socket connections on an IMS-to-IMS TCP/IP connection in a local IMS Connect by issuing an IMS Connect command, such as the WTOR format command STOPSCLN.

Usually, the STOPSCLN command and the equivalent IMS type-2 and z/OS MODIFY commands are used only to stop send client socket connections that are used for IMS Open Transaction Manager Access (OTMA).

Recommendation: Do not stop IMS Connect send client socket connections if they are used for Multiple Systems Coupling (MSC). If you need to clean up sockets that are used for MSC, use the STOPLINK command or the equivalent IMS type-2 and z/OS MODIFY command.

In IMS Connect, the send client socket connection is represented by a client ID that is automatically generated by IMS Connect when the socket connection is established.

When you stop a send client socket that is connected but idle, as indicated by a CONN status in IMS Connect, the following actions occur:

- IMS Connect disconnects the send client socket and cleans up the associated control blocks.
- IMS Connect notifies the remote IMS Connect and the remote IMS Connect disconnects the related socket and cleans up its control blocks.

When you stop a send client that is waiting for an acknowledgement from the remote IMS Connect, as indicated by a RECV status in IMS Connect, the following actions occur:

- IMS Connect disconnects the send client socket and cleans up the associated control blocks.
- IMS Connect issues a negative acknowledgement (NAK) to OTMA that direct OTMA to reroute unacknowledged message to the IMS Connect dead letter queue HWS\$DLQ.
- IMS Connect issues message HWST3570E to the console.
- When the remote IMS Connect attempts to return the acknowledgement, the remote IMS Connect receives a socket error and disconnects the related socket and cleans up its control blocks.

To stop an IMS Connect send client socket of an IMS-to-IMS TCP/IP connection:

Issue any one of the following IMS Connect commands in the local IMS Connect instance:

- In the IMS type-2 format, UPDATE IMSCON TYPE(SENDCLNT) NAME(*sendclient_name*) RMTIMSCON(*rmtimscon_name*) STOP(COMM)
- In the WTOR format, STOPSCLN
- In the z/OS MODIFY format, DELETE RMTIMSCON NAME (*rmtimscon*) SENDCLNT(*clientid*)

IMS Connect issues message HWS3525I after the connection is terminated.

Related reference:

 IMS Connect commands (Commands)

 UPDATE IMSCON commands (Commands)

Stopping IMS Connect communication with an IMSplex when MSC TCP/IP links are supported

You can stop communications between a local IMS Connect instance and a local IMSplex by issuing an IMS Connect command, such as the WTOR format command STOPIP.

Stopping IMSplex communications breaks all MSC TCP/IP physical links that are supported by the local IMS Connect in the local IMSplex.

When a STOPIP or similar command is issued, the following actions occur for MSC TCP/IP links:

- The local IMS Connect disconnects the parallel send and receive socket that support the physical link and sends a shutdown directive to MSC in the local IMS system.
- When MSC receives the shutdown directive, MSC cleans up the physical link and issues message DFS3176E.
- IMS Connect changes the status of the MSC TCP/IP physical link to DISCONNECTED.

To stop communication between a local IMS Connect instance and a local IMSplex from the local IMS Connect instance issue any one of the following IMS Connect commands:

- In the IMS type-2 format, UPDATE IMSCON TYPE(IMSPLEX) NAME(*imsplex_name*) STOP(COMM)
- In the WTOR format, STOPIP *imsplex_id*
- In the z/OS MODIFY format, UPDATE IMSPLEX NAME (*imsplex_name*) STOP(COMM)

Related reference:



IMS Connect commands (Commands)



UPDATE IMSCON commands (Commands)

Starting IMS Connect communication with an IMSplex when MSC TCP/IP links are supported

You can start communications between a local IMS Connect and a local IMSplex by issuing an IMS Connect command, such as the WTOR format command STARTIP.

When a STARTIP or similar command is issued, communication resumes between IMS Connect and the IMSplex.

IMS Connect changes the status of the MSC TCP/IP links that it supports in the IMSplex to ACTIVE.

To start communication between a local IMS Connect and a local IMSplex from the local IMS Connect instance issue any one of the following IMS Connect commands:

- In the IMS type-2 format, UPDATE IMSCON TYPE(IMSPLEX) NAME(*imsplex_name*) START(COMM)
- In the WTOR format, STARTIP *imsplex_id*
- In the z/OS MODIFY format, UPDATE IMSPLEX NAME (*imsplex_name*) START(COMM)

Related reference:



IMS Connect commands (Commands)



UPDATE IMSCON commands (Commands)

Cleaning up MSC logical link resources in IMS Connect

If an MSC logical link has terminated in IMS, but the IMS Connect resources associated with the logical link did not clean up automatically, you can clean up the resources by using an IMS Connect command, such as the WTOR format command STOPLINK.

Recommendation: Under normal circumstances, use the IMS command /PSTOP to terminate MSC logical links. Stop MSC links in IMS Connect only when the IMS Connect resources associated with an MSC logical link that has already terminated have failed to clean up correctly.

A logical link on one MSC physical link can have the same name as a logical link on another physical link. To ensure that a logical link with a duplicate name is not stopped on another physical link, include the ID of the physical link that the logical link is assigned to on the command when you stop a logical link. The ID of the physical link can be found on the LCLPLKID parameter of the MSC configuration statement.

When you clean up MSC logical links in a local IMS Connect instance by using a command like STOPLINK, IMS Connect takes the following actions:

- Stops any communication on the MSC logical link
- Notifies the local IMS that communication has stopped on the logical link
- Deletes the control blocks associated with the logical link and frees the associated storage

- Issues message HWSF3310I

To stop an MSC logical link and clean up its associated resources in a local instance of IMS Connect issue any one of the following IMS Connect commands in the local IMS Connect instance:

- In the IMS type-2 format, UPDATE IMSCON TYPE(LINK) NAME(*linkname*) MSC(*lclplk_id*) STOP(COMM)
- In the WTOR format, STOPLINK *linkname lclPlkid*
- In the z/OS MODIFY format, DELETE LINK NAME (*linkname*) LCLPLKID(*lclPlkid*)

IMS Connect issues message HWS3310I after the link is stopped.

Related reference:

 [IMS Connect commands \(Commands\)](#)

 [UPDATE IMSCON commands \(Commands\)](#)

Stopping an MSC physical link in IMS Connect

Do not stop an MSC physical link in IMS Connect, unless you need to clear the affinity of the link with an IMS system when TCP/IP generic resources are used.

Recommendation: Under normal circumstances, use the IMS command /PSTOP to terminate MSC physical links. Stop MSC physical links in IMS Connect only when you need to clear the affinity of an MSC physical link with a particular IMS system that uses TCP/IP generic resources.

When you stop an MSC physical link in a local IMS Connect instance by using a command like STOPMSC, IMS Connect takes the following action:

- Stops communication on the specified MSC physical link, including communications on all the MSC logical links that are assigned to the physical link.
- Informs IMS that communication has stopped on the physical link so that IMS can also terminate the physical link and any logical links that are assigned to the physical link.
- Changes the status of the MSC physical link and its assigned logical links to NOT ACTIVE.
- For TCP/IP generic resources, clears affinity of a physical link to the IMS system
- Issues message HWSF3305I

Use the VIEWMSC command or the QUERY MSC command to display information about MSC physical links that are defined to IMS Connect.

To stop an MSC physical link in a local instance of IMS Connect:

Issue any one of the following IMS Connect commands in the local IMS Connect instance:

- In the IMS type-2 format, UPDATE IMSCON TYPE(MSC) NAME(*lclPlkid*) STOP(COMM)
- In the WTOR format, STOPMSC *lclPlkid*
- In the z/OS MODIFY format, UPDATE MSC NAME (*lclPlkid*) STOP(COMM)

IMS Connect issues message HWS3310I after the link is stopped.

Related concepts:

➡ MSC TCP/IP generic resources (Communications and Connections)

Related reference:

➡ IMS Connect commands (Commands)

➡ UPDATE IMSCON commands (Commands)

MSC operations

In a non-sysplex environment, each IMS system in a Multiple Systems Coupling (MSC) configuration is operationally an independent unit. Each IMS system exclusively owns its own communication resources, and is controlled by its own master terminal.

Related concepts:

“IMS-to-IMS TCP/IP connection operations” on page 257

Related tasks:

“Modifying Multiple Systems Coupling resources” on page 35

MSC initialization

To start Multiple Systems Coupling (MSC) communications between two IMS systems, issue either the IMS type-1 command `/RSTART LINK` or the IMS type-2 command `UPDATE MSLINK NAME(linkname) START(COMM)`.

If the physical link type is either a channel-to-channel connection or a main-storage to main-storage connection (MTM), you must issue the command in both IMS systems at each end of the link.

If the physical link type is either TCP/IP or VTAM, you need to issue the command in only one IMS system. The normal procedure is for the operator to issue this command when a system has started. Communication is allowed only if the characteristics of the specified links are compatible. If a required link is not successfully started, messages wait until the links are reassigned.

The SYSID of each IMS system running in a shared queues group can either be cloned across all systems, or not. If the SYSIDs are not cloned, initialize all IMS systems in the shared queues group before starting regions to process the transactions on the shared queue to avoid occurrences of pseudoabend U0830.

After all IMS systems have been initialized, transaction processing can begin, and individual IMS systems can be brought down and later restarted as needed. Initialization allows all IMS systems to exchange MSC SYSIDs and MSNAMEs, and creates a SYSID table that contains all of the SYSIDs for that shared queues group. If the SYSIDs for all IMS systems in a shared queues group *are* cloned, then you do not need to perform this initialization.

Attention: If you cold start a non-shared queues system that has messages queued, the queued messages are lost. Because the messages that were lost can be from or to terminals and programs in other systems, the impact of a cold start is not limited to the cold-started system.

In a shared-queues environment, IMS systems in the IMSplex exchange SYSIDs and MSNAMEs at initialization. During this SYSID exchange, each IMS system creates dynamic MSNAMEs for any MSNAMEs that are defined in other IMS

systems that it does not define itself. The dynamic MSNAMES result in paths to all of the IMS systems in the IMSplex. The SYSIDs are merged to create a common SYSID routing table. The table is the same in each IMS. Therefore, any local SYSIDs are local in all IMS systems and override any remote or undefined SYSIDs. Remote SYSIDs override only undefined SYSIDs.

MSC termination

You terminate Multiple Systems Coupling (MSC) links from either of the two linked IMS systems by issuing either a type-1 command `/PSTOP LINK` or a type-2 command `UPDATE MSLINK NAME(linkname) STOP(COMM)`. When transmission terminates on one side, its partner in the other system terminates its own transmission and notifies the MTO.

Use the command forms shown in the following table to terminate links between partner MSC systems.

Table 39. Commands to terminate links

Command	Keyword	MSC use and effect
/IDLE	LINK	IMS forces termination of all transmissions on the physical link associated with the named logical link. Use only after a shutdown checkpoint.
	NOSHUT	For TCP/IP and VTAM link types only, IMS forces the idle condition without taking a shutdown checkpoint.
/PSTOP	LINK	IMS halts transmissions associated with the logical link between two partner systems, but continues queuing for the remote resources. IMS enqueues, but does not send, broadcast messages that would use the link. IMS stops the logical link in the partner system, and sends message DFS2161 to its MTO.
	LINK PURGE	IMS forces PSTOP condition even if other system failed while channel-to-channel (CTC) I/O was in progress. CTC link type only.
	MSPLINK	Stops logons to an MSC TCP/IP or VTAM physical link and enables the operator to issue the /MSASSIGN command to reassign logical links to the physical link. Any links in sessions that have not been stopped by the /PSTOP command are not affected by an /MSASSIGN command. For physical links in a TCP/IP generic resource group, after the physical link is stopped, all logical links must terminate normally to clear link affinity before they can be moved to a different IMS system.

Table 39. Commands to terminate links (continued)

Command	Keyword	MSC use and effect
UPDATE MSLINK NAME(<i>linkname</i>)	STOP(COMM)	Halts transmissions associated with the logical link between two partner systems, but queuing for the remote resources continues. Broadcast messages that would use the link are queued but not sent. This command is equivalent to the type-1 command /PSTOP LINK. The logical link is stopped in the partner system and its MTO is notified by message DFS2161.
	OPTION(FORCE)	This command applies to TCP/IP and VTAM link types and forces a /PSTOP condition. OPTION(FORCE) leaves the link in PSTOPPED IDLE ERE mode. The subsequent restart of the link is an emergency restart. For VTAM link types, STOP(COMM) without OPTION(FORCE) must be issued before OPTION(FORCE) is used. For TCP/IP links, OPTION(FORCE) can be used any time. OPTION(FORCE) might need to be issued on both sides of the link.
	STOP(LOGON)	Stops logons to the physical link. This applies to MSC TCP/IP and VTAM links only. It stops logons to the physical link and enables the operator to reassign logical links to the physical link.
	STOP(GENLOGON)	For MSC TCP/IP links in a TCP/IP generic resource group only, stops logons to the IMS system on the physical link. After the command stops logons to the physical link and all logical links on the physical link have terminated normally to clear link affinity with the IMS system, the logical links can be reassigned to a physical link on another IMS system.

Changing logical link assignments

You define initial logical link assignments (logical link to physical link) as part of the IMS system definition process. However, you can change logical link assignments dynamically.

To dynamically make or change a logical link assignment, use either the type-1 command /MSASSIGN or the type-2 command UPDATE MSLINK NAME(*linkname*) SET(MSPLINK(*msplinkname*)).

Recommendation: Use the /MSASSIGN command or the UPDATE MSLINK NAME(*linkname*) SET(MSPLINK(*msplinkname*)) command to create logical links only for unscheduled reassignments resulting from failing physical connections or systems.

Because a logical link must always communicate with its partner, the operators for the two systems must coordinate their assignments of corresponding physical links. You can replace any type of physical link with any other type of physical link.

Changes to logical link assignments remain in effect until you change them using an /MSASSIGN command, UPDATE MSLINK NAME(*linkname*) SET(MSPLINK(*msplinkname*)) command, or cold start IMS. The commands only alter the relationship of MSC resources in the local system.

The logical link relationships you can change and the commands you can use to change them include:

- Logical link to physical link
 - Type-1 command /MSASSIGN.
 - Type-2 command UPDATE MSLINK NAME(*linkname*) SET (MSPLINK(*msplinkname*)).
- Remote SYSID to logical link
 - Type-1 command /MSASSIGN.
 - Type-2 command UPDATE MSNAME NAME(*linkname*) SET (MSPLINK(*msplinkname*)) assigns an MSNAME with the desired remote SYSID to a logical link. Alternately, you can modify the remote SYSID of the MSNAME already associated with the logical link by issuing the type-2 command UPDATE MSNAME NAME(*msname*) SET(SIDR (*remote_sysid*)).
- Logical link path to logical link:
 - Type-1 command /MSASSIGN.
 - Type-2 command UPDATE MSNAME NAME(*msname*) SET(MSLINK(*mslinkname*)).

You can also use the /MSASSIGN command to complete the following tasks:

- Change a remote program to local
- Change a local program to remote
- Assign a logical path

When you modify logical link assignments:

- Operating logical or physical links must be assigned one-to-one, except for TCP/IP and VTAM link types, which can have parallel sessions in effect. In this case, multiple logical links can be assigned to one physical link.
- If the link uses TCP/IP or VTAM, before you assign a logical link either to or from an MSPLINK name you must stop the physical link by executing one of the following commands:
 - The type-1 command /PSTOP MSPLINK
 - The type-2 command UPDATE MSPLINK NAME(*msplinkname*) STOP(LOGON)
 - In a TCP/IP generic resource group, the type-2 command UPDATE MSPLINK NAME(*msplinkname*) STOP(GENLOGON)
- Before reassigning a logical link, you must stop the logical link by using either a type-1 command /PSTOP or a type-2 command UPDATE MSLINK NAME(*linkname*) STOP(COMM) and the logical link must be idle.
- A destination SYSID cannot be reassigned to a logical link unless its currently assigned logical link is idle after a type-1 command /PSTOP or a type-2 command UPDATE MSLINK NAME(*linkname*) STOP(COMM).
- MSC communication can occur only when logical links in two IMS systems share the same partner identification and have assignment to an operating communications facility between the systems.
- IMS does not determine whether the requested logical link assignment is reasonable or results in a valid configuration for MSC communication. This is only done by communication with appropriate remote systems. You can accomplish this using the /MSVERIFY command after making assignment changes.

Restarting a logical link

If a restart is pending on a logical link because of a physical link failure, reestablish communications between both systems through an alternative physical link.

Restriction: In a TCP/IP generic resource group, where a physical link is defined in multiple IMS systems, if a logical link is in a PSTOPPED ERE state on any IMS system in the group, do not start a logical link that uses the same physical link on any other IMS system in the group. Before you can start a logical link on the same physical link on another system, you must clear the affinity of the ERE link either by shutting it down normally or by resetting it to a COLD state.

As a precaution, display the affinity status of an MSC link on its current IMS system before moving the link by restarting it on another IMS system. If the affinity is still active, shutdown the link normally, which resets the affinity status. If the link cannot be shut down normally, set the link to COLD status by issuing either the type-1 command `/CHANGE LINK linknum FORCSESS | SYNCSESS COLDSESS` or the type-2 command `UPDATE MSLINKNAME(linkname) SET(SYNCOPT(COLDSESS))`

Attention: Changing the link status from ERE to COLD or moving a link in ERE status and restarting it on another IMS system prevents the synchronization of the message sequence numbers during restart, which can cause the duplication or loss of messages.

1. If the link uses TCP/IP or VTAM, before you assign a logical link either to or from an MSPLINK name, stop the physical link by issuing one of the following commands:
 - The type-1 command `/PSTOP MSPLINK`
 - The type-2 command `UPDATE MSPLINK NAME(msplinkname) STOP(LOGON)`
 - For links in a TCP/IP generic resource group only, the type-2 command `UPDATE MSPLINK NAME(msplinkname) STOP(GENLOGON)`
2. Reassign the logical link to the alternate physical link by using either of the following commands:
 - The type-1 command `/MSASSIGN`
 - The type-2 command `UPDATE MSLINK NAME(linkname) SET(MSPLINK(msplinkname))`
3. If the link uses TCP/IP or VTAM, restart the physical link by issuing one of the following commands:
 - The type-1 command `/RSTART MSPLINK`
 - The type-2 command `UPDATE MSPLINK NAME(msplinkname) START(LOGON)`
 - For links in a TCP/IP generic resource group only, the type-2 command `UPDATE MSPLINK NAME(msplinkname) START(GENLOGON)`
4. Start the logical link by issuing one of the following commands:
 - The type-1 command `/RSTART LINK`
 - The type-2 command `UPDATE MSLINK NAME(linkname) START(COMM)`

Switching TCP/IP and VTAM physical link types

Because of the operational similarities between MSC physical links that use TCP/IP and those that use VTAM, the VTAM- and TCP/IP-type physical links are well-suited to serve as a back-up link types for each other.

Prerequisite: In advance of any failure, you must define a redundant set of physical links that use the alternative link type and that duplicate the primary set of MSC physical links. Both TCP/IP and VTAM physical links are defined by using MSPLINK system definition macros.

Any supporting components required by the backup physical link type must also be pre-configured and ready for use. For example, if TCP/IP is your backup link type, the Structured Call Interface (SCI) component of the Common Service Layer (CSL) and IMS Connect must also be configured and ready for use.

The following steps assume a scenario in which the primary physical links in use are VTAM physical links and the back up physical links are TCP/IP. However, the steps are also a general guide for switching from TCP/IP to VTAM.

The steps reassign logical links on a failed VTAM physical link between a front-end IMS system and a back-end IMS system to a backup TCP/IP physical link that was previously defined between the two systems.

In a failure scenario like this, any terminals on the front-end IMS system that were waiting for a response are likely hung and message DFS3222I has been issued to the front-end master terminal operator (MTO). On the back-end IMS system, the IMS applications have stopped receiving messages for processing and message DFS3222I has also been issued to the back-end MTO.

The operators at the front-end and back-end IMS systems must coordinate the timing of the steps to ensure that messages from the front-end IMS system do not start flowing before the backend IMS system is ready to receive them.

1. Optional: At the front-end IMS system, stop the remote transactions that route messages on the failed physical link. When the remote transactions are stopped, any terminal that submits requests for the remote transactions receive message DFS065.
2. At the back-end IMS system:
 - a. Confirm that a local instance of IMS Connect is active and appropriately configured to support the backup physical link from the back-end IMS system. For each MSC TCP/IP physical link that the IMS Connect instance supports, at least one RMTIMSCON statement must be defined and one MSC statement for each physical link must be defined in the IMS Connect configuration member in the IMS.PROCLIB data set.
 - b. Confirm that a local instance of SCI is active and available for use by the back-end IMS system and IMS Connect.
 - c. Reassign the logical links from the failed VTAM physical link to the TCP/IP physical link by issuing either of the following commands:
 - The IMS type-2 command `UPDATE MSLINK NAME(linkname) SET(plinkname)`
 - The IMS type-1 command `/MSASSIGN LINK xxx to MSPLINK plinkname`
3. At the front-end IMS system:
 - a. Confirm that a local instance of IMS Connect is active and appropriately configured to support the backup physical link from the front-end IMS system. For each MSC TCP/IP physical link that the IMS Connect instance supports, at least one RMTIMSCON statement must be defined and one MSC statement for each physical link must be defined in the IMS Connect configuration member in the IMS.PROCLIB data set.

- b. Confirm that a local instance of SCI is active and available for use by the front-end IMS system.
- c. Reassign the logical links that use the failed VTAM physical link to the TCP/IP physical link by issuing either of the following commands:
 - The IMS type-2 command UPDATE MSLINK NAME(*linkname*) SET(*plinkname*)
 - The IMS type-1 command /MSASSIGN LINK *xxx* to MSPLINK *plinkname*
- d. After the logical links have been reassigned to the same TCP/IP physical link at both the front-end and back-end IMS systems, start the logical links by issuing either of the following commands:
 - The IMS type-2 command UPDATE MSLINK NAME(*linkname*) START(COMM)
 - The IMS type-1 command /RSTART LINK *xxx*
- e. If the remote transactions were stopped, restart them.

Related concepts:


 MSC physical links (Communications and Connections)

“Recovery considerations for multiple systems” on page 279

Related tasks:

 Defining IMS-to-IMS TCP/IP connections for MSC (System Definition)

Related reference:

 MSPLINK macro (System Definition)

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

MSC TCP/IP link operations

Operating procedures for MSC TCP/IP links are generally the same as the operating procedures for MSC VTAM links; however, a few differences do exist.

IMS Connect manages the TCP/IP connections that are required for MSC TCP/IP. Communication between MSC and IMS Connect is managed by the Structure Call Interface (SCI) component of the IMS Common Service Layer (CSL).

Under normal circumstances, the use of TCP/IP, IMS Connect, and SCI is not apparent during the operation of an MSC TCP/IP link. MSC TCP/IP links are started and stopped by using the same IMS commands that are used to initialize and terminate VTAM-type links. When you issue an UPDATE MSLINK NAME(*linkname*) START(COMM) command for a TCP/IP link in a local IMS system, the link is started in both the local and remote IMS and IMS Connect pairs.

IMS Connect also provides commands to control MSC TCP/IP links. Although you can stop a TCP/IP link by using IMS Connect commands, only do when a link has failed to clean up properly in IMS Connect after terminating in IMS.

Other differences between the operating procedures for MSC TCP/IP links and the other MSC link types are found in the error recovery procedures and in the management of affinities when TCP/IP generic resources are used.

Related concepts:

“Monitoring IMS Connect connections” on page 103

➡ MSC and IMS-to-IMS TCP/IP communications (Communications and Connections)

Related tasks:

➡ IMS-to-IMS TCP/IP connections (System Definition)

Related reference:

➡ /PSTOP command (Commands)

➡ QUERY IMSCON commands (Commands)

➡ QUERY MSLINK command (Commands)

➡ QUERY MSPLINK command (Commands)

➡ /RSTART command (Commands)

➡ UPDATE MSLINK command (Commands)

➡ UPDATE MSPLINK command (Commands)

➡ IMS Connect WTOR commands (Commands)

➡ IMS Connect z/OS commands (Commands)










Restarting MSC TCP/IP links after errors

If an IMS system or an MSC TCP/IP link terminates unexpectedly, you need to confirm that the MSC TCP/IP link has also been cleaned up properly in the remote IMS system.

In the event of a failure, IMS attempts to clean up the MSC TCP/IP links at both the local and remote site; however, if the failing IMS system is unable to notify the remote site of the failure, the remote site might not clean up the link.

- To determine if a link has been cleaned up in the remote IMS system, you can issue either the type-1 format /DISPLAY LINK command or the type-2 format QUERY MSLINK NAME(*linkname*) command.
- After an MSC TCP/IP link terminates at a local site, if the remote site does not also clean up the link, issue a /PSTOP LINK FORCE or an UPDATE MSLINK STOP(COMM) OPTION(FORCE) in the remote IMS system. If the link is restarted in the local IMS system before either the link or socket connection is cleaned up at the remote site, the link or socket connection is cleaned up at the remote site during restart processing.
- To restart a TCP/IP link after an IMS system is started, issue either an /RSTART command or the UPDATE MSLINK START(COMM) command.
- To set a stopped MSC TCP/IP link to its coldstart status, issue either the type-1 format command /CHANGE LINK COLDSESS or the type-2 format command UPDATE MSLINK SET (SYNCOPT(COLDSESS)). The link must be stopped first. When specifying the COLDSESS keyword, the SYNCSESS or FORCSESS keyword must also be specified.

Related reference:

-  [/CHANGE LINK command \(Commands\)](#)
-  [QUERY IMSCON commands \(Commands\)](#)
-  [QUERY MSLINK command \(Commands\)](#)
-  [QUERY MSPLINK command \(Commands\)](#)
-  [/RSTART command \(Commands\)](#)
-  [UPDATE MSLINK command \(Commands\)](#)
-  [UPDATE MSPLINK command \(Commands\)](#)
-  [IMS Connect WTOR commands \(Commands\)](#)
-  [IMS Connect z/OS commands \(Commands\)](#)

Commands that help control resources in an MSC environment

Uses and effects of Multiple System Coupling (MSC) commands are described so that you can assess their effects in your MSC environment.

If an MSC error occurs when the MSC trace is not operational, check the log for X'67' records, because some error information is logged even when the link trace is not set on.

When you suspect link problems, you can use either of the following two commands to start a trace:

- Type-1 command /TRA SET ON LINK X LEVEL 3 MODULE ALL.
- Type-2 command UPDATE MSLINK NAME(mslinkname) START(TRACE). This command uses the same level and module settings that were used the last time the /TRACE SET (ON) LINK command was issued. If a /TRACE SET (ON) LINK command has not been issued since the last cold start, this command defaults to MODULE=ALL and LEVEL=4.

Table 40. MSC environment commands

Command	MSC use and effect
/ASSIGN	Change transaction priorities.
/BROADCAST	Use of the LTERM keyword value ALL sends the broadcast message to all terminals of the local system only. No general broadcast capability exists for remote MSC terminals. Their LTERMs must be defined in the broadcasting system and specified in the command. Broadcasts to MTOs of remote MSC systems are possible even when the MTO terminals are not defined in the broadcasting system; use the keyword MASTER. The system must have defined the SYSIDs of such remote systems by MSNAME macros.
/CHANGE	Change link mode table and session restart.
/PSTOP	Queue for remote processing, but do not send.
/PURGE	Reject subsequent primary requests for remote processing. Accept secondary or continued conversational requests. Continue sending those requests that are allowed to queue.

Table 40. MSC environment commands (continued)

Command	MSC use and effect
/STOP	If remote program or terminal stopped in input system: Return DFS065 to input terminal if a primary request. Queue but do not send secondary requests. If local program or terminal accessed by remote MSC systems: Queue secondary requests. Return message DFS065 to input terminal in input system for primary requests. Message DFS065 identifies, by SYSID, the MSC processing system that rejected the message. Message is logged, but canceled.
/START	Reset previous /START and /PSTOP effects.
/TRACE	Using the LINK keyword, you can invoke a trace of MSC operation. Trace data is recorded within type X'67' records on the IMS system log.
UPDATE MSLINK	Using the UPDATE MSLINK, you can set or change various logical link attributes, including starting and stopping the logical link and starting and stopping traces of the logical link. The command requires the Operations Manager (OM) API. The command syntax for this command is defined in XML and is available to automation programs that communicate with OM. Using the UPDATE command, you can also modify MSLINK system definition parameter values BUFSIZE=, MODETBL=, PARTNER (ID)=, and the MSLINK name (MSLINK label field).
UPDATE MSNAME	Using the UPDATE MSNAME, you can set or change various logical link path attributes, including starting and stopping the queuing or sending of messages to logical link path. The command requires the Operations Manager (OM) API. The command syntax for this command is defined in XML and is available to automation programs that communicate with OM. Using the UPDATE command, you can also modify MSNAME system definition parameter value SYSID=.
UPDATE MSPLINK	Using the UPDATE MSPLINK, you can set or change various physical link attributes, including preventing or allowing logons to the physical link. The command requires the Operations Manager (OM) API. The command syntax for this command is defined in XML and is available to automation programs that communicate with OM. Using the UPDATE command, you can also modify MSPLINK system definition parameter values MODETBL=, and the physical link name (MSPLINK label field).

Displaying information about an MSC network

There are two commands that you can use to view information about an MSC network: the type-1 command /DISPLAY and the type-2 command QUERY.

/DISPLAY command and MSC

The IMS type-1 /DISPLAY commands show available information that helps you manage an MSC environment. These facilities operate only within the domain of the local system, however, and do not provide information that is available only from remote systems. The following table summarizes the information that is available to the MTO with the /DISPLAY command.

Table 41. MSC information returned by the /DISPLAY command.

Command	Keyword	MSC use and effect
/DISPLAY AFFIN	LINK	Displays the current link affinities that logical links have within a TCP/IP or VTAM generic resource group. Also displays the name and number of the logical links and for VTAM links, the node.

Table 41. MSC information returned by the /DISPLAY command (continued).

Command	Keyword	MSC use and effect
/DISPLAY ASMT	LINK	The physical link, the local and remote system IDs, and logical link paths assigned to the specified logical link.
	MSNAME	Logical link number, physical link, local and remote system IDs, MSNAME label, and IMS ID.
	MSPLINK	Logical link number, physical link, physical link type (CTC, MTM, TCP/IP, VTAM), physical link address, the maximum number of allowed sessions, and the identifier used by the remote IMS system. For TCP/IP physical links, the command output displays the IMS ID of the remote IMS system. For VTAM physical links, the output displays the VTAM node name of the remote IMS system.
	SYSID	The physical link, logical link, and logical link path assignments associated with the specified system identification.
/DISPLAY LINK	[None]	Logical link number.
		Partner ID.
		The local totals for the numbers of messages received, sent, enqueued, dequeued, currently queued for link, and link status.
	MODE	For VTAM links, displays various modes in which VTAM terminals can operate.
	OPTION BUFSIZE	Displays the link number, link name, bandwidth, buffer size, and the link status.
/DISPLAY MSNAME	QCNT	The global totals for the numbers of messages received, sent, enqueued, dequeued, currently queued for link, and link status.
	[None]	Logical link path name (MSNAME) and the local totals for the numbers of messages enqueued, dequeued, and currently queued to the MSNAME.
		Logical link path name (MSNAME) and the global totals for the numbers of messages enqueued, dequeued, and currently queued to the MSNAME.

QUERY command and MSC

The IMS type-2 QUERY commands display available attribute and status information to help you manage an MSC environment. The following table summarizes the MSC information returned by the QUERY command.

Table 42. MSC information returned by the QUERY command.

Command	Keyword	MSC use and effect
QUERY MSLINK	BANDWIDTH(ON OFF)	Identifies logical links that are or are not in bandwidth mode.
	SHOW(ALL attribute)	Returns the user-defined attributes of logical links. You can limit the information returned by specifying a valid attribute as a parameter of the SHOW() keyword.
	STATUS(status_type)	Identifies the logical links that have the specified status type.

Table 42. MSC information returned by the QUERY command (continued).

Command	Keyword	MSC use and effect
QUERY MSNAME	QCNT()	Identifies logical link paths that have queue counts relative to a number you specify. For example, you can view all logical link paths with a queue count that is greater-than-or equal-to 100.
	SHOW(ALL attribute)	Returns the user-defined attributes of logical link paths. You can limit the information returned by specifying a valid attribute as a parameter of the SHOW() keyword.
	STATUS(status_type)	Identifies the logical link paths that have the specified status type.
QUERY MSPLINK	SHOW(ALL attribute)	Returns the user-defined attributes of physical links. You can limit the information returned by specifying a valid attribute as a parameter of the SHOW() keyword.
	STATUS(status_type)	Identifies the physical links that have the specified status type.
	TYPE(link_type)	Identifies physical links by their link types: CTC, MTM, TCP/IP, and VTAM. You can query one or more link types at a time.

Related tasks:

“Viewing MSC connection information for IMS-to-IMS TCP/IP communications” on page 257

“Viewing MSC link information in IMS Connect” on page 259

Logical link path control

You can use several commands to control logical link paths. The logical link path is the lowest level of control across a Multiple Systems Coupling (MSC) environment because it is the lowest level that is necessarily defined in intermediate MSC systems.

Logical link paths are defined by SYSID pairs that identify the sending and destination systems. You define logical link paths using the MSNAME system definition macro. The following table summarizes the commands you can use for link control.

Table 43. Commands used to control MSC link paths

Command	Keyword	MSC use and effect
/START	MSNAME	Start a previously stopped MSNAME.

Table 43. Commands used to control MSC link paths (continued)

Command	Keyword	MSC use and effect
/STOP	MSNAME	<p>Stop the sending or receiving of primary request messages associated with the logical link path.</p> <p>When stopped by an input system, primary requests for remote programs or terminals associated with the stopped logical link path are canceled and IMS returns message DFS065 to the input terminal. Conversations in progress are allowed to continue.</p> <p>When stopped by a destination system, messages received from other systems over a stopped logical link path cause a logical link path to be stopped in the sending system (input or intermediate); IMS issues messages DFS2140 and DFS2142 (respectively) to the MTOs of the sending and receiving systems. Messages remain enqueued in the sending system until the logical link is subsequently started in both systems.</p>
/PURGE	MSNAME	<p>Halt enqueueing of primary requests for all remote terminals and programs represented by the MSNAME. Continuing conversations and secondary requests are still handled. Primary requests entered through an input terminal receive message DFS065. Requests from other systems that require use of the logical link path for a response are not accepted, but remain enqueued in the sending system (see /STOP MSNAME above).</p>
UPDATE MSNAME	SET(), START(), and STOP()	<p>Using the UPDATE MSNAME command, you can set or change various logical link path attributes, including starting and stopping the queuing or sending of messages to logical link path. The command requires the Operations Manager (OM) API. The command syntax for this command is defined in XML and is available to automation programs that communicate with OM.</p>

Recovery considerations for multiple systems

Each system in the multisystem configuration uses the full recovery capabilities of IMS. These capabilities assure that messages are not lost or duplicated within the single system.

You will not lose messages under the following conditions:

- You do not cold start the subsystem or emergency restart (BUILDQ) if using an earlier checkpoint.
- No log records are lost.

VTAM provides message integrity for SDLC links in addition to the MSC control functions.

Related tasks:

“Switching TCP/IP and VTAM physical link types” on page 271

Message recovery

IMS assures that messages are not lost or duplicated across a Multiple Systems Coupling (MSC) link by logging information about a message in both the sending and receiving systems.

IMS restores this information during restart and exchanges it between the systems after the link is established. The sending system can then dequeue a message that was received by the receiving system but for which the acknowledgment was lost due to a link or system failure. The sending system can also re-send a message that

was sent but not enqueued by the receiving system. If an IMS subsystem fails to recover, the messages for which it has recovery responsibility are lost.

Because IMS provides commands to dynamically change link assignments, you can set up an alternate processor for an inoperable one. The IMS system that resides in the inoperable processor can run in the alternate processor after all involved links are properly reassigned by the master terminal operators (MTOs).

Stopped transactions

If IMS stops a destination transaction, the action taken by the destination system varies based on the type of request. Requests can be conversational or nonconversational.

For a primary request that is not conversational or that starts a conversation, IMS sends an error message to the input terminal and cancels the message.

For a primary request that continues a conversation or a secondary request, IMS enqueues the message. If the request is the first one received for that stopped transaction, IMS also sends a message to the MTO at that transaction's local system.

Application program abnormal termination

When an application program abnormally terminates, and the abnormal termination is not the result of a deadlock situation, IMS issues message DFS554 to the master terminal of the system in which the abnormal termination occurred. This message includes the logical terminal name of the input terminal. If the input message is still available, IMS issues message DFS555, which includes the first portion of the input message, to the input terminal. When IMS sends the DFS554 message, the message text includes the logical terminal name of the input terminal.

Dequeuing messages, responses, and transactions

You can use the /DEQUEUE command with the PURGE or PURGE1 keyword to dequeue messages, responses, and transactions that are in error for a specific MSC link.

When you enter the /DEQUEUE MSNAME PURGE command, the Message Control/Error exit routine (DFSCMUX0) is called before the command executes. The exit routine protects you from inadvertently entering the command in error and potentially destroying a message, response, or transaction by dequeuing it before it reaches its destination. The exit routine can suppress the dequeue of a message.

The system programmer must be aware of each of:

- The existence of the /DEQUEUE command, and the consequences of using it
- The interaction of the /DEQUEUE command with the Message Control/Error exit routine

The MTO should be aware of each of these facts:

- Your installation policy for using the /DEQUEUE command
- The fact that any messages that are dequeued are destroyed
- The /DEQUEUE command should be reserved for emergency use only, or as dictated by your site

Recommendations:

- Only authorize the MTO to issue the /DEQUEUE command.

- Establish installation-specific standards for using the /DEQUEUE command. Be sure to identify in advance and validate the scenarios where the MTO should use this command. Consider the /DEQUEUE command when coding the Message Control/Error exit routine.

Related reference:

 /DEQUEUE command (Commands)

 Message Control/Error exit routine (DFSCMUX0) (Exit Routines)

Establishing maintenance procedures

Running IMS utilities on a regular, prescribed basis is one of the main tasks for maintaining IMS. Maintaining IMS might also involve additional tasks such as reorganizing databases and installing service regularly.

Maintaining IMS might involve the following activities (done manually or automated):

- Running reports, analyzing the reports, and taking action based on the analysis
- Reorganizing databases on a regular basis
- Performing preventive maintenance on the message queues
- Running IBM-supplied or vendor-supplied tools on a regular basis
- Installing service on a regular basis
- Running IMS-supplied utilities on a regular basis

The frequency with which utilities are run is an installation-dependent decision.

Tip: You can determine the current maintenance level of your IMS system by using the /DIAGNOSE SNAP MODULE(*modname*) command.

Setting up standard JCL

You can set up standard JCL for the jobs that must be run (such as the utility jobs) as part of establishing procedures for your operations people.

Using DBRC can greatly assist you in this area, even though having completely “canned” JCL is not possible. For example, the MTO must supply the required log data sets when running the Database Batch Backout utility. If you use DBRC for recovery control, you may only have to set up JCL for DBRC itself.

Restriction: DBRC does not develop JCL for backout.

Operator test procedures

After you have prepared procedures for operating IMS, all those who operate the system should be given an opportunity to familiarize themselves with the procedures. After that, all procedures should be thoroughly tested.

Simulating system failures

To test some of the recovery procedures, certain types of system failures must be simulated. The following table shows possible failures and how they can be simulated.

Table 44. Simulating system failures

System failure	Simulated by
BMP or CCTL thread abended or canceled in error	Use MTO command /STOP REGION n ABDUMP.
IMS control region abended or canceled in error	Use z/OS MODIFY or CANCEL command.
MPP abended, looping, or in wait state	Use one of your own programs, or use the DL/I test program DFSDDLTO; use the special calls ABEND and ZING.
CCTL abended, looping, or in wait state	Use one of your own programs.
I/O error on SLDS or RLDS	Use an SLDS or RLDS that has previously been damaged. Reset, rewind, and reload tape drive.
z/OS error (loop or abend)	Unplug or switch off z/OS system residence drive. Cancel IMS control region.
Hardware error, no loss of real storage	Unplug or switch off z/OS system residence drive.
Power [®] failure, or hardware error with loss of real storage	Press System Reset, and re-IPL.

Qualities of good tests

The specific nature of your IMS use and installation setup will determine the details of your testing program. However, all good tests have these qualities in common:

- They are realistic.
- They are specific—addressing each one of the many possible operating situations.
- They involve the same people who operate IMS in actual, production-mode situations.
- Their results are verifiable.

Improvement through feedback

Although formal testing of your procedures ends when production begins, each actual instance of operation and recovery can be considered a test—a chance to improve your procedures. If problems develop or procedures are found to be vague or inaccurate, you should use this experience to update and improve your procedures.

You might also want to consider conducting periodic reviews of your procedures by your operations staff and system administration staff. Such reviews are particularly important as you add maintenance to your IMS system or upgrade to a new release of IMS.

Retesting when the system changes

Any time you significantly change your system, you should do at least some retesting. Changes could be in the form of either a new release or major upgrade of IMS itself, or a redesign or expansion of your production setup.

The retesting should verify that your old procedures still function properly in the new environment.

Related concepts:

 [DL/I test program \(DFSDDL0\) reference \(Application Programming APIs\)](#)

Chapter 7. Developing user procedures

End users (remote terminal operators) are the people who use the terminals connected to IMS. End-user procedures should focus on the needs of end users, not on IMS operation.

It is important that your end-user procedures be good ones, because the success of your online system depends largely on its acceptance by your end users.

The procedures should be prepared as an integral part of the application-development process. The layout and organization of the procedures should support the business function. Procedures differ for different types of applications. Procedures for a bank teller, for example, should cover more than just the entry of an IMS transaction. Procedures for a data-entry clerk, however, may consist largely of instructions for entering transactions.

End-user procedures are generally packaged in a guide. You can develop the guide yourself, or it might be appropriate for individual application-development teams to develop the guide. In the latter case, you should consider preparing “boiler plate” material to distribute to the application-development teams for basic procedures that are always included in end-user guides. Do not forget to update the boiler plate information whenever you make major modifications to the online IMS system. For example, if you convert from using non-VTAM to using VTAM, you probably need to change terminal logon procedures.

In general, your procedures need to familiarize end users with how to:

- Operate their terminal
- Establish a connection to IMS
- Communicate with IMS
- Disconnect from IMS
- Respond to error conditions

Related concepts:

“Operations personnel” on page 229

Procedures for user terminal operators

The application-development cycle should include the development of end-user procedures, along with education about the transactions and terminal hardware.

The writers of the procedures need to consider the following points when developing the end-user guide:

- Consistency between remote and central control
- Which services should be provided by the online system rather than the terminal or workstation
- Coordination of support for end-users
- Definition of problem-reporting procedures

Related concepts:

“Planning the content of procedures” on page 234

User operator tasks

An end user performs one or more tasks as part of his or her job. You need to provide your end users with instructions that are clearly written and accurate for the task that they will perform.

These tasks can involve many of your IT tools and systems, of which IMS is one. You should review those tasks to be sure your procedural instructions do both of the following:

- Help the user perform application-related tasks
- Safeguard the primary task

You should review existing instructions that have implications for administration services. The following table shows a partial checklist illustrating the characteristics of this review. You should expect other topics depending on the scope of the user procedure.

Your review contributes to the usability of the procedures. You should plan to correct them based on testing and user feedback during actual production.

Table 45. Checklist for administrative services

Review item	Related administration task
Are availability statements clear?	Check scheduling algorithm and operations procedures.
Are response claims reasonable and not misleading?	Reconcile with performance criteria. Ensure unusual delays are reported.
Are terminal connection procedures described?	Check whether the MTO or other operators help in connection.
Are security requirements explained?	Check security definition and maintenance procedure.
Is interpretative material included for output?	Check accuracy of IMS message and command output information.
Are instructions given for an alternative connection?	Check that the MTO has alternative configuration instructions.
Are instructions given for response to system and broadcast messages?	Check with developer on potential messages.
Are instructions given for service problems for the terminal, transaction input, and online system failure?	Coordinate failure reporting and information flow. Have MTO issue standard information messages.

Related concepts:

“Operating instructions for terminal operators”

Operating instructions for terminal operators

The process and goals of developing end-user documentation include developing information in several stages and to tailor the documentation to the experience level and information requirements of the end users.

The success and acceptance of an online system often depends on the quality of the end-user documentation. The operating procedures for terminal operators form an important part of this documentation, contributing to ongoing education and ease of use for the applications.

Develop and refine the procedures in several stages:

- Define initial structure and content
- Analyze existing procedures
- Test the new procedures
- Gather education feedback
- Gather production feedback
- Make any necessary modifications to the procedures

It is generally better to extract information from existing sources and tailor it to an application rather than point to total descriptions or general operation manuals. For example, describe the 3270 operational characteristics rather than assume that the operator is familiar with the content of the 3270 Data Stream Programmer's Reference.

Also, show representative screen formats rather than showing MFS output. You can obtain working copies of the screen layouts by using the copy feature on a remote 3277 terminal.

Do not ignore the last stage, the actual production cycle. In many cases, you can obtain useful feedback by selecting a group of end users to be pilot operators. Their mission is to validate and improve the end-user service. After making any necessary changes, publish the final procedures.

Do not neglect to update the procedures when you make modifications to the online IMS system. Sometimes the operational emphasis changes when you install subsequent application programs or IMS features. Your same pilot group could renew their activities when hardware changes or other significant application changes occur.

Some guidelines for the content of end-user operator documents follow. The goal is to have the document tuned to the experience level and information requirements of the end users:

- Make the main section procedural for the applications.
Ensure the instructions are sequenced in the order in which they are performed. Place transactions that are more frequently used, or more important, ahead of others. If there are numerous transactions, an alphabetic order might be better, if you include an introduction (or graphic) that indicates which are critical or used frequently.
- Illustrate your procedures with simple examples.
Show both entry and response formats. Where IMS commands are used, show only relevant keywords and explain the replies.
- Include instructions for error handling.
Address not only input errors but also responses to warning messages from the application program. Where IMS connection or system problems might occur, provide alternative instructions. If this material is lengthy, you could place it in an appendix. Arrange the topics in the appendix by symptom. Provide an explanation of error-reporting forms.
- Optionally, include a brief introduction to IMS and describe the MTO.

The emphasis should not be on how IMS works, but on how it supports data access and processing at your site. Distinguish, as necessary, between what IMS provides and what the operating system and hardware provide.

- Identify the main hardware operational characteristics.
Include advice that highlights efficient ways of using the terminal and warnings of any user actions that are prone to error. If appropriate, merge copies of relevant operating instructions into this section or appended to the document.
- Include only relevant IMS command operation details.
Make sure that commands and options described are allowed for end-user use. Discourage the use of commands that reserve IMS resources, such as /HOLD for conversations or /IDLE which enqueues output messages.
- Include a list of contacts.
Show names in preferred order and describe the conditions under which the user should make contact.

Related concepts:

“User operator tasks” on page 286

Potential use of IMS commands

Although the primary purpose of end-user procedures is to describe transaction input, you might also allow the users to enter a subset of IMS commands so that they can make more informed use of IMS resources.

For example, if the end users can choose the order in which transactions are submitted, they can use a /DISPLAY command to gauge the current queue levels.

Recommendation: If you do allow end users to enter IMS commands, consider disallowing the use of the ALL keyword to minimize the effect of these commands.

The following table presents a list of IMS commands and evaluates their potential use by a remote terminal operator. The /DISPLAY TRAN and QUERY TRAN commands are especially appropriate. However, because of the many keyword choices for these commands, you should specify individual keywords. For convenient entry, include command abbreviations.

The commands in the table that are marked with an asterisk (*) are available to the remote terminal operator using default command security. To allow the remote terminal operator to use other commands, you need to declare the full list of commands for that terminal through RACF (or an equivalent product). You can also use the DFSCCMD0 exit routine for command security.

Table 46. Command suitability for use by a remote terminal operator.

Command group	Command	Default security	End-user use	Comments
End-user specific	/RCLSDST	*	YES	Acts as remote VARY offline (for VTAM terminals).
	/RCOMPT	*	YES	Decentralizes control of VTAM component.
	/BROADCAST	*	YES	Provides convenient user interaction; use of message switching is preferred.
	/SET	*	YES	Provides convenient entry for batched transactions.
	/RESET	*	YES	Used with /SET.
	/RDISPLAY	*	YES	Identifies MTO location; used with /BROADCAST.
	/CANCEL	*	YES	Cancels both single- and multisegment messages.
	/EXCLUSIVE	*	YES	Safeguards discrete application usage.
Connection commands	/OPNDST		NO	Allows access to VTAM network.
	/CLSDST		NO	Restricts access to VTAM network.
	/QUIESCE		NO	Restricts access to VTAM network.
	/SIGN	*	YES	Allows access to IMS resources; required for signon verification.
	/RSTART		YES	Starts IMS resources; allowed if /STOP allowed.
	/MSASSIGN		NO	Alters assignments for IMS resources; required for MFS transactions.
	/FORMAT	*	YES	Formats a screen; required for conversation terminations.
	/EXIT	*	YES	Terminates conversations; required for non-VTAM switched terminals.
	/IAM	*	YES	Suspends conversations; gives flexible response.
	/HOLD	*	YES	Suspends conversations.
	/RELEASE	*	YES	Resumes conversations.
Connection commands for MTO	/START		NO	Starts IMS resources.
	/STOP		NO	Stops IMS resources.
	/CHANGE		NO	Allows changes to mode table and automatic session restart. Immediate termination best handled by the MTO.
	/IDLE		NO	Allows control of multiple systems.
Device control	/COMPT		NO	Sets terminal components; also available through /RCOMPT.
	/MONITOR		NO	Stops programmable remote stations; best handled by MTO.

Table 46. Command suitability for use by a remote terminal operator (continued).

Command group	Command	Default security	End-user use	Comments
Monitoring	/DISPLAY		YES	Displays information about IMS resources; convenient for status of queues and terminals.
	/TRACE		NO	Allows IMS to trace resource activity; IMS Monitor and traces affect system performance.
	/LOG	*	YES	Records user text on IMS system log.
Serviceability	/DIAG		NO	Retrieves diagnostic information for system resources (for example, IMS control blocks) without taking console dump.
Testing	/TEST	*	YES	Sets test mode; limit use for remote terminal operators.
	/LOOPTEST	*	NO	Tests I/O for terminal.
	/END	*	YES	Terminates test, looptest, or exclusive mode.
	/MSVERIFY		NO	Verifies IMS resources; used by administrator of multiple systems.
Resource control	/ASSIGN		NO	Assigns IMS resources; prime control of resources for MTO.
	/PSTOP		NO	Restricts use of resources.
	/PURGE		NO	Restricts input to resources.
	/LOCK	*	NO	Restricts resource use; use only for supervisor remote terminal operators.
	/UNLOCK	*	NO	Unlocks IMS resources.
	/RMxxxxxx		NO	Allows control of DBRC resources; database availability should be controlled by MTO.
	/SMCOPY		NO	Controls the logging of system messages, commands, and command responses to the secondary master terminal.
	/MODIFY		NO	Allows online changes to IMS system.
Recovery	/DBRECOVERY or UPDATE		NO	Controls database availability; restrict to MTO.
	/DBDUMP or UPDATE		NO	Controls database availability; restrict to MTO.
	/DEQUEUE		NO	Controls integrity of message queues; restrict to MTO.
	/CHECKPOINT		NO	Shuts down IMS; restrict to MTO; simple checkpoints can have a performance impact.
	/ERESTART		NO	Restarts IMS; restrict to MTO.
	/RECOVER		NO	Recovers full-function databases and Fast Path areas; restrict to a DBA or the MTO.
/SWITCH			NO	Switches data sets used in an XRF environment; restrict to MTO.

Related concepts:

“IMS commands used to communicate with IMS” on page 296

Related reference:

- ➡ IMS Command Language Modification facility (DFSCKWD0) (Exit Routines)
- ➡ /DISPLAY TRAN command (Commands)
- ➡ QUERY TRAN command (Commands)
- ➡ Command Authorization exit routine (DFSCCMD0) (Exit Routines)

Problem reporting for a remote terminal operator

The main flow of the end-user guide should present the normal operations sequence, but the operator needs to be prepared for unexpected or error responses.

When you examine the sequence of instructions for a remote terminal operator, you need to ask a series of “what-if” questions. The following are examples of these “what-if” questions:

- What if the terminal is not operational, behaves erratically, or loses contact with the system?
- What if the application program returns warning messages about the status of a database or of processing results?
- What if IMS system messages are received?
- What if master terminal or other operator actions interrupt processing?
- What if the z/OS or IMS system goes down?

You need to assess the importance of such unusual events and include appropriate procedures for them. The end-user guide could describe recovery actions inline with normal operations or include the information in symptom-based appendixes.

If your installation has a problem-reporting center, having a user-liaison group enables you to concentrate recovery skills rather than educating all end users. Many installations report terminal problems to network control personnel, and they have appropriate details about IMS and the application programs it runs. It is often a good idea to standardize problem reporting. You can prepare the answers to a series of questions, or elicit a problem's symptoms using a structured script.

Occasionally, the master terminal operator can be a point of contact. You should ensure such contact is only be on a formal, prearranged basis, because the master terminal operator has a primary duty of system control and monitoring. An example of acceptable contact is to request the scheduling of a BMP.

A good controlling technique is to require an error incident report. You can make preformatted forms, with examples of the information required, available to the end user. The data on the forms should indicate what action was taken or is pending. Administration personnel can then review these events on a regular basis. In this way, you can detect common problems or differences in operational procedures.

Related concepts:

“Administration support for errors” on page 298

Connecting to IMS

Before a user can use a terminal, the master terminal operator (MTO) must initialize the communication network using one or more of the following commands: /START DC, /START NODE, and (possibly) /OPNDST for VTAM terminals.

Terminal startup for nonswitched lines

For startup to be successful, the remote terminal must be turned on. IMS notifies the remote terminal of a successful /START or UPDATE command with the DFS059 terminal started message. Communication with IMS can only begin after the startup message is received.

Exception: A remote terminal can communicate with the following terminals before receiving message DFS059: the master terminal, physical terminals on the master terminal line, and the system console.

Terminal startup for switched lines

When a switched line starts, it can accept calls from remote terminals. The data set attached to the remote terminal must be used to establish a connection with IMS. After the call is made and a connection is established, the terminal operator uses the /IAM command to identify the terminal to IMS.

IMS notifies the remote terminal of a successful /IAM command with the DFS058 command completed message. Communication with IMS can begin only after this message is received.

Restriction: Fast Path transactions cannot be entered from a switched terminal.

VTAM terminals (nodes)

SLU type P, FINANCE (3600), and ISC terminal systems do not receive IMS-generated connect messages. Instead, they communicate with a user-written program that resides in the terminal system's storage. Type-1 SLUs in unattended mode also do not receive connect messages (because there is no operator to receive the message).

VTAM terminals receive message DFS3649, DFS3650, or both. The DFS3649 message requests signon information, and the DFS3650 message indicates successful signon (or connection) and session status. Message DFS3649 also appears in the following situations: when failures occur during the signon process, when the operator signs off, or (for ETO only) automatic session timeout due to inactivity.

Releasing VTAM terminals hung in I/O

Use the /TRACE command to specify what action, if any, IMS should take when it determines that a VTAM response has been outstanding for a longer time than you specify. This facility applies only to VTAM devices that have VTCBs or MSC links.

IMS can handle outstanding responses in one of the following ways:

- No action.
- Issue a message that I/O has been outstanding for a longer time than you have specified; that is, the device has timed out.
In this case, the operator must determine what action, if any, to take. For example, the MTO or an AOI program could issue an /IDLE command followed by an /ACTIVATE command and an /OPNDST command.
- Issue a message that I/O is outstanding followed by a VTAM VARY NET,INACT command to deactivate the terminal, and a VTAM VARY NET,ACT command to reactivate it. If the device is operable and defined to IMS as non-shared, and IMS is not shutting down, IMS issues a /OPNDST command to establish a session with the node.

Use the /DISPLAY command to display all the nodes that have I/O which has been outstanding for a longer time than was specified during installation. The /DISPLAY TRACE command displays the status of all nodes that have I/O outstanding for a longer time than you specified using the /TRACE SET TIMEOUT command.

Fulfilling security requirements

Depending on the security options specified for the system, the end user might need to use the /SIGN command to provide a user ID for verification by IMS or RACF and supply an IMS or RACF password. Users with ETO terminals must use the /SIGN command. You need to document the security procedures the end user must follow.

Related concepts:

 IMS terminal network (Communications and Connections)

Related reference:

 Terminals and equipment supported by IMS Version 12 (Release Planning)

Communicating with IMS

You can communicate with IMS using transactions, sending messages to other users, and using IMS commands.

Transactions for communicating with IMS

When using a terminal (a non-programmable workstation), the operator must enter a transaction in a certain format in order for it to be recognized by IMS.

When the operator uses the /FORMAT command, IMS formats the screen using MFS. This format is defined by the installation.

Transaction code format for entering transactions

Use the transaction code format shown below to send a message to an application program:

```
transaction_code      (password)      text
```

where:

transaction_code

Is a 1- to 8-byte alphanumeric code defined during IMS system definition. Special characters (embedded blank, slash, dash, equal sign, comma, and

period) are not allowed. A blank or left parenthesis must follow the transaction code. You can use the COMM system definition macro to eliminate the need for a blank in a segment consisting of only the transaction code.

(password)

Is a password that must be entered with the transaction code (when required by installation security). Enclose the password in parentheses and concatenate it to the transaction code. If a password is not required, IMS ignores any password you enter.

For terminals using MFS, password entry is controlled by installation-defined formats.

text

Is the single-segment or multi-segment input message. Only the first segment of a multi-segment message can contain the transaction code and password.

Use the /SET command to send a continuous series of messages to the same application program. The /SET command puts the terminal in preset mode and eliminates the need to enter the transaction code (and password, if required) for each message. Preset mode fixes the destination for all messages entered from a terminal.

The text included in the transaction is application dependent, and should be clearly documented for each transaction type an end user might use.

Response mode for transactions

When the response mode is in effect, IMS does not accept input from the line, terminal, or user until IMS has sent an output response to the previous input. Response mode can be defined for either full function or Fast Path.

Response mode describes a connection between IMS and a communication line, terminal, or user that occurs for certain types of terminals or users under conditions specified during system definition.

You can specify communication lines, terminals, or users as:

- Always operating in response mode
- Never operating in response mode
- Operating in response mode only if the transaction code being processed specifies it

Any transaction can be defined as a response-mode transaction. When IMS receives an input transaction that requires response mode, IMS makes the terminal unusable (locks the keyboard) until IMS processes the transaction and sends the response.

While in response mode, if the user enters a single transaction for which the application program generates no response, IMS automatically sends message DFS2082 as the response to the transaction. This message removes the terminal from response mode and unlocks the keyboard, if necessary.

While in response mode, if the user enters multiple transactions for which the application program generates no response, the keyboard remains locked and the terminal is unusable. The terminal remains in response mode even if the user attempts to unlock the keyboard (by, for example, pressing the RESET key on a

3270 terminal); any subsequent input is rejected. The first output from an application program whose input originated from this terminal satisfies the response.

User IDs defined with ETO user descriptors remain in response mode after abnormal session termination or autosignoff. In this situation, the end user should contact the MTO (using a means other than the terminal or user ID). If the terminal is non-operational, the MTO can enter an /RSTART command to restore terminal operation.

Attention: If you restart IMS while users are in response mode, what happens to transactions depends on whether they have reached a sync point:

- If the transaction is complete (it has reached a sync point), and the end user is waiting for a response, the user does not receive the normal response because IMS restart resets response mode. After the restart completes, the user might reenter the transaction, which could adversely affect data integrity.
- If the transaction is incomplete (it has not reached a sync point), IMS discards the transaction.

You should develop operating procedures to cover this situation for all end users who could be affected.

To reset static nodes and ETO dynamic users that are hung in Fast Path input response mode, issue the /STOP and /START commands in sequence with the appropriate parameter. For static nodes, issue the commands /STOP NODE and /START NODE. For ETO dynamic users, issue the commands /STOP USER and /START USER.

Conversation mode for transactions

Conversational processing allows an end user to have a continuing dialogue with an application program.

A user initiates a conversation by entering a transaction code that is defined as conversational. A conversation can follow an ETO user from one terminal to another. The maximum number of conversations that can be active for one user or one terminal at a time is 65535.

The following commands affect conversations:

- /DISPLAY displays the status of conversational processing in the system
- /EXIT terminates an active or suspended conversation
- /HOLD suspends an active conversation
- /RELEASE restarts a held conversation
- /RSTART LINE | NODE | USER recovers an active or suspended conversation on a stopped line, node, or user
- /START LINE | USER restarts the line or user but terminates the conversation

Using the appropriate keywords of the /DISPLAY command, the MTO can determine:

- The identification of all active (BUSY) or suspended (HELD) conversations in the system
- The number of conversations active or suspended in the system or related to a specific line, line and physical terminal, or user

- The number of inactive available conversations (the maximum number of conversations defined for the system minus the number of currently active or suspended conversations)

When you use the /DISPLAY command to display the status of a line or physical terminal, the displayed status includes, if applicable, the status of conversations related to that line.

Operators-to-operator messages

An end user can send messages to other operators by using the IMS message switch facility or by using the /BROADCAST command.

If your end users need to send messages, document clearly how they are to do this.

When using message switching, the input format is as follows:

logical_terminal_name *text*

where:

logical_terminal_name

Is the name of the logical terminal to which the message is to be sent. The logical terminal names allowed are specified during IMS system definition.

A blank must follow the logical terminal name. You can use the system definition COMM macro to eliminate the need for a blank in a segment consisting of only the logical terminal name.

text

Is the single-segment or multi-segment input message

For a multi-segment message, only the first segment should contain the logical terminal name. If the logical terminal name is specified in subsequent segments, it is treated as text.

You can use the /SET command to send a continuous series of messages to the same logical terminal. The /SET command puts the terminal in preset mode and eliminates the need to enter the logical terminal name for each message. Preset mode fixes the destination for all messages entered from a terminal.

IMS commands used to communicate with IMS

Although the primary purpose of an end-user guide is to describe transaction input, you might allow end users to use certain IMS commands. If you do, you should document their use in the guide.

Related concepts:

“Potential use of IMS commands” on page 288

Commands for special operating modes

IMS provides commands to enter and exit special operating modes, such as preset mode, test mode, MFSTEST mode, looptest mode, and exclusive mode. These modes alter the relationship between the terminal and the system.

A special operating mode remains active until one of the following events occur:

- The user issues a command to exit from the special operating mode
- The MTO stops and restarts the line, physical terminal, or user by issuing the /START command

To restart a terminal or user without affecting its special operating mode, use the /RSTART command.

The /IAM command removes a terminal from any special operating mode.

Preset mode

Preset mode fixes the destination for all messages entered from a terminal. Use the /SET command to enter preset mode. To leave preset mode, use the /RESET command.

Test mode

Test mode (also called “echo” mode) performs input and output operations between a terminal and IMS. The /TEST command places a terminal in test mode. Messages entered from a terminal in test mode are transmitted back to that terminal. The /END command terminates test mode. In an ETO environment, the /TEST command applies only to ETO-defined users, not to ETO-defined terminals.

MFSTEST mode

MFSTEST mode tests MFS control blocks without interrupting normal production activity. The /TEST MFS command places MFS-supported terminals in MFSTEST mode. The /END command terminates MFSTEST mode. In an ETO environment, the /TEST command applies only to ETO-defined users, not to ETO-defined terminals.

Looptest mode

Looptest mode tests output operations between a terminal and IMS. Use the /LOOPTEST command to enter this mode. This command creates an output write loop that repeatedly transmits a single-segment message to the terminal being tested. The /END command terminates looptest mode.

Exclusive mode

Exclusive mode prevents output messages from being transmitted to the terminal.



Exception: Output messages that are responses to messages entered at the terminal during exclusive mode are transmitted to the terminal.

After the terminal is removed from exclusive mode, the output messages are held for transmission. Use the /EXCLUSIVE command to enter exclusive mode, and the /END to leave exclusive mode. In an ETO environment, the /EXCLUSIVE command applies only to ETO-defined users, not to ETO-defined terminals.

Disconnecting from IMS

If an ETO-defined terminal or user ID is idle (no transactions are performed) for a given amount of time, IMS automatically disconnects the terminal or user ID. This amount of time is determined by the MTO. The end user can sign back on by issuing the /SIGN command.

Related tasks:

-  Autosignoff (ASOT) (Communications and Connections)
-  Autologoff (ALOT) (Communications and Connections)

Administration support for errors

Administrative support for end users can be provided in a variety of ways, one of which is to use a user-liaison group. By using a user-liaison group, you can concentrate recovery skills, rather than educating all end users.

It is a good idea to standardize end-user problem reporting. You can have the end user answer a series of questions before calling the user-liaison group, or you can have the user-liaison group elicit problem symptoms from the end user using a structured script. Among other things, each report should include which system and terminal experienced the error.

No matter how you decide to provide administrative support, the end-user guide should include a contacts list. This list identifies the people to be called for application, system, or terminal errors. Names, functions, and telephone numbers should be included in this list.

Related concepts:

“Problem reporting for a remote terminal operator” on page 291

Chapter 8. Operations and IMS-supported devices

This topic describes the devices that IMS supports from an operations perspective. This information should help you develop those parts of your end-user procedures that are device related.

This topic assumes that you are familiar with the general features and operation of each device. The focus is on device operation when you use IMS.

3270 Information Display System

During IMS system definition, you can specify that a 3270 system attached to IMS should operate in forced-response mode, nonresponse mode, or transaction-dependent response mode. For an Extended Terminal Option (ETO) terminal, specify these modes in the ETO user descriptor.

When 3270 terminals are not attached to IMS through VTAM, the MTO can initiate their use by entering the appropriate /START LINE command for the lines to which the 3270 components are attached.

When 3270 terminals are attached to IMS through VTAM, they are initiated when one of the following occurs:

- The z/OS VTAM network operator enters a VARY NET command (for example, VARY NET, ID=*a*, LOGON=*b*).
- The IMS MTO enters the /OPNDST command.
- If the 3270 terminal components are defined in the VTAM definition as belonging to IMS, VTAM makes them automatically available.
- The 3270 terminal operator enters a request to VTAM to be connected to IMS.

Each operational terminal receives a DFS059I message indicating that the terminal has been started. The MTO can use a /DISPLAY command to determine which lines are attached.

Connecting to IMS (3275 with dial feature-switched line)

When IMS accepts the /IAM command signon, the 3275 is physically and logically attached to IMS. From this point on, operate the switched 3275 in the same way as a nonswitched 3275.

To connect a 3275 with the dial feature-switched line to IMS:

1. Make the 3275 ready (if the terminal is off, power it on; if it is already on, power it off and back on).
2. Dial the IMS telephone number on the data phone.
3. Switch to data mode when you receive the correct answer tone.
4. Press the 3275 Clear key.
5. Enter the /IAM command when you receive the DFS073I message.

Disconnecting from IMS (3275 with dial feature-switched Line)

To physically disconnect the 3275 from IMS, enter the /IAM DONE command, or raise and release the 3275 Disconnect switch, then place the handset on the hook when the System Ready light goes off.

Interacting with IMS

If you enter data with a keyboard, a selector light pen, a magnetic card reader, or program function key special features, use the /FORMAT command to display a format on the screen before you begin working with IMS.

After you receive either message DFS059I or DFS3650I telling you that the terminal or session is started, you can enter a transaction code, a command, or an IMS message switch.

When you enter a transaction code, IMS returns and displays the output data from the application program that processes that transaction. This output might include a prompt to solicit additional input. You can fill in only the blank fields on the screen.

After entering data, press Enter to send it to IMS for processing by the application program. Continue in this manner until the interaction with IMS or the application program is complete.

3270 terminal components that operate with IMS

You must understand the various functions of 3270 terminal to work with IMS.

Terminal display size

The 3270 display station can display from 480 to 3564 characters of information, depending on the model. These characters can be alphanumeric characters, special characters, certain device-control characters (visible on the screen), and attribute bytes (not visible on the screen). In addition, you can prevent display of information in certain fields on the screen.

The 3275 or 3277 terminal can display:

- 480 characters (Model 1)
- 1920 characters (Model 2)

The 3276 or 3278 terminal can display:

- 480 or 960 characters (Model 1)
- 1920 characters (Model 2)
- 1920 or 2560 characters (Model 3)
- 1920 or 3440 characters (Model 4)
- 3564 characters (Model 5)

The 3279 terminal can display:

- 1920 characters in 4 colors (Model 2A)
- 2560 characters in 4 colors (Model 3A)
- 1920 characters in 7 colors (Model 2B)
- 2560 characters in 7 colors (Model 3B)

3270 terminal page-advance function

A transaction's output can be more than one full screen of data. You request the additional output by pressing the PA1 (program access 1) key on the terminal keyboard. If you continue to press the PA1 key, the terminal displays successive pages until you reach the end of the current message. If you press PA1 when there are no more messages waiting to be sent to the display station, the display remains unchanged or IMS sends an appropriate error message.

3270 terminal message-advance function

If you press the PA2 key at any time during a message, IMS deletes the current output message and displays the next message. If no messages are waiting when you press the PA2 key, IMS does not modify the screen.

How to enter passwords from 3270 through MFS

Entering a password through Message Format Service (MFS) can be different from password entry on a cleared screen or on other IMS terminals.

The password can be created from:

- Data read from a 3270 ID card reader
- Data that you key into one or more fields or a portions of fields in the device format
- Data that you key into a field that is not displayed on the screen and is designated as the password field only in the device format
- Literals
- Combinations of the above

Terminal operator logical paging

If you do not specify operator logical paging, the terminal operator can advance pages only one at a time, by using the page advance function. If you do specify operator logical paging in the application program's message output description, the terminal operator can page forward or backward to any page in a message.

With operator logical paging, the terminal operator can enter any of the functions shown in the following table to display a particular logical page.

Table 47. Functions for Operator Logical Paging

Enter	To display
=	The next logical page
=n	The <i>n</i> th logical page of the message
=nn	
=nnn	
=+n	The <i>n</i> th logical page after the current page
=+nn	
=-n	The <i>n</i> th logical page before the current page
=-nn	
=L	The last logical page of the message

Enter these functions in the screen area reserved for command and page request information. If no area is reserved, either clear the screen (using the Clear key) and

then enter the paging request, or use the next logical page (NEXTLP) function if it has been defined as a PF key in the format definition.

If you request a page outside the range of the pages in the message, IMS issues an error message. After the error message displays, press the PA2 key to continue the message that was being displayed when the error occurred. If a logical page in the message is larger than one physical page, use the PA1 key to step from one physical page to the next within a logical page.

Modified data tags and the 3270 keyboard

IMS resets modified data tags each time it unlocks the 3270 keyboard.

Exception: IMS does not reset modified data tags when:

- IMS sends a message using the SYMSG field.
- IMS sends output to its master terminal using the Message Format Service (MFS) master-terminal format. IMS resets modified data tags only after you enter input.

3270 terminal screen protection

Under certain conditions, IMS avoids sending output to a 3270 terminal that could possibly destroy the operator's current work. IMS protects the screen so that you can safely enter data or control requests.

You can:

- Set the screen mode to protect and:
 - IMS successfully sends output to the screen.
 - You press the Clear key and IMS responds by unlocking the keyboard.
 - You or the application program perform a successful Copy function.
- Requested the NEXTMSGP function (the PA3 key or a PF key so defined) when no output message was available.

The screen is not protected after you:

- Press one of the following keys:
 - Enter
 - PA1 or PA2
 - Any PF key (except a PF key defined as NEXTMSGP)
- Operate either of the following devices:
 - The selector light pen
 - The ID card reader

Attention: If the screen is protected, an IMS shutdown checkpoint cannot complete.

During IMS system definition (using the TERMINAL and TYPE macros or ETO logon descriptors) you can specify that a 3270 screen be left in unprotected mode when IMS sends an output message to it. In unprotected mode, IMS can send output to the device at any time without requiring input from the terminal (an input message or a PA1 or PA2 action). IMS provides this option on a terminal-by-terminal or message-by-message basis. The terminal option of unprotected mode applies to all messages for that terminal. If you do not select unprotected mode, MFS-formatted messages leave the screen protected or unprotected on a message-by-message basis, and messages that bypass MFS leave the screen protected.

By allowing the terminal to be unprotected, IMS can continually update the 3270 display terminal with current information without requiring the terminal operator to initiate any action to receive output messages. Use this option only with terminals that require limited input and require the most current IMS output message.

To avoid losing data, press the PA3 key (if the PA3 key is not used for copy) before entering data. When you press the PA3 key, you put the terminal in protected mode, and you can enter data. After you enter the data, IMS returns the terminal to unprotected mode.

3270 terminal lock and unlock option

During IMS system definition, by using the `TERMINAL` and `TYPE` macros or Extended Terminal Option (ETO) logon descriptors, you can specify whether the application program or IMS should be responsible for unlocking the terminal keyboard and resetting the modified data tags.

If you select `LOCK`, the application program unlocks the keyboard after it receives input. If you select `UNLOCK`, IMS controls the unlocking of the keyboard. This option only applies to messages that bypass Message Format Service (MFS) on output with a MOD name of `DFS.EDTN`.

If the application program terminates without unlocking the keyboard, you can unlock the keyboard by pressing the Reset key.

Attention: If the application has not terminated but is slow in responding, do not press the Reset key and start entering data, because the data might be lost when the application program sends another output message to the terminal.

3270 MFS bypass option

If you select the 3270 Message Format Service (MFS) bypass option (by using a MOD name of `DFS.EDTN`), the application program controls the functions performed by the Clear key, selector pen attention, PA1, PA2, or PA3 keys.

3270 terminal programmed symbols

Programmed symbols (PSs), an optional feature of the 3270 terminal, allow storage and access of up to six sets of 190 characters, each of which can be user-defined and loaded by the program. You can load PS buffers with either a VTAM application or a user-defined procedure. After the PS buffers are loaded, you can use them on a field-by-field basis on the screens formatted by the Message Format Service (MFS).

If a hardware error occurs while you are loading a PS buffer, IMS:

1. Returns the message used to load the PSs to the IMS message queue.
2. Takes the terminal out of service (except for SLU-2 devices).
3. Logs the error to the IMS log.
4. Sends a message to the IMS master terminal.

After you correct the hardware error and the terminal is in service, IMS re-sends the message used to load the PSs.

If the PS load function failed because of an error in the message used to load them:

1. Use the `/DEQUEUE` command to dequeue the message. You might have to enter the `/DEQUEUE` command from the master terminal.
2. Correct the error.

3. Reenter the transaction to send the message to reload the PSs.

For SLU-2 devices, if you have not loaded the PS buffers and IMS sends a message requiring the use of a PS buffer, the 3270 terminal rejects the message and IMS:

1. Returns the invalid message to the IMS queue.
2. Logs the error to the IMS log.
3. Sends a DFS2078I message to the node, telling you that the output was rejected. However, if the message being rejected is in storage, IMS sends the DFS2078I message to the IMS master terminal and closes the node.

IMS leaves the terminal in protected mode after it receives the DFS2078I message.

3270 terminal screen formats

The 3270 can operate with either an unformatted or a formatted screen. Additionally, there are also input formats supplied by IMS, and installation-defined formats.

Unformatted mode

The screen is in unformatted mode:

- When you turn on the display station
- Whenever you press the Clear key on the keyboard

While the screen is in unformatted mode, you can enter the following types of messages:

- IMS commands
- Paging requests (when a message using logical paging is in progress)
- Transaction or message switch data that does not require Message Format Service (MFS)

Formatted mode

IMS MFS defines the 3270 screen to have various fields with predetermined uses. The layout of the fields on the screen and the data displayed depend on the application program. A format consists of all the fields on a screen, and each message sent to a 3270 terminal must indicate a specific format.

An application program might use several types of formats. The format on the screen can change between messages, between input and output sequences of a message, or between transactions.

Different formats define input areas at different locations on the screen. When the 3270 terminal receives a message, the format usually places the 3270 cursor at the beginning of the first input area. The input can consist of commands, transactions, or message switches.

IMS-supplied input formats

Some IMS-supplied formats allow one or two segments of input from the 3270 screen (one segment equals one line on the screen). When you want more than two segments of input, you can use the /FORMAT command to specify the DFSM04 format. This format is an IMS-supplied format that allows up to four segments of input data. Each segment in the format consists of only the data entered without new line characters or other cursor-control characters.

Installation-defined formats

The installation can develop a variety of formats. Each format definition determines the information that you can enter and the way that you can use the format.

Installation-defined formats allow any of the following (either single-segment or multi-segment):

- Transactions
- Page requests
- Commands

If the current installation-defined format does not allow the necessary input, try using the page-advance function (PA1), the message-advance function (PA2), or the Clear key to produce a screen format that allows entry of the desired data or command. Remember that using PA2 discards the remaining portion of the current output message.

Related concepts:

“Operation using the MFS master-terminal formatting option” on page 308

Multiple physical page input

An input message can consist of more than one physical page. For example, a 3275 or 3277 Model 1 display (480 characters) requires four physical pages for an input message designed to be equal to the 3275 or 3277 Model 2 screen size (1920 characters).

Use the following operational sequence as a guide when preparing procedures for 3270 operators who use Message Format Service (MFS) and enter multiple physical pages of data to create one message.

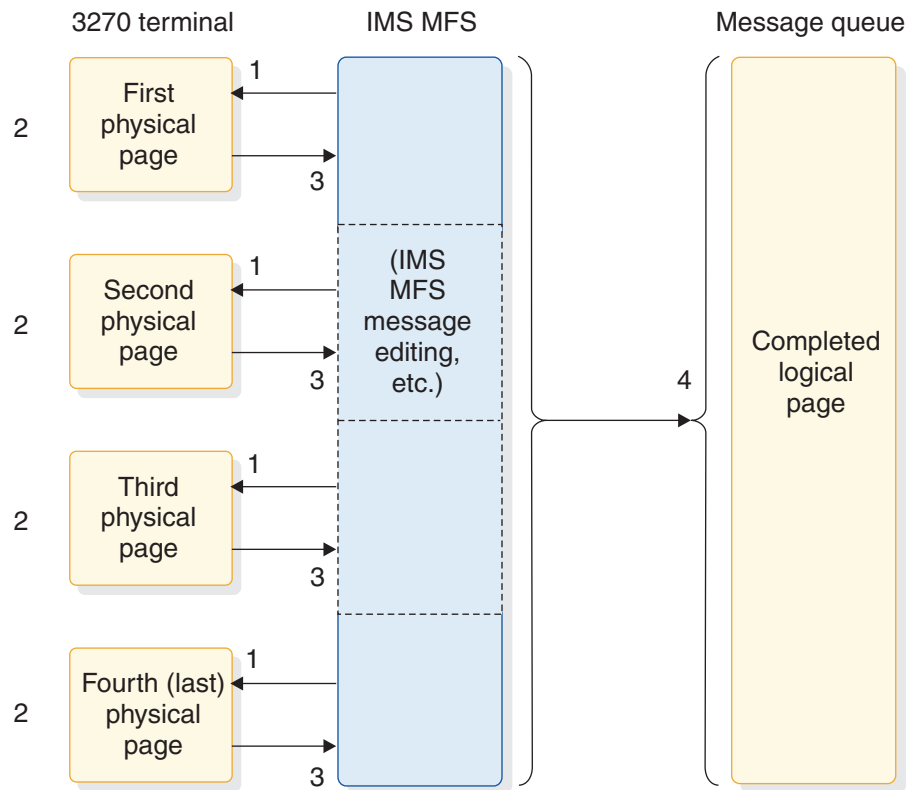


Figure 13. Data entry sequence for multiple physical pages

Note: The following notes describe the processing that occurs in the previous figure.

1. IMS sends the data-entry format for the physical pages to the 3270 screen.
2. The 3270 terminal is ready to accept your input whenever it displays the format of a physical page. You can input data (to start or continue creating a message) or use control functions.

While the 3270 terminal displays the first physical page of an output message, your action determines what operation IMS performs:

- a. You can enter a message or choose not to enter a message.
 - When the 3270 terminal displays the first physical page and before you signal data entry, the following control functions are valid:
 - The PA1, PA2, and PA3 keys
 - The Clear key
 - The NEXTTPP, NEXTMSG, NEXTMSGP, and PAGEREQ functions

After you signal data entry, the rules for normal message creation apply.

- b. When you do not signal data entry on the first physical page, IMS assumes a display-only operation. When the 3270 terminal is displaying the first physical page, MFS can create an input message only if an output message is still available. An output message is not available if your previous request dequeued the current message and no other messages are enqueued. When IMS displays any page other than the first physical page on a display-only operation, only the following control functions are valid:

- The Clear key
- The PA1, PA2, and PA3 keys

Attempting any other function results in an error message.

- c. You can create an input message and call for entry of that data:
Enter the required data into the physical page format displayed on the screen and then press Enter. Pressing a PF key that has not been assigned another specific function also causes the data to be entered.
You may want to advance to the next physical page without entering data from the currently displayed page. When the displayed physical page is not the first physical page, press the PA1 key.
 - d. You can signal the end of the current input message by entering the ENDMPPPI function.
MFS edits the data keyed onto the screen and then performs end-of-message procedures to send the message to the message queue.
 - e. You can cancel the current input message by pressing one of the Clear, PA2, or PA3 keys, or by entering one of the NEXTLP, NEXTMSGP, or PAGEREQ functions.
IMS performs the respective function and discards all previous input for the current message.
3. You signal that the data on the screen is to be entered by pressing the Enter key.
MFS accepts the data from the 3270 terminal and performs any message editing that is required (inserting fill characters, literals, and so forth).
 4. MFS sends the completed message to the IMS message queue.
You signal the end of the current message by entering the ENDMPPPI function.
MFS performs all required message editing and end-of-message procedures.

Additional rules that you must be aware of when you operate a 3270:

- You can enter the following from any physical page, but only once per message:
 - A PFK literal
 - An immediate pen-detect literal
 - An Op-ID card
- If you define DFLDs (device fields) with the same label on more than one physical page, MFS accepts data from only the first one defined as a part of the message.

3270 master terminal support

A 3270 terminal can be the IMS master terminal. When the IMS master terminal is a 3270 terminal, IMS sends all system messages to the primary device. IMS also sends selected system-generated messages critical to IMS operation to the secondary device.

A 3270 master terminal consists of two 3270 devices: a 3270 display (3275, 3276, 3277, or 3278) and a 3270 printer (3284, 3286, 3287, 3288, or 3289). IMS defines the 3270 display as the primary master terminal and the 3270 printer as the secondary master terminal. IMS does not support the 3284 Model 3 as a master terminal.

IMS responds only to the primary device in the following situations:

- When you enter erroneous information, such as a nonexistent transaction code or an invalid parameter, from the master terminal
- When another operator switches a message specifically to the primary device
- When an application program inserts a message to the primary device
- When IMS issues a command-completed message or output from a /DISPLAY command

If the secondary device is non-operational, IMS saves enqueued messages for transmission when it becomes operational.

The /ASSIGN command allows independent assignment of the 3270 display and 3270 printer.

Related reference:

➡ COMM macro (System Definition)

Operation using the MFS master-terminal formatting option:

When the IMS master terminal is a 3270 device, IMS displays system messages either with IMS-supplied formats or with the MFS master-terminal format

The MFS master-terminal format is optional and is not available for all 3270 display devices. Choosing a 3270 display with 24 lines by 80 columns (a 1920 character display such as a 3275, 3276, 3277, or 3278 Model 2) as the IMS master terminal allows use of the MFS master-terminal formatting option.

When you use this option, MFS divides the 3270 master-terminal screen into four areas: a message area, a display area, a warning area, and a user input area. The following figure shows the relationship of the areas on the screen, followed by a description of each area.

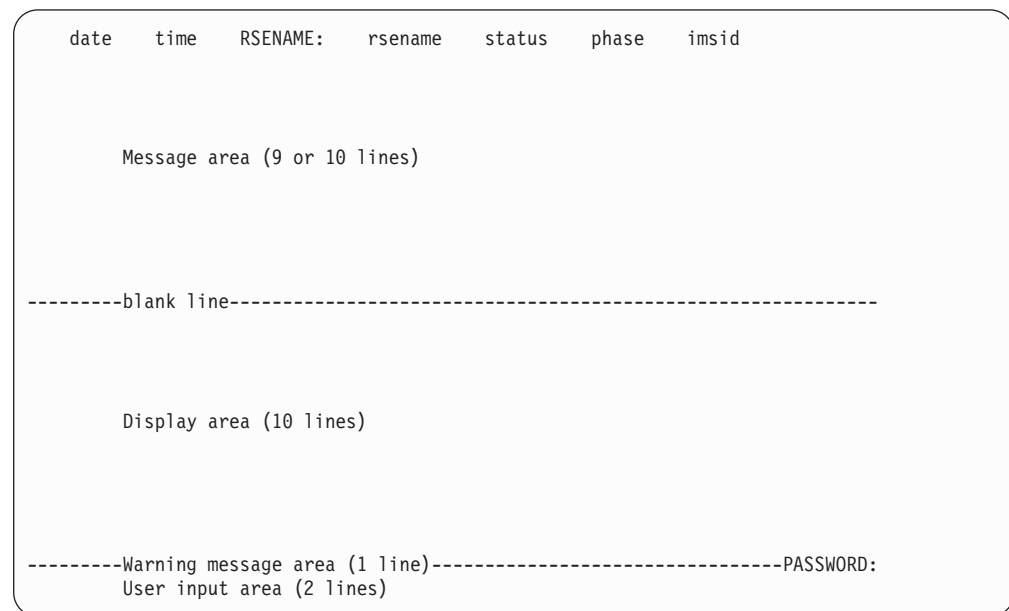


Figure 14. MFS screen format of 3270 (24-line, 80-column) master terminal

Related concepts:

“3270 terminal screen formats” on page 304

Status line:

If the IMS subsystem is XRF capable, the top line of the screen displays fields in high intensity.

date The current date

time The current time

RSENAME: *rrename*

The name specified for the RSENAME parameter in the DFSHSBxx member of IMS.PROCLIB

status The current status, which can be null, ACTIVE, or BACKUP

phase The current XRF phase, which can be any of the following:

Null status	ACTIVE status	BACKUP status
INITIALIZATION	null AWAITING I/O PREVENTION	AWAITING SNAPQ SYNCHRONIZATION TRACKING IN PROGRESS TAKEOVER REQUESTED TAKEOVER IN PROGRESS

imsid The IMS ID of the active subsystem

Message area:

Message Format Service (MFS) designates the first ten lines of the screen as the message area if the subsystem is not Extended Recovery Facility (XRF) capable, or the first nine lines of the screen under the status line, if it is XRF capable. When you signal readiness to receive output, IMS displays the new data on the top line in the area, overriding the previously displayed data, and inserts a blank line following the new line to separate it from the old messages.

These lines receive such messages as:

- Responses to IMS commands (except /DISPLAY, QUERY, and /RDISPLAY)
- Messages sent from other terminals or applications
- IMS system messages

When the message area is full and more data is ready for output, MFS sends a message in the warning message area.

Display area:

Message Format Service (MFS) displays the output from the /DISPLAY, QUERY, and /RDISPLAY commands here, one page at a time. Each new page of the displayed data completely replaces the previously displayed page.

The ten lines below the message area are the display area.

Warning message area:

The warning message area contains a literal, PASSWORD, followed by a field in which you can enter an IMS password. To enter a password, position the cursor immediately after the colon (:). You can press the back tab key to move the cursor from the user input area to the password field.

Line 22 is the warning message area. The four possible warning messages issued by Message Format Service (MFS) are:

MASTER LINES WAITING

Displayed when you receive a multisegment message that is larger than the currently available message area, or when the message area is full and additional lines need to be displayed. To display the remaining lines, press the PA1 key.

MASTER MESSAGE WAITING

Displayed when the message area is full and you receive a message from another terminal or from an application program. To display the message and reset the warning, press the PA1 key.

DISPLAY LINES WAITING

Displayed after the display area is full because the output of a /DISPLAY, QUERY, or /RDISPLAY command is more than 10 lines. To continue displaying the output, press the PA1 key.

USER MESSAGE WAITING

Displayed when you receive a message for the master terminal that is not formatted in one of the IMS default formats. To display the message (in its format), press the PA1 key. To return the screen to the master format, enter any IMS command, or place the terminal in unprotected status and wait to receive a message in the master format.

User-input Area:

The two bottom lines (lines 23 and 24) on the screen are the user-input area. You can enter up to two 79-character segments of input. When you are not entering data, the cursor is located at the beginning of this area (column 2 of line 23).

Message Format Service (MFS) returns the cursor to this location when a response is received for the entered data. MFS leaves the entered data on the screen and you overlay it with your next input.

A 3270 master terminal using the MFS master-terminal formatting option is normally in unprotected status. This means that IMS can send output to the terminal at any time. You can put the terminal in protected status by pressing the PA3 key or the PFK11 key before entering input data.

The MFS-supplied master-terminal format defines literals for nine of the 3270 program function (PF) keys. The PF keys, when used with the 3270 master terminal formatting option, generate specific IMS commands. Pressing a PF key inserts a command into the first segment, in front of the data. The commands associated with the nine PF keys are:

PF Key**Command/Function**

- | | |
|---|----------------------|
| 1 | /DISPLAY |
| 2 | c |
| 3 | /DISPLAY STATUS |
| 4 | /START LINE |
| 5 | /STOP LINE |
| 6 | /DISPLAY POOL |
| 7 | /BROADCAST LTERM ALL |

- 11 NEXTMSGP
- 12 Copy (3270 remote)

When pressing the PF7 key to enter the /BROADCAST LTERM ALL command, make sure the first line of input in the user-input area contains only the command. The second line of input should contain the data to be broadcast.

When you use the master-terminal format, MFS displays any message whose MOD name begins with DFSM0 (except DFSM03) in the message area. MFS displays any message whose MOD name is DFSDSP01 in the display area. Messages with other MOD names cause the warning message USER MESSAGE WAITING.

3270 terminal copy function

You can use the IMS Copy function to obtain a printed copy of the currently displayed screen on a remote terminal. You can define the Copy function for one or more terminals on a remote 3270 line.

You or an application program can use the Copy function through the use of the system control-area field. If you want to copy a screen from a display station, and the Copy function is allowed for your terminal, press the PF12 key.

The rules that govern the Copy function are:

In the SNA environment:

For SNA-defined nodes (in other words, SLU 2), IMS sends a copy request to the display node. The controller that selects the printer processes the copy command. The printer need not be defined to IMS.

In the non-SNA or bisync environment:

- If you request a copy from a 3275 display station, and a 3284 Model 3 printer is attached, the 3284-3 produces the hard copy.
- If you attempt a copy from a 3275 display station, and the 3284-3 printer is not ready, IMS issues an appropriate error message.
- If you request a copy from a 3277 display that is attached to a 3271 or 3274 control unit, or from a 3278 display attached to a 3274 control unit, or from a 3276 display (Models 1 through 4), IMS selects one of the following printers to receive the listing:
 - The printer that, during IMS system definition, you selected as eligible to receive copies from a specific 3270 display.
 - The first available printer of a group of 3270 printers that, during IMS system definition, you selected as eligible to receive copies from a specific 3270 display. If all printers are currently busy, the keyboard on the display station remains locked until a printer that can be used for the copy becomes free.
- If you attempt a copy from a 3270 display station attached to a 3271 or 3274 control unit or from a 3276 display station and no printers are currently available (all printers are stopped, pstopped, in exclusive status, or require intervention), IMS issues an appropriate error message.
- If you want all copy-function listings to be directed to a specific printer, you can use the /STOP command to prohibit the use of the other printers, or you can turn off the other printers.

Restriction: The IMS copy function in the non-SNA/Bisync environment is not supported for ETO terminals.

To terminate a Copy function before it has completed, press the Reset key followed by either the Clear key or the Enter key. IMS terminates the copy function and unlocks the keyboard.

Related concepts:

 Printer Operations

3270 terminal function keys

Certain keys on the display station keyboard are reserved for use by IMS. If you press any of these keys (or the selector light pen or operator ID reader) when the terminal is not screen protected, IMS might not be able to recognize your action.

Recommendation: Use these keys only when the screen is protected from receiving output.

These keys and their functions are:

PA1 – Page advance

The page advance key requests IMS to display the next screen of the current output message. The next screen can be the next physical page or the next logical page. The page-advance function is the only method of advancing from one physical page to the next physical page when a logical page consists of more than one physical page.

PA2 – Message advance

The message-advance key requests IMS to display the next message enqueued for the terminal. IMS deletes the current message, even if you have not displayed every page of it. If no more messages are waiting for the terminal, IMS leaves the current message on the screen.

PA3 – NEXTMSGP

The PA3 key usually performs the NEXTMSGP (message-advance protect) function, but could instead perform the Copy function.

PF12 – Copy

PF12 performs the Copy function for terminals that allow this function. The Copy function causes IMS to print the current display. IMS does not restore keyboard operation until it has either transferred the message to the 3270 printer (3276, 3277, or 3278) or completed the printing (3275). After a copy has completed on a 3275, reposition the cursor to an appropriate location before attempting any input.

Clear Key

To set the display screen to unformatted mode, press the Clear key once. While the screen is in unformatted mode, it is protected from receiving other output until you perform an input operation.

If IMS currently displays an output message, pressing the Clear key does not dequeue the output message.

If you press the Clear key when the terminal is not screen protected, IMS might not recognize the Clear request. However, if IMS does recognize the request, it unlocks the keyboard, displaying no data. If you press Clear and IMS responds by displaying data, the screen is invalid for input. You must press the Clear key again.

Enter Key

The Enter key sends to IMS either the fields currently modified by the operator or the last output message. The keyboard remains locked until IMS receives the message and restores keyboard operation.

3284 Model 3 printer operation

When attached to a 3275 terminal, IMS supports the 3284-3 printer as a component of the 3275 terminal. IMS considers the display portion of the terminal to be component 1 (COMPT 1) and the printer portion to be component 2 (COMPT 2).

As with any component-type terminal, IMS sends messages to the two components on a rotating basis.

If IMS cannot send messages to the printer component, it sends messages continuously to the display component, as if no printer component existed. The operator must intervene to allow printer output.

If IMS cannot send messages to the display component, it sends messages to the printer as if the display component did not exist. As long as IMS can send messages to the printer component, no operator intervention is required.

Any situation (such as a stopped LTERM or non-operational printer) that prevents sending messages to the LTERM or LTERMs assigned to a particular component causes IMS to stop sending messages to that component.

To suspend messages to either the printer or the display component, use the /RCOMPT command with the NOTRDY keyword. This command sets the component that you specify as non-operational (not ready) if the 3270 is attached using VTAM.

Using programmed symbols for IBM 3270

Programmed Symbols (PS), an optional feature on the IBM 3270 terminal, lets you store and access up to six character sets of 190 characters each that are user-definable and program-loadable.

After the coded graphic character sets have been loaded into the PS buffers, they can be used on a field-by-field basis by using Message Format Service (MFS).

If the PS buffers are desired and have not been loaded, they must be loaded by a VTAM application or by some other means.

Determining whether the buffers are loaded

If you know that the buffers have already been loaded with the desired programmed symbols, you do not need to resend the entire load programmed symbols data stream.

To determine if the buffers are loaded with the desired programmed symbols, use one of these methods:

- Maintain a handwritten log at the device and record the current status of the PS buffers.
- Run a test transaction that attempts to use the desired programmed symbols. Either the desired transaction's output is successfully displayed or an error message (DFS2078) is sent to the terminal and the message is returned to the IMS message queue, with the terminal left in protected mode. The message must be dequeued and a transaction started to load the programmed symbols buffer.

Loading the PS buffer

If the operator knows the PS buffers need to be loaded (because the device was just powered on), enter a response-mode transaction that loads programmed symbols by using an installation-written application program that loads PS buffers.

The load PS data stream should be the first part of the message sent by the application program and a notification, such as THE PROGRAMMED SYMBOLS LOAD FOR psname COMPLETE should be the remainder. The notification informs the terminal operator when the programmed symbols have been loaded. This transaction requires the Message Format Service (MFS) bypass option.

When the output device for which the PS buffers are to be loaded includes a printer or a different display, another technique must be used.

Example: These events illustrate one example of using an automated operator application program:

1. The master terminal operator at Display_A enters a transaction (response or conversational) requesting that programmed symbols be loaded for Display_A, Printer_B, and Display_C.
2. In processing the transaction, an AOI program assigns LTERMs for Printer_B and Display_C, temporarily, to a special PTERM reserved for loading the PS buffers. The AOI program assigns dummy LTERMs to Printer_B and Display_C.
3. The AOI program inserts the message containing the load programmed symbols messages to the dummy LTERMs of Printer_B and Display_C.
4. The AOI program sends a programmed symbols message to Display_A.
5. The master terminal operator visually verifies messages on both displays and the printer and confirms to the transaction that the load was successful.
6. The transaction reassigns the LTERMs back to their original status.

Related reference:

 MFS bypass for the 3270 or SLU 2 (Application Programming APIs)

Solving load problems

If a hardware error occurs while loading a programmed symbols buffer, IMS performs series of predefined actions.

If a hardware error occurs while loading a programmed symbols buffer:

- The message used to load the programmed symbols is returned to the IMS message queue.
- The terminal is taken out of service.

Exception: SLU 2 devices are not taken out of service.

- The error is logged on the IMS log.
- Message DFS2078 is sent to the IMS master terminal. For AOI purposes, provide the message number, if possible.

After the hardware error is corrected and the terminal is in service, the message to load the programmed symbols is resent.

If the loading of the programmed symbols fails because of an error in the message to load the programmed symbols, the operator must:

- Dequeue the message. The master terminal operator might need to issue the dequeue (/DEQ) command.
- Correct the error.
- Reenter the transaction to resend the programmed symbols load message.

If a method is available for informing the next user of the programmed symbols buffer status, the terminals with loaded PS buffers should not be powered off. When a power failure occurs or a terminal is powered off, the contents of the PS buffers might be lost.

When a terminal is powered on, if no IMS messages are waiting to be sent to the display, an IMS response mode transaction should be entered to load all required PS buffers or some non-IMS method should be used to load the PS buffers. However, if IMS messages are waiting to be sent, and these messages require the use of one or more PS buffers, sending the queued messages must be delayed until the PS buffers can be reloaded.

For SLU 2 devices, if the PS buffers are not loaded and a message requiring the use of a programmed symbols buffer is sent to the terminal, the message is rejected and IMS:

- Returns the invalid message to the IMS queue
- Logs the error on the IMS log
- Sends message DFS2078 to the IMS master terminal and to the node, indicating an output failure.

The terminal is left in protected mode after receipt of message DFS2078.

A user application program can be designed to queue to an LTERM an unsolicited message requiring a particular programmed symbols buffer to be loaded. The first part of the message can include a load programmed symbols data stream; however, this message cannot be processed by Message Format Service (MFS).

Reloading the PS buffer

When a message is waiting on the IMS queue for a terminal and requires programmed symbols that are not loaded, you must load the programmed symbols.

To load the programmed symbols, take one of the following actions:

- If the terminal is attached by VTAM, load the PS buffers using a VTAM application.
- If a queued message requires a programmed symbols buffer, and it is normal user output (for example, the output is not response mode or conversational), then use a response-mode transaction to load the programmed symbols buffer, which loads the programmed symbols buffer so that the queued message is properly displayed.
- Dequeue (/DEQ) the message (or have the master terminal operator dequeue the message) that requires the use of a programmed symbols buffer; enter a transaction to load the programmed symbols buffer required; and then reenter the transaction that originally generated the queued message.

Application design considerations

A message that uses programmed symbols is different from other IMS output messages; it relies on the correct programmed symbol (PS) buffer content for its presentation.

One application might require a PS buffer content that is incompatible with a different application. Further, the PS buffer can be erased by external action (such as powering off the terminal). IMS helps in managing this PS buffer by keeping messages that cannot be displayed on the IMS message queue; this allows the PS buffer to be loaded and the message resent.

Multiple output messages can be queued for the same device. IMS uses a set algorithm to determine which message to send to the terminal next. For example, response-mode output is always sent before any nonresponse-mode output messages are sent; this is very useful for programmed symbols, because it allows a transaction to be entered to load the PS buffer before other queued messages are sent to the terminal.

Provide a response-mode transaction to load the PS buffer. This transaction allows its output to overtake any nonresponse-mode output that requires the PS buffer to be loaded. A database keyed by LTERM name can be used to record the program symbol set name that is to be loaded into the buffer. All IMS applications using program symbols need to update this value whenever the PS buffer is loaded. In this way the terminal operator never needs to know the program symbol name.

Output messages using PS should be sent as nonresponse-mode, allowing the PS buffer to be reloaded using a special response-mode transaction (described in the previous paragraph). An alternative is to use a non-IMS program to load the PS buffer.

An Automated Operator Interface (AOI) routine can be triggered by the DFS2078 message, indicating a possible need to reload the PS buffer. This avoids the need for any human intervention. If you choose this technique, provide for invalid or corrupted 3270 data streams that produce the DFS2078 message, rather than an actual PS buffer problem. Be careful to design the AOI routine to detect and avoid looping if it automatically reactivates the terminal.

Remote 3270 errors (VTAM)

A remote 3270 can inform IMS about certain unusual conditions that it detects during its operation. When IMS receives a message of this type, called a sense-status message, from a 3271, 3274, 3275, or 3276, IMS determines if it indicates an error or some normal condition. If the message indicates an error, IMS formats the message for printing and sends it to the z/OS system console.

The sense-status messages, in general, indicate the following 3270 conditions:

- Device end (for example, operation complete on printer)
- Intervention required (for example, out of paper or 3270 display power off)
- Busy (operator action with unprotected screen)
- Hardware-detected errors (for example, data check or control check)

An intervention-required condition can be caused when:

- A display station is turned off (not on 3275)
- A printer is turned off or has its covers open
- A printer is without paper
- A security keylock is in the off position

When IMS detects an intervention-required condition, it suspends output to the unit. When you make the unit ready (by turning power on or by putting paper in the printer, for example), IMS resumes output. If a message is in progress when IMS detects the intervention-required condition, IMS resends that message from the beginning. For a 3270 printer, IMS can only detect an intervention-required condition that is caused by a power-off condition when IMS attempts to select that device for printing.

To determine which printers and displays are in an intervention-required state on a remote 3270 line, enter a /DISPLAY LINE PTERM or /DISPLAY NODE command. Those units in an intervention-required state display as COMPINOP (component inoperative).

IMS handles the printer on the 3275 (3284-3) differently when it detects the intervention-required condition. IMS marks the printer as having a component inoperable. However, the 3275 display continues to receive output messages and retains all input capabilities. Operations on the 3284-3 printer resume when you make the printer ready.

IMS retries 3271, 3274, 3275, and 3276 hardware-detected errors several times before it considers the error to be permanent. If you do not correct the error, IMS sets the physical terminal to pstopped and inoperative status and considers it unformatted.

IMS tries to pinpoint the cause of any hardware failures. It tries to eliminate the smallest component possible (for example, a terminal rather than a control unit, or a control unit rather than a line) when a problem forces elimination of resources. For a VTAM-attached 3270 terminal, IMS issues a /CLSDST command to disconnect the terminal. After you make appropriate repairs, you can restore IMS operation by issuing a /START LINE PTERM or /RSTART LINE PTERM command for those terminals that are part of the failing unit.

IMS tries to ensure that a terminal correctly receives all data sent to it. However, certain printer operations, including the Copy function, cause mechanical actions on the printer after the data has been transferred. For some errors occurring under these circumstances, no error recovery is possible, and none is attempted.

33/35 teletypewriter

You can use a 33/35 teletypewriter to connect and communicate with IMS.

Connecting to IMS using the 33/35 teletypewriter

During IMS system definition, you can specify that a 33/35 teletypewriter attached to IMS should operate in forced-response mode, nonresponse mode, or transaction-dependent response mode.

Establish a connection to IMS using the 33/35 teletypewriter as follows:

1. Press the Orig button to receive a dial tone at the originating terminal.
2. Dial the number. The computer should answer after no more than three rings and respond with IMS READY.
3. Enter the /IAM command, unless you defined an automatic signon option during IMS system definition.

A successful signon results in the DFS058I IAM COMMAND COMPLETE message. You can enter messages from the keyboard. To enter data from paper tape, first enter:

```
/RDR CR,LF,X-OFF
```

Communicating with IMS using the 33/35 teletypewriter

You can send messages to IMS using a teletypewriter. IMS sends responses to entered messages immediately only if you enter a message consisting of only an

X-OFF character. Otherwise, the terminal remains in input mode and IMS sends no output until 28 seconds have passed without any input's being received.

To send single-segment messages to IMS, use the following format:

- Begin the message with a transaction code.
- Terminate the message with the CR, LF, and X-OFF characters.

Example 1

```
transaction_code message_text CR,LF,X-OFF
```

Use the following format to send multisegment messages to IMS:

- Begin the message with a transaction code.
- Terminate each segment with the CR, LF, and X-OFF characters.
- Terminate the message with a period (.), and the CR, LF, and X-OFF characters.

Example 2

```
transaction_code segment_1 CR,LF,X-OFF
                  segment_2 CR,LF,X-OFF
                  segment_3 CR,LF,X-OFF
                  . . .
                  segment_n CR,LF,X-OFF
                           .,CR,LF,X-OFF
```

If you enter X-OFF from paper tape, follow it with at least three RO (rub-out) characters.

Disconnecting from IMS

There are several ways to disconnect from IMS. You should choose the way that applies to your situation.

To terminate a call to IMS, enter an /IAM DONE command or press the EOT (end-of-transmission) key after you receive all output.

To terminate while you are sending data, press the EOT key. This causes IMS to ignore the message being entered.

To terminate while output is being sent, repeatedly press the Break key or press the Clear key to disconnect the line. IMS saves all output not sent, unless the terminal is signed on as the inquiry-logical terminal; in this case, IMS discards output.

System console

IMS communicates with the z/OS system console using z/OS write-to-operator (WTO) and write-to-operator-with-reply (WTOR) messages.

Enter messages from the console in WTO or WTOR format. z/OS passes only the data portion of the reply to IMS. Each reply constitutes one segment of an input message. Each segment can contain either uppercase or lowercase characters. Multi-segment messages must end with a period to identify end-of-message. A period can appear either at the end of the last segment (before the end quote), or as a segment itself. Because of WTO and WTOR restrictions, IMS translates any embedded carriage return (CR) characters in output messages to blanks and ignores them.

The following data formats are acceptable. IMS WTO and WTOR reply formats vary depending on the z/OS system and the level of JES being used.

Table 48. IMS WTO and WTOR Formats

Input Location	Data Format
Command Line	R nn, '/ERESTART CHKPT 0 '
Second Line	R nn, 'BUILDQ FORMAT RS.'
Transaction Line	R nn, 'FIRST SEGMENT'
Second Line	R nn, 'LAST SEGMENT.'
Message Switch	R nn, 'DESTNAME TEXT'
Second Line	R nn, 'TEXT 2ND LINE.'

IMS does not print output messages enqueued for the system console on the terminal until you enter a message. This ensures that you are present when output messages are printed. If you do not need to enter an input message, enter a blank or period as a single-character message to cause enqueued output to be printed without input-message processing.

Table 49. IMS WTOR Example

Example	Explanation
27 DFS996I *IMS READY	IMS is ready for an output segment.
:	
:	
R 27, '.'	Reply to cause enqueued output to be printed.
DFS000I WTOR OPERATOR ARE YOU THERE?	Message 1 from output queue.
DFS000I PLEASE START TRANS XYZ...OK?	Message 2 from output queue.
28 DFS996I *IMS READY	IMS is ready for next input segment.
:	
:	

Local card reader

IMS supports a local card reader (for example, 2540 or 2502) as an input device. Before operation, however, you must enter a /START LINE command from the master terminal to activate the card reader.

You can then place the card deck in the hopper and make the reader ready. The reader remains active and able to read cards until you enter a /STOP LINE command or until an I/O error occurs.

If you press the end-of-file button, IMS enqueues the last message read for processing regardless of whether it was terminated by a slash-period (/.) delimiter. IMS then issues a read command to allow you to add another deck of cards when needed. If you do not press the end-of-file button, IMS does not enqueue the last message for processing unless it is a single-card input or IMS reads a slash-period. This allows you to fill the hopper with a partial tray of cards without determining

whether the last card is also the end of the IMS message. In this case also, a read is left outstanding so you can continue with the input deck by filling the hopper and making the reader ready.

Communicating with IMS

IMS delineates a segment by the presence of an end-of-media (EM) character (X'19', card punch 11-1-8-9) in any column of the input card. IMS ignores any text following the EM character. If no EM character occurs, IMS assumes a segment to be 80 characters in length. The character string slash-period (/.) immediately **preceding** the EM character (if present) or as the first two characters in the following input card terminates a message. This allows you to input variable-length data, even though BSAM is reading fixed-length records.

When you assign a SYSIN data stream to a local reader line, IMS reads the entire file of data and closes the device. You can restart the line using a /START command after IMS notifies you that it has processed the file.

Responding to errors

If an I/O error occurs, follow standard I/O error-recovery procedures for the device. If IMS detects invalid data, it stops the line. Correct the condition, reenter the deck beginning with the transaction that caused the failure, and restart the line from the master terminal. You should use the /SET command to identify the input destination in order to prevent data errors in the middle of input streams.

Local printer

IMS supports a local printer (for example, 1443 or 1403) as a high-speed online output terminal. After you enter a /START LINE command from the master terminal, the printer prints output as received from the application program without your intervention.

If a permanent I/O error occurs, IMS stops the terminal. Correct the error and restart the line from the master terminal using the /START LINE PTERM or /RSTART LINE PTERM command.

Replace the magnetic tape

If you use a tape for SYSIN or SYSOUT data sets, you must change tapes when z/OS notifies you of end-of-volume (EOV).

To remove a tape before EOV for offline processing:

1. Issue a /STOP LINE PTERM command from the master terminal.
2. Remove the tape, mount a new one, and to continue processing, enter one of the following commands
 - /START LINE PTERM
 - /RSTART LINE PTERM

Disk data sets

You can allocate a disk data set for use as an intermediate data file accessed by an IMS application program or any other user program. The operation is the same as for a tape file, but IMS provides no means of dismounting disk volumes.

When you allocate a group of disk data sets for the spooled SYSOUT option, the IMS Spool SYSOUT Print utility (DFSUPRT0) can copy one or more of these data sets to SYSOUT. IMS can automatically schedule this utility at end-of-volume on one of the data sets. An end-of-volume condition is also simulated:

- If an I/O error occurs while writing to a spool data set.
- During system initialization, if a spool data set was in use when the system was last running and the spool data set was not printed.

In this case, you must issue a `/START LINE` command to complete the scheduling procedure.

Recommendation: Define at least two data sets for your spool line.

You can reschedule the utility with the `/STOP LINE PTERM` command. You can also start the utility at any time by issuing a `/START REGION procname` command from the master terminal. IMS generates line numbers and PTERMs in the sequence of the input definition statements during IMS system definition. After the system is generated and ready to run, you can display line numbers to determine what line number was generated for a spool data set.

You must establish your own conventions for controlling special forms. One possible method is to allocate a logical terminal for each report. After you set up the forms for a specific report, lock the other logical terminals using the `/LOCK LTERM` command, then place the terminal in exclusive mode using the `/EXCLUSIVE` command. Thereafter, IMS prints only the data enqueued for the report you want. If you do not use a technique like this one, messages alternate among the logical terminals associated with the line.

Programmable remote systems

Operators of programmable remote systems do not communicate directly with IMS, but rather with a user-written program that resides in the remote system's storage. If you use programmable remote systems, you must provide procedures for terminal operators.

SLU-1 devices

SLU-1 devices can include communication terminals and printers.

These topics describe the following SLU-1 devices as examples:

- IBM SLU-1 Communication Terminal.
- IBM SLU-1 Data Communication System (Programmable Models). The SLU-1 supports input from a transmit data set (TDS) or a user data set (UDS), in addition to print data set (PDS) and UDS output if you define them to IMS as unattended.
- IBM SLU-1 Data Communication System (Nonprogrammable Models).

MFS is optionally available for TDS input and PDS output. You can define a SLU-1 device for either attended or unattended operation.

Printers attached to IMS as SCS SLU-1 devices

Most of the printers have Cancel, PA1, and PA2 keys that can be used when the device is operating in SNA character string (SCS) mode and is attached to an

appropriate controller, such as the 3174, 3274, 3601, or 4701. When these keys are pressed while IMS is printing a message, the following results occur:

- If the Cancel key is pressed, IMS stops transmission and discards the current output message (but does not dequeue it).
- If a PA key is pressed, a string of characters is sent to IMS. You can define this string of characters as transaction code APAK 01 for PA1 and APAK 02 for PA2, or you can use the appropriate IMS user exit routine to modify this string, for example to set it to a command or message switch.

Restriction: You must specify `CONSOLE` instead of `PRINTER1` for the `COMPT1` keyword of the `TERMINAL` macro (or the logon descriptor for the printer). In the `TYPE` macro or logon descriptor, specify the printer as `UNITYPE=SLUTYPE1`.

SLU-1 communication terminal

These topics describe IMS operations for a SLU-1 Communication Terminal or a SLU-1 Data Communication System (nonprogrammable models).

Connecting to IMS using SLU-1 devices

Secondary Logical Unit (SLU-1) devices can connect to IMS using a telephone data set.

Prepare the SLU-1 device for communication before attempting to log on. See the SLU-1 operator's guide for the proper switch settings.

You should log on to terminals that are connected with a switched line using the procedure below. Log on to terminals connected through a nonswitched line using the same procedure, but start with step 2.

1. Establish connection with the computer through a telephone data set.
2. Press the Sys Req key. The Proceed light should go on.
3. Enter the logon message. For example, to log on to an IMS system named IMSA:
 - Enter: `LOGON applid (IMSA)`
 - Response: `DFS3650I SESSION STATUS`

If you want host-initiated logon for a SLU-1 Data Communication System, you must leave the terminal in a line-to-console job. To prepare the SLU-1 terminal for host-initiated logon:

1. Press the Start Job key. For a 3774 or 3775, also press the S key.
2. Press the Line key.
3. Press the Console key.
4. Press the EOM (end of message) key.

Communicating with IMS using SLU-1 devices

You can begin communication with IMS after you complete the logon procedure. If an output message is enqueued for the terminal, IMS attempts to send it immediately after the `DFS3650I` message. If you begin to enter an input message before you receive the output message, IMS sends the output message later.

Entering single-segment input

You must remember to press the EOM key to let IMS know that you have entered a segment.

Enter single-segment input as follows:

1. Enter the input data.
2. Press the carriage-return key.
3. Press the EOM key.

Entering multi-segment input

When you set the AUTO switch to AUTO, you send data in the buffer plus an EOM to IMS when you press the carriage-return key or the Form Feed key. For a SLU-1 Communication Terminal, set the AUTO switch to OFF when you enter multi-segment input.

Enter multi-segment input as follows:

1. Enter the first segment of the input.
2. Press the carriage-return key.
3. Press the EOB (end of buffer) key to send the buffer contents, if required. If you want to enter another segment in this buffer, do not press the EOB key.
4. Enter the next input segment.
5. Press the carriage-return key.
6. If there are more segments to enter, return to step 3. If not, press the EOM key.

Receiving an output message

You can use the Cancel and Attn keys to interrupt output messages destined for SLU-1 devices.

IMS sends all output messages for a SLU-1 device in their entirety unless you interrupt them at the device using the Cancel and Attn keys as follows:

- The Cancel key stops the current message, which IMS then discards. If another output message is presently enqueued, IMS sends it.
- The Attn key causes IMS to finish sending the current message. When the message completes, you can enter an input message.
- The Attn key followed by the Cancel key stops the current message, which IMS then discards. When the message is stopped, you can enter an input message.

SLU 1 with MFS

During IMS system definition, you can define a SLU-1 terminal to work with Message Format Service (MFS). Like a 3270 terminal, system-supplied formats are available.

Even when defined to operate with MFS, the SLU-1 terminal operates in unformatted mode (using IMS basic edit) until one of the following occurs:

- You enter *//midname*.
- MFS processes an output message to the terminal using an MFS message output descriptor (MOD) that names an MFS message input descriptor (MID) to be used to process subsequent input data.

If you enter data on the same line as *//midname*, separate *//midname* from the data with a blank. IMS discards the *//midname* and the blank, and formats the data according to the named MID. If you enter *//midname* without any data on the same line, IMS considers the next line received from the terminal to be the first line of the message.

When IMS processes an output message to a SLU-1 terminal using a MOD that names a MID, MFS uses that MID to format the next input from that terminal. An

application program, the IMS /FORMAT command, a message switch, or some IMS functions can create this output message.

When in formatted mode (using MFS), a SLU-1 terminal remains in formatted mode until one of the following occurs:

- You enter two slashes (//) or two slashes followed by a blank. The terminal returns to unformatted mode, and IMS discards the two blanks. The two slashes are escape characters.
- You enter two slashes, a blank, and data. MFS returns the terminal to unformatted mode, IMS discards the slashes and the blank, and formats the data using basic edit.
- IMS sends an output message that does not name a MOD to the terminal.
- IMS sends an output message whose MOD does not name a MID to the terminal.

IMS uses MFS default formats for many output messages. When a terminal receives such a message, the terminal goes into formatted mode, and MFS processes the next input message unless you use the escape characters. If in doubt, use the escape characters.

Disconnecting SLU-1 devices from IMS

To disconnect from IMS, press the Sys Req key and then enter a logoff command. For a SLU-1 Data Communication System, you can also press the Stop Job key.

SLU-2 devices

SLU-2 devices can initiate a session with IMS in several ways.

Related reference:

 [Terminals and equipment supported by IMS Version 12 \(Release Planning\)](#)

Using SLU-2 devices to connect and disconnect to IMS

When you connect to or disconnect from IMS using a SLU-2 device, the procedure is the same regardless of the device.

If you define the terminal as the master terminal, IMS automatically enables the terminal for a session.

Initiate a session by using any one of the following three methods:

- The VTAM network operator can enter the VARY NET command (VARY NET, ID=*a*, LOGON=*b*).
- You can enter the /OPNDST command from the IMS master terminal.
- VTAM can make the unit automatically available if it is defined in the VTAM definition as belonging to IMS.

Close a terminal session by issuing a /RCLSDST command.

Related reference:

 [Terminals and equipment supported by IMS Version 12 \(Release Planning\)](#)

3290 information panel

When communicating with IMS using a 3290, the terminal operator might need to know that how the PA1 key operates is based on what option is selected for presenting pages to a partition.

If an error occurs, the 3290 session usually terminates. The 3290, not IMS, handles device errors.

Related concepts:

 [3290 information panel in partitioned format mode \(Application Programming APIs\)](#)

NTO logical units

IMS allows start-stop devices such as TTY and TTY-compatible devices (for example, the 3101) to attach to IMS through VTAM.

The Network Terminal Option (NTO) Program Product handles the device characteristics associated with line control and communication protocols, and translates them into SNA-SLU 1 communication protocols handled by VTAM and IMS. The operational characteristics of each device do not change, but session initiation and termination are modified as described below.

Related reference:

 [Terminals and equipment supported by IMS Version 12 \(Release Planning\)](#)

Connecting to IMS using NTO devices

Before you can initiate a Network Terminal Option (NTO) session, IMS must recognize the NTO device. You need to define the device to IMS and enter the /START DC from the master terminal.

For IMS to recognize the NTO device, you must do both of the following:

1. Define the NTO device to IMS during IMS system definition using the PU keyword in the TERMINAL macro. This keyword specifies the actual hardware that this NTO device represents.
2. Enter a /START DC command from the IMS master terminal.

IMS initiates a session with an NTO device when one of the following occurs:

- You enter the LOGON APPLID(*nnnnnnnn*) command at the NTO terminal.
- You enter an /OPNDST command (for nonswitched devices only) from the IMS master terminal.
- The VTAM network operator issues a VARY command with the LOGON= keyword (for nonswitched devices only).

The NTO device receives a DFS3650I message when session initiation has completed successfully.

How to communicate with IMS from an NTO device

The Network Terminal Option (NTO) handles the device characteristics associated with line control and communication protocols, and translates them into SNA-SLU 1 communication protocols handled by VTAM and IMS.

You can communicate with IMS from a network terminal option (NTO) device using control characters, single-segment messages, and multi-segment messages.

Control characters

IMS uses the following control characters to delineate messages and message segments:

. (period)

End-of-message. This character signifies the end of a multisegment message when you do not use MFS.

CR Carriage return (X'0D').

LF Line feed (X'25').

Single-segment messages

To send data to IMS, use the following format:

- Begin the message with a transaction code.
- End the message with the CR and LF characters.

Example:

```
transaction_code text CR,LF
```

Multi-segment messages

To send data to IMS, use the following format:

- Begin the message with a transaction code.
- End each segment with the CR and LF characters.
- End the message with a period (.), followed by the CR and LF characters.

Example:

```
transaction_code segment_1 CR,LF
                  segment_2 CR,LF
                  . . .
                  segment_n CR,LF
                        .,CR,LF
```

You can reverse the order of the CR and LF characters.

Operation of the NTO device is the same as for those devices without NTO and VTAM. For additional information on entering messages, receiving messages, using response mode, using MFS, and using IMS commands from the terminal, see the preceding sections on each type of physical device.

Restriction: The /IAM command is not supported for NTO devices.

Disconnecting NTO devices from IMS

You can disconnect the session between a network terminal option (NTO) device and IMS by issuing certain commands. NTO devices will also disconnect when IMS terminates or when VTAM terminates.

The session between an NTO device and IMS terminates when any of the following occurs:

- IMS terminates.
- VTAM terminates.
- The network operator enters a VARY INACT command.
- You do one of the following at the NTO terminal:
 - Issue a /RCLSDST command.
 - Enter a logoff command.
 - Disconnect the phone connection for a switched device.
- You enter the /CLSDST command from the IMS master terminal.
- Data or transmission errors force termination of the session.

Restriction: You must not stop the NTO logical unit using the IMS /STOP command. The network operator must issue the VTAM VARY command (with ANS=ON for switched devices).

Related tasks:

 [Connecting to IMS using NTO devices \(Operations and Automation\)](#)

Special operational notes and restrictions

TTY devices and devices acting as TTY devices use blank lines in the form of carriage returns and linefeeds both before and after a message to set off the message for easier identification.

For TTY devices, IMS system messages cause a carriage return (CR) and linefeed (LF) before and after the message. Other output does not cause this to happen.

For 3101 devices acting as TTY devices, IMS sends a carriage return (CR) as X'0D'. This has the same effect on the 3101 screen for IMS system messages as the carriage return (CR) and linefeed (LF) do for the TTY.

There are no component attachments to an NTO device. You must define each component as a unique NTO device to VTAM, the NTO Program Product, the NCP, and IMS.

An NTO device cannot be used as an IMS master terminal.

Master terminal

If you use a 3270 device, IMS requires that you use a 3270 display station with a 3270 printer for the master terminal (IMS does not allow the 3284 Model 3 printer). The master terminal cannot be an ETO terminal or an LU 6.2 device.

You can define any of the following devices as the master terminal for an IMS subsystem:

- System console
- IBM 3270 Information Display System (nonswitched)

- SLU-1 device (attended mode only)
- SLU-2 device

Chapter 9. Automated operations

Automated operations are tools and techniques that help you improve your installation's productivity. As your system grows more complex and your message traffic increases, automating certain tasks can increase efficiency of your system.

IMS operations can also be automated using Tivoli NetView for z/OS. Because Tivoli NetView for z/OS functions independently of IMS, it can gather information and issue commands that are not available to IMS.

You can use automated operations to:

- Minimize errors
- Increase availability
- Expedite problem diagnosis and prevention

IMS provides the following tools to help you automate your operations:

- Time-controlled operations (TCO)
- Automated Operator Interface (AOI)
- REXX SPOC API

Related tasks:

“Recovering terminal sessions” on page 221

Related reference:

 Tivoli Software Information Center

Advantages of automation

Many of the jobs that operators perform are simple, repetitive tasks, such as monitoring the system and issuing recovery commands. You can often automate these jobs and thereby free the operator for more complex activities, such as implementing operational procedures.

In the IMS environment, TCO and AOI can improve your operator productivity by:

- Improving operator productivity and accuracy. Automating operator tasks simplifies procedures, reduces operator input, and minimizes operator errors. For example, TCO can reduce operator input by automatically monitoring system status, starting message regions and telecommunication lines, and notifying users of system status.
- Expediting problem determination. An operator's primary job is to bypass or fix problems quickly. If adequate information about a problem is not available, it might not be possible for the operator to fix it quickly. Automated operations are especially suited for problem determination. For example, TCO can automatically gather necessary information and do diagnostic analysis so that a problem can be quickly identified and corrected.

Deciding what to automate

Identifying operations that can be automated and implementing them requires a sound understanding of your installation's operations. You must collect and analyze various resources to identify which tasks are good candidates for automation. Repetitive and predictable tasks are good candidates for automation.

Tip: The process of automating operations should be an iterative one. After you have developed and used automated procedures, you should evaluate them and, in the process, consider whether any new tasks can be automated.

Resources you should analyze include:

- System logs
- Problem management reports
- Record of calls to the help desk
- Operators' notes

Tools for automated operations

You can use the Automated Operator Interface (AOI), the REXX SPOC API, and the time-controlled operations (TCO) to help you automate the IMS subsystem.

IMS operations can also be automated using Tivoli NetView for z/OS. Because Tivoli NetView for z/OS is independent of IMS, it can gather information and issue commands that are not available to IMS.

IMS provides the following tools to help you automate your operations:

- Automated Operator Interface (AOI), of which there are two types:
 - Type 1, available for the IMS DB/DC and DCCTL environments
 - Type 2, available for all IMS environments (DB/DC, DCCTL, and DBCTL)
- REXX SPOC API, available in an IMSplex environment
- Time-controlled operations (TCO), available for IMS DB/DC and DCCTL environments

Related concepts:

Chapter 10, “Type-1 automated operator (AO) application program (GU, GN, CMD, and GCMD calls),” on page 355

Chapter 11, “Type-2 automated operator (AO) application program (GMSG, ICMD, and RCMD calls),” on page 365

Related reference:

 [Tivoli Software Information Center](#)

IMS Automated Operator Interface (AOI)

The IMS Automated Operator Interface (AOI) lets an application program, using DL/I calls, issue most IMS operator commands and receive command responses.

AOI lets you write an exit routine to intercept:

- IMS system messages routed to the MTO
- Commands
- Command responses

Using AOI, you can examine these messages, commands, and command responses and issue a message to any terminal or to the IMS message queue to be processed by an application program.

You can use AOI to develop commands tailored to your installation and to automate some portions of the recovery process. AOI can relieve the MTO of numerous controlling and monitoring duties. This is especially useful if you manage a large terminal network. IMS can send single segment or multi-segment messages. Your AOI exit must be able to handle both types of messages in the event that messages run over one line.

Using AOI:

- Relieves operators from repetitive command entry.
- Increases accuracy for command entry, especially when the commands have LTERM and telecommunication line control parameters.
- Helps you control IMS resources at prearranged times, such as during region startup and database image copy.
- Lets you prompt the MTO to take specific actions.
- Allows entry of commands with long operand lists. This can be important for commands that are larger than can be entered from a terminal.

Automated operator programs can:

- Supervise terminals
- Control IMS monitor and trace operations
- Help detect database problems
- Assist in recovery by issuing the commands /STOP DB, /LOCK DB, and either the /DBRECOVERY DB or UPDATE DB STOP(ACCESS)

Related concepts:

Chapter 10, "Type-1 automated operator (AO) application program (GU, GN, CMD, and GCMD calls)," on page 355

Chapter 11, "Type-2 automated operator (AO) application program (GMSG, ICMD, and RCMD calls)," on page 365

 Security for AO application programs (System Administration)

AOI functions

The IMS AOI consists of an exit routine and application program DL/I calls. The exit routine is called the *Automated Operator (AO) exit routine*. An application program that issues the DL/I calls is called an *AO application program*. The term AOI refers to the overall function, rather than to either of its two parts.

IMS calls the AO exit routine continually for system messages, commands, and command responses. The AO exit routine can:

- Examine IMS messages, commands, and command responses
- Manipulate IMS messages
- Send messages to an AO application program

By using an AO application program, you can:

- Retrieve messages sent from an AO exit routine
- Issue various IMS operator commands
- Retrieve responses to those commands

You can use the AO exit routine and application program together or independently of each other. For example:

- You could use just an AO exit routine to monitor IMS messages and delete unwanted messages from the master terminal (z/OS system console in the DBCTL environment).
- You could use just an AO application program to issue a set of IMS commands, perhaps to start programs or databases. When used alone, you start the AO application program just like any other IMS application program.
- You could use both the AO exit routine and AO application program together. The exit routine could, for example, direct a message to the application program. The application program could then do specific processing based on the message it receives from the exit routine. In this case, the AO exit routine starts the application program.

AOI can be used in any IMS environment (DB/DC, DBCTL, or DCCTL). But, depending on the environment, there are some differences in what AOI can do and what you must do to use it.

Related concepts:

“AOI and the IMS environments” on page 335

How to use AOI to complete different tasks

You can use the automated operator interface (AOI) to discard unwanted messages, set up a batch window, restart application programs after an abend, customize messages, create an audit trail, or supervise terminals.

Discarding unwanted messages

You can write an AO exit routine to prevent the master terminal (z/OS system console in the DBCTL environment) from being flooded with unwanted IMS messages. You can still log discarded messages to the secondary master terminal (by issuing the /SMCOPY command from an AO application program) in case, for example, the messages are subsequently needed for problem determination.

You can also prevent the logging of IMS system messages to the secondary master terminal by using the /SMCOPY MSG OFF command.

Messages that are logged to the secondary master can also be controlled using AOI exit DFSAOE00.

Setting up a batch window

You can write an AO application to set aside time (usually each night) for running batch jobs, that is, to set up a batch window. When you run batch jobs, you usually take the databases offline. To set up your batch window using an AO application program, you can write a multiple step job that:

- Issues a /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command to take the database offline (BMP step)
- Issues a /DISPLAY DB or QUERY DB command to determine when the database is actually offline (BMP step)
- Initiates your batch work and does routine operational tasks, such as making image copies (batch step)
- Issues a /START DB or UPDATE DB START(ACCESS) SET(ACCTYPE()) command to bring the databases online (BMP step)

Restarting application programs after an abend

When an application program abends, IMS issues a message and immediately stops the program. You can write an AO exit routine that sends the message to an AO application program that can then issue the /START PROGRAM command to allow work to continue immediately, rather than waiting for the MTO to restart the program.

Customizing messages

You can write an exit routine that customizes IMS system messages before they are routed to the master terminal. You can delete parts of messages that have no meaning for the MTO. You can reformat a message to conform to a format established at your installation. Or, you can add text to a message to tell an MTO what action to take in response to the message.

Creating an audit trail

You can use AOI to track system or user actions. For example, if you find that certain commands are not being used appropriately at your installation, you can intercept them with an AO exit routine. The AO exit routine can send a copy of the command to an AO application, which can in turn send the command, with pertinent user information, to the MTO, or copy it to the log for later analysis.

Supervising terminals

If your network is extensive, you might make one or more terminal operators responsible for terminal availability. These operators also watch for communication problems and can reassign LTERMs to alternate devices. You can write an AO application program that performs the same tasks, and in effect, functions as a supervisory terminal. Either the MTO or someone entering the appropriate transaction from a remote terminal can schedule the AO application program.

Processing of messages, commands, and command responses

The automated operator interface (AOI) process system messages, command, and command responses. AOI lets an application program, using DL/I calls, issue most IMS operator commands and receive command responses.

Because AOI can manipulate system messages, commands, and command responses, these topics define these terms and explains how AOI processes them.

Processing system messages

A *system message* is any IMS message that is not a direct response to a command. When IMS issues a system message, it sends a copy of the message to one or more of the following:

- IMS master terminal (z/OS system console in the DBCTL environment)
- z/OS system console
- IMS secondary master terminal

IMS sends all system messages destined for the IMS master terminal to the AO exit routine. IMS does not send copies of system messages destined for the z/OS system console (in a non-DBCTL environment) or the secondary master terminal to

the AO exit routine. The AO exit routine can change or delete a system message sent to the master terminal. In the DBCTL environment, the original copy of a changed or deleted message is lost.

Copies of a system message destined for the z/OS system console or the secondary master terminal, if the logging of system messages is not disabled, are sent before they are sent to the master terminal.

The logging of IMS system messages to the secondary master terminal is controlled by using the `/SMCOPY MSG ON|OFF` command.

Processing commands

A *command* is any IMS operator command. IMS sends a copy of all IMS operator commands entered from a terminal to the AO exit routine, but does not send the following:

- `/FORMAT`
- `/LOOPTEST`
- `/MSVERIFY`
- `/RELEASE`

In addition, in the IMS DB/DC and DCCTL environments, IMS does not send the following commands:

- `/NRESTART`
- `/ERESTART`

In addition, in the IMS DBCTL environment, IMS does not send the two restart commands listed above to the AO exit routine.

IMS also sends copies of the following types of commands to the AO exit routine DFSAOE00:

- Commands issued by ICMD calls in an AO application program
- IMS internally-generated commands

IMS does not execute the command until after IMS sends it to the AO exit routine.

Processing command responses

After a command is executed, IMS generates a *command response*. If an AO exit routine expressed interest in the command, IMS sends a copy of the command response to the routine. Any asynchronous system messages IMS produces as a result of a command are not command responses; they are system messages.

IMS does not send copies of command responses for the following types of commands to the AO exit routine DFSAOE00:

- Commands issued by ICMD calls in an AO application program
- IMS internally-generated commands

To process a command response, the AO exit routine must check for message segments after the first one. The `/DISPLAY`, `/RDISPLAY`, and `/RMxxxx` commands all generate multi-segment response messages.

AOI and the IMS environments

You can use one of two Automated Operator Interface (AOI) functions, depending on which IMS environment you are operating in and what you want to do. The AOI function for the IMS DB/DC and DCCTL environments is called *type-1 AOI*. The AOI function for the DB/DC, DCCTL, and DBCTL environments is called *type-2 AOI*.

Type-1 AOI uses the following DL/I calls:

- CMD to issue operator commands
- GCMD to retrieve responses to those commands
- GU and GN to retrieve messages routed to the AO application program from the AO exit routine

The AO exit routine you write in this environment is named DFSAOUE0. Type-1 AOI uses the IMS message queues for communication between an AO application and an AO exit routine.

Type-2 AOI uses the following DL/I calls:

- ICMD to issue operator commands
- RCMD to retrieve responses to those commands
- GMSG to retrieve messages routed to the AO application program from the AO exit routine

The AO exit routine you write in this environment is named DFSAOE00. Type-2 AOI does not use the IMS message queues for communication between an AO application and an AO exit routine; instead, it uses an AOI token (an eight-byte name that associates messages with AO application programs).

Although this topic pairs the AO exit routine and application program together based on the environments in which they operate, you could use a single AO application program to issue calls and retrieve messages from both AO exit routines, DFSAOUE0 and DFSAOE00. You can write such an application program to run in the IMS DB/DC or DCCTL environment, but not in the DBCTL environment.

Related concepts:

“AOI functions” on page 331

“Restart and recovery considerations” on page 370

Type-1 AOI

Type-1 AOI consists of the AO exit routine, DFSAOUE0, and AO application programs that use the GU, GN, CMD and GCMD calls. Type-1 AOI is applicable for the IMS DB/DC and DCCTL environments.

Type-1 AO exit routine (DFSAOUE0)

IMS calls the type-1 exit routine, DFSAOUE0, for:

- System messages destined for the master terminal
- A subset of operator-entered commands
- Associated command responses

DFSAOUE0 can examine these messages and commands before IMS sends or executes them. IMS does not call DFSAOUE0 for system messages destined for terminals other than the master terminal or for certain commands and command responses.

DFSAOUE0 can handle both single-segment and multi-segment messages. You can write DFSAOUE0 to do the following:

- Modify IMS system messages. In addition to modifying a message, DFSAOUE0 can add up to 20 bytes of text to the end of a message. A copy of the modified message can be sent to an alternate destination.
Messages issued by IMS after a command has been entered (for example, the DFS058"command in progress" message) are not really messages; they are command responses and cannot be modified.
- Delete an IMS system message and, optionally, send a copy of the message to an alternate destination.
- Ignore selected segments of a message or an entire message.
- Send a copy of system messages, commands, or command responses to an alternate destination.
- Send a new message to an alternate destination in response to receiving a system message, command, or command response.
- Modify a copy of a command or command response and send the copy to an alternate destination. Note that, unlike system messages, the exit routine can only modify a **copy** of commands and command responses.
- Enter a transaction to be placed on a message queue. This function lets DFSAOUE0 schedule an AO application program. For example, after receiving startup messages DFS680 and DFS994I, DFSAOUE0 could invoke an AO application program to enter a series of commands to start the network.

Restriction: In a shared-queues environment, any transactions to be processed by the local IMS subsystem (rather than using the shared queues) must be defined as SERIAL. These transactions cannot be processed by other IMS subsystems in the sysplex.

- Request the edited command buffer (when the input is a command) to determine which command parameters succeeded or failed.

DFSAOUE0 can modify or delete system messages only. It can modify or delete the system message that the original destination (the master terminal) receives. It can also send a modified copy to an alternate destination. You cannot, however, use DFSAOUE0 to modify or delete the original command or command response. You can use it to modify a copy of a command or command response that it sends to an alternate destination, but not the copy that the primary destination receives.

Your AO exit routine must be a standalone, 31-bit, reentrant module.

Specify your exit requirements by linking the exit routines you want to use in the IMS.SDFSRESL concatenation. You can specify both DFSAOUE0 and DFSAOE00.

Type-1 AO application program (GU, GN, CMD and GCMD)

An AO application program can:

- Issue DL/I calls to retrieve messages from DFSAOUE0 (or DFSAOE00)
- Issue a subset of IMS commands
- Retrieve the responses to those commands

Retrieving messages (GU and GN call)

An AO application program receives control when IMS enqueues a message with the appropriate transaction code. The AO application program retrieves the message from DFSAOUE0 by issuing GU and GN calls.

Issuing commands and retrieving command responses (CMD and GCMD)

The AO application program issues a command using the CMD call. You can only use the CMD call if you also use the I/O PCB. PCB status codes indicate the results of a CMD call.

If there are more than one command response segments, the first one is returned to the I/O area. Use the GCMD call to retrieve subsequent segments. This call, similar to a GN call, places subsequent command response segments in the I/O area.

Security

You can secure the CMD call by using RACF (or an equivalent product), the IMS Command Authorization exit routine (DFSCCMD0), or a combination of RACF and the exit routine. With the AOI= parameter on the TRANSACT system definition macro, you can specify which programs can issue operator commands. With RACF or DFSCCMD0, you can define which commands an AOI application program can issue.

Recovery

IMS puts messages sent from DFSAOUE0 to the AO application program on the IMS message queue. Therefore, all messages are recoverable at IMS restart.

Supported application program environments

The following table shows the types of application programs that can issue CMD and GCMD calls, by IMS environment.

Table 50. CMD and GCMD call support by application region type

Application region type	IMS environment		
	DB/DC	DBCTL	DCCTL
DRA thread	No	No	N/A
BMP (non-message-driven)	No	No	No
BMP (message-driven)	Yes	N/A	Yes
MPP	Yes	N/A	Yes
IFP	No	N/A	No

Related concepts:

 Type-1 AO command security (System Administration)

Type-2 AOI

Type-2 Automated Operator Interface (AOI) consists of the AO exit routine, DFSAOE00, and AO application programs that use the GMSG, ICMD, and RCMD calls. Type-2 AOI is applicable for all IMS environments (DB/DC, DCCTL, and DBCTL).

Type-2 AO exit routine (DFSABOE00)

IMS calls the type-2 exit routine, DFSABOE00, for:

- IMS system messages destined for the master terminal
- Commands entered from a terminal and the responses to those commands
- Commands issued from an AO application program using the ICMD call
- IMS internal commands

IMS does not call DFSABOE00 for:

- Command responses to internal commands
- Command responses to an AO application program ICMD call
- Commands issued from an AO application using the CMD call
- System messages for which the destination is not the master terminal

You can write DFSABOE00 to complete the following tasks:

- Modify the text of IMS system messages. DFSABOE00 can also add up to 20 bytes of additional text to the end of a message.
- Delete IMS system messages. In a DBCTL environment, if a message is deleted or modified, IMS does not keep the original message sent to the z/OS system console.
- Direct any system message, command, or command response to an AO application program.
- Filter the system messages sent to the secondary master terminal.
- Start a BMP job (for example, an AO application program). DFSABOE00 can issue SVC 34 to start a BMP. You can use the APARM parameter of the EXEC statement to pass information to the AO application program. The AO application program can issue the DL/I INQY call to retrieve the value specified for the APARM parameter. You can use the START command to override the APARM parameter of the EXEC statement.

DFSABOE00 must use IMS AOI callable services to communicate with an AO application program. AOI callable services let you:

- Direct messages, commands, and command responses to an AO application program
- Cancel messages previously directed to an AO application program but not yet sent

Your AO exit routine must be a standalone, 31-bit, reentrant module.

Specify your exit requirements by linking the exit routines you want to use in the IMS.SDFSRESL concatenation. You can specify both DFSABOE00 and DFSABOE00.

Restriction: You cannot use DFSABOE00 to modify or delete commands and command responses, including commands from the z/OS system console or an application program, or internally generated commands.

Type-2 AO application program (GMSG, ICMD, and RCMD)

An AO application program can issue DL/I calls to:

- Retrieve messages from AO exit routine DFSABOE00 (or DFSABOE00)
- Issue a subset of IMS commands

- Retrieve the responses to those commands

Retrieving messages (GMSG call)

DFSAOE00 does not use the IMS message queues to route messages to an AO application program. Thus, your AO application program can retrieve messages when you cannot or do not want to use the message queues.

The AO application program retrieves messages using the GMSG call, which associates an AOI token name with a message. The application program passes an 8-byte AOI token to IMS, and IMS returns to the program only those messages associated with the AOI token. By using different AOI tokens, DFSAOE00 can direct messages to different AO applications programs. This use of the AOI token is similar to the use of transaction names using the other AO exit routine, DFSAQUE0.

The GMSG call, specifying an 8-byte AOI token, retrieves the first segment of a message associated with the token. Subsequent GMSG calls, with no AOI token specified, retrieve the second through the last segments of a multi-segment message, one segment at a time. If a program needs all segments of a multi-segment message, it must issue GMSG (with no AOI token) until all segments are retrieved before issuing another GMSG call with an AOI token. When the program issues a new GMSG call with an AOI token, IMS discards all remaining message segments from the previous call.

An application program can specify a wait in the GMSG call. If there are no messages currently available for the specified AOI token, the application is put into a wait state until a message is available. The decision to wait is specified by the program, unlike a WFI (wait-for-input) transaction where the wait is specified in the transaction definition.

Restriction: In a shared-queues environment, any transactions enqueued to an AOI token are processed by the local IMS subsystem (rather than using the shared queues) and cannot be processed by other IMS subsystems in the sysplex.

Issuing commands and retrieving command responses (ICMD and RCMD calls)

The ICMD call issues an IMS command and retrieves the first command response segment if one exists. The RCMD call retrieves subsequent command response segments. When an ICMD call results in a multi-segment command response, you might not want to retrieve all segments. If you need all segments, you must issue RCMD calls until all segments are retrieved before issuing another ICMD call. When the program issues a new ICMD call, IMS discards all remaining message segments from the previous call.

Using AIBTDLI

The GMSG, ICMD, and RCMD calls do not require the AO application program to pass a PCB to IMS. The calls use the application interface block (AIB) interface (AIBTDLI), which lets an application issue DL/I calls without a PCB address. The AIBTDLI interface also lets IMS pass unique return and reason codes to the application in the AIB.

Security

You can secure the ICMD call using RACF (or an equivalent product), the IMS Command Authorization exit routine (DFSCCMD0), or both. RACF lets you specify which commands an application program can issue. DFSCCMD0 and RACF also allow you to check authorizations during ICMD call processing.

Recovery

The messages directed from DFSAOE00 to the type-2 AO application program are not written to the IMS message queues. Therefore, messages directed to the AO application program cannot be recovered during IMS restart.

Supported application program environments

The following table shows the types of application programs that can issue GMSG, ICMD, and RCMD calls for different IMS environments. In addition, these calls are supported for CPI Communications driven applications programs.

Table 51. GMSG, ICMD, and RCMD call support by application region type

Application region type	IMS environment		
	DB/DC	DBCTL	DCCTL
DRA thread	Yes	Yes	N/A
BMP (non-message-driven)	Yes	Yes	Yes
BMP (message-driven)	Yes	N/A	Yes
MPP	Yes	N/A	Yes
IFP	Yes	N/A	Yes

Related concepts:

 Type-2 AO command security (System Administration)

Related reference:

 IMS callable services (Exit Routines)

Summary of exit routine processing

In an IMS DB/DC or DCCTL environment, you can use both Automated Operator exit routine (DFS AOUE0) and Type 2 Automated Operator exit routine (DFS AO E00).

If both exit routines are loaded, IMS calls DFS AO E00 first. DFS AO E00 can process the message or command, or it can pass control to DFS AOUE0 to do the processing.

Processing for a system message using both AO exit routines

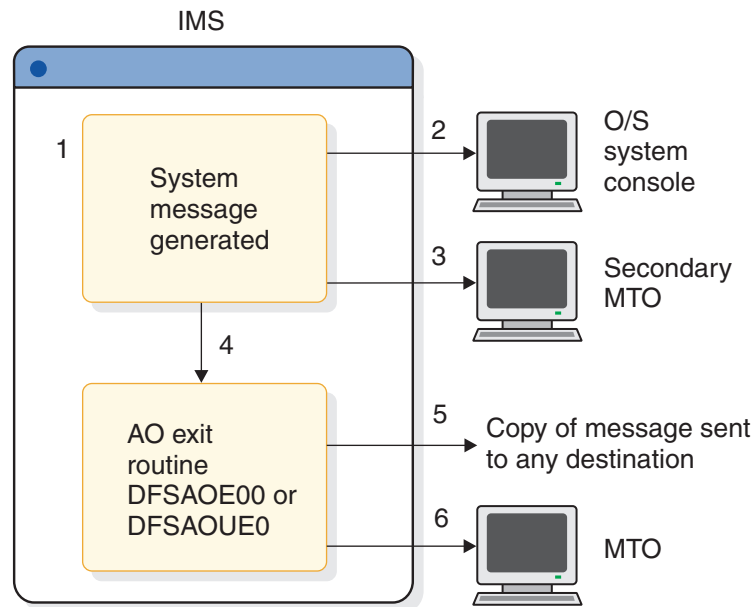


Figure 15. Processing for a system message using both AO exit routines

These steps describe the processing that occurs in the previous figure.

1. IMS generates a system message destined for the master terminal. Both AO exit routines have been specified.
2. Depending on the specific message, IMS sends a copy of the message to the z/OS system console.
3. If you specify that a copy of the message be sent to the secondary master terminal, IMS sends a copy of the message (assuming the secondary master terminal exists and that the logging of system messages to the secondary master terminal is not disabled).
4. IMS passes the copy of the message destined for the master terminal to DFSABOE00. DFSABOE00 determines which exit routine should process the message. When both exit routines are specified, IMS only calls DFSABOUE0 if DFSABOE00 determines that DFSABOUE0 should process the message.
5. The exit routine that processes the message can send a copy of it to any destination. If the exit routine is DFSABOE00, the destination is the AOI token, to which the message is enqueued. If the exit routine is DFSABOUE0, the destination is an LTERM or a transaction. The exit routine can alter or delete any segment of the message.
6. Unless the message has been deleted, IMS sends it to the master terminal.

Processing for a command using both AO exit routines

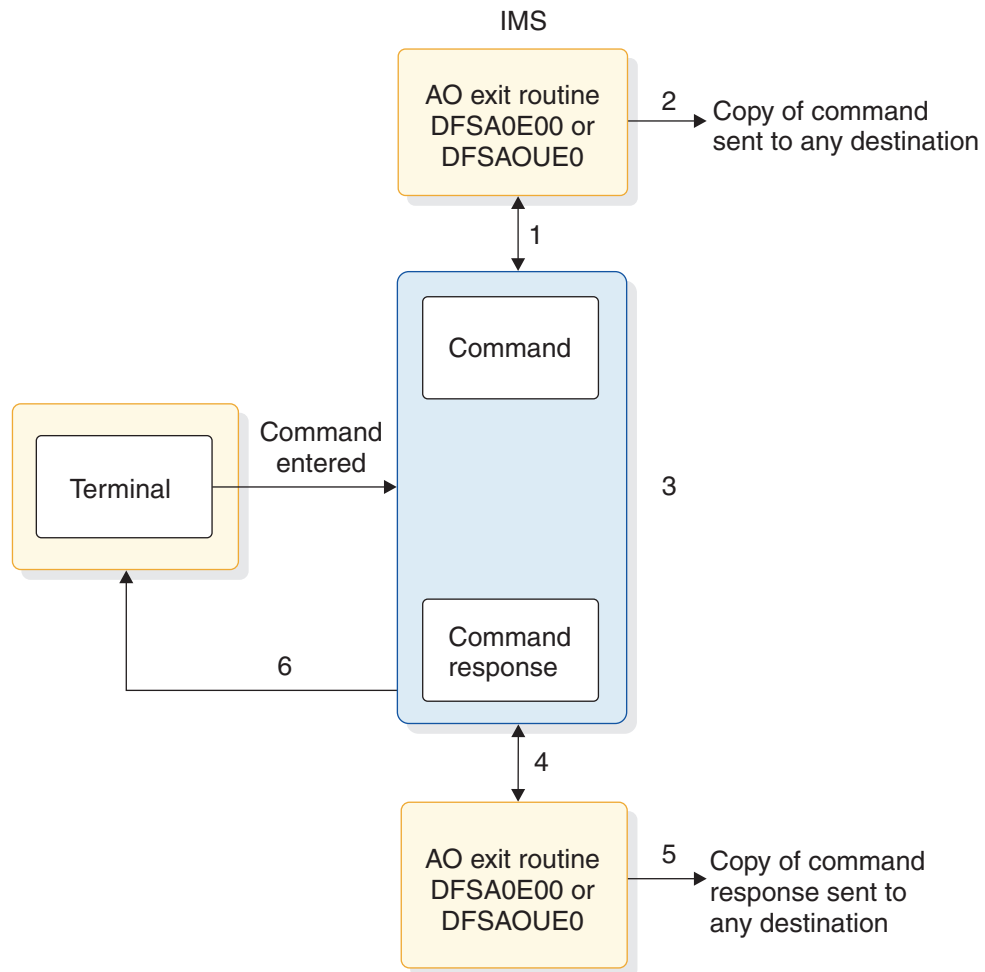


Figure 16. Processing for a command using both AO exit routines

These steps describe the processing that occurs in the previous figure.

1. When a command is entered from a terminal, IMS sends a copy of the command to DFSAOE00 before executing the command. DFSAOE00 determines whether to process the command itself or have DFSAOUE0 process the command.
2. The exit routine that processes the command can send a copy of the command to any destination. If the exit routine is DFSAOE00, the destination is the AOI token, to which the message is enqueued. If the exit routine is DFSAOUE0, the destination is an LTERM or a transaction.
3. IMS executes the command and generates a command response.
4. If the exit routine requested a copy of the command response, before sending the command response, IMS passes it to whichever exit routine processed the command.
5. The exit routine that handles the command response can send a copy of it to any destination. If the exit routine is DFSAOE00, the destination is the AOI token, to which the message is enqueued. If the exit routine is DFSAOUE0, the destination is an LTERM or a transaction.
6. IMS sends the command response to the terminal that originated the command.

Comparison of type-1 AOI and type-2 AOI

There are two different types of the automated operator interface (AOI): type-1 and type-2. You can use type-1 AOI in the IMS DB/DC and DCCTL environments. You can use type-2 AOI in all IMS environments, DB/DC, DBCTL, and DCCTL.

The following comparisons show the differences between type-1 AOI and type-2 AOI:

- Type-1 AOI uses the IMS message queues; type-2 AOI does not.
- DFSAOE00 can be used for control if messages are logged to the secondary master; DFSAOUE0 cannot.
- DFSAOE00 can process internal commands and commands from a type-2 AO application; DFSAOUE0 cannot.
- DFSAOUE0 can send a message to an alternate destination and request the edited command buffer when the message is a command; DFSAOE00 cannot.
- DFSAOE00 uses IMS AOI callable services; DFSAOUE0 does not.
- The ICMD call uses RACF (or equivalent), the Command Authorization exit routine (DFSCCMD0), or both, for security. The CMD call uses RACF (or equivalent), the Command Authorization exit routine (DFSCCMD0), or both for security.
- An AO application program issues GU and GN calls to retrieve a message from DFSAOUE0; it issues MSGG calls to retrieve messages from DFSAOE00.
- MPP and message-driven BMP regions can issue CMD and GCMD calls in the IMS DB/DC and DCCTL environments. MPP, BMP, and IFP regions, DRA threads, and CPI Communications driven applications can issue ICMD and RCMD calls in the IMS DB/DC, DBCTL, and DCCTL environments.
- Only specifically identified IMS applications are permitted to issue type-1 AOI CMD and GCMD calls. To identify IMS applications to issue these type-1 AOI calls, include the AOI= parameter on the system definition TRANSACT macro. All IMS applications are permitted to issue type-2 AOI ICMD and RCMD calls.

REXX SPOC API

The REXX SPOC API allows REXX programs to submit IMS operator commands to an IMSplex, to retrieve command responses, and to subscribe to the Operations Manager (OM) for unsolicited system messages, which can include responses to commands issued from any SPOC in an IMSplex that are routed back through OM.

The REXX language is a powerful command procedure language. Many operator automation programs are written in REXX and they can run in various environments. REXX programs, using the REXX SPOC API, have the following advantages over other IMS automation tools like AOI and TCO:

- Can issue type-2 commands to IMSplexes. REXX programs can also issue other commands, such as z/OS operator commands or TSO commands to operate the system.
- Can examine command responses, unlike TCO, and take further action based on those responses. Command responses that are returned to the REXX SPOC API are in the form of XML statements; these responses are easier to parse than the command responses that are returned to the AOI interface. You can write REXX routines to improve the self-healing and self-operating abilities of your system.

For example, a REXX program issues the following command to query the system:

```
QUERY TRAN NAME(prod001) SHOW(QCNT)
```

If the command response from the REXX SPOC API shows that work is delayed (the queue count is greater than you want), the REXX program can issue additional operator commands to resolve the problem.

- Are independent of IMS (AOI runs as an IMS application program and TCO runs as an IBM subtask), and are therefore easier to implement and run.
- Can run in Tivoli NetView for z/OS.
- Can issue /NRE and /ERE commands, unlike TCO.

Related reference:



Tivoli Software Information Center



REXX SPOC API and the CSL (System Programming APIs)

IMS time-controlled operations

The time-controlled operations (TCO) allows you to initiate time-driven procedures for any IMS operation.

TCO relieves the operator of any task that is performed routinely or at a specific time each day. TCO can issue any IMS command that an operator can issue manually, except the /NRESTART and /ERESTART commands. TCO can start and stop telecommunication lines and message regions, monitor periodic system status, and notify users of system status. TCO uses scripts that call specific exit routines that schedule transactions, execute commands, execute message switches, and deliver deferred output.

TCO is upwardly compatible with and replaces IBM's Time Initiated Input Facility (TIIF) Field Developed Program.

TCO can generate any IMS input that an IMS operator can. Specifically, TCO can:

- Execute time-initiated commands, transactions, and message switches.

Operators usually execute transactions and message switches manually at specific times. TCO can perform these functions and can also execute commands at specified times. Using TCO eliminates errors caused by incorrectly typing commands.

Restriction: TCO cannot execute the IMS restart commands, /NRESTART and /ERESTART, or initiate conversational transactions. TCO also does not support full-function response mode or Fast Path input transactions.

- Start automatically during IMS initialization and terminate when IMS shuts down.

In addition, you can:

- Start, stop, and change TCO scripts without disrupting operations.

After starting IMS with the time-initiated function, you can:

- Stop TCO processing
- Add, delete, change, or replace a TCO script
- Start processing an existing script
- Begin processing a new script
- Start TCO if it was not started automatically at IMS restart

These actions do not affect IMS operations.

- Restart the initial script DFSTCF during IMS restart.

When IMS is restarted, TCO loads the initial script, called DFSTCF, for processing. If you stop TCO before IMS terminates, TCO will not start processing when IMS restarts.

- Test TCO scripts without interfering with online execution or shutting down IMS.

Using the TCO Verification utility ensures IMS processes only error-free scripts and can do so without disrupting normal operations.

The following are possible implementations of TCO facilities:

- Starting IMS resources
- Handling peak loads
- Scheduling low-priority and non-IMS jobs
- Handling operations across multiple time zones
- Monitoring the system
- Updating user status
- Scheduling different procedures for different days
- Starting large IMS TM networks in phases to prevent resource contention or deadlocks
- Shutting down IMS

TCO components

TCO (time-controlled operations) allows you to initiate time-driven procedures for any IMS operation. TCOs consist of many components and at least one exit routine.

TCO consists of:

- An initialization routine that performs initial housekeeping and starts the TCO timer service
- A TCO timer service that processes startup requests and timer requests, calls an exit routine, and waits for future requests
- A device-dependent module that passes the input generated by the timer service to the IMS control region
- An offline script-member verification utility that ensures that IMS only processes error-free scripts
- A termination routine that performs cleanup and tells the timer service when to terminate

You also need an exit routine. This can be the exit routine IMS provides or an exit routine you write.

TCO and the exit routines execute as a subtask of the IMS control region. An error in the exit routine terminates TCO only, and does not disrupt IMS online operation.

Installing time-controlled operations

There are no IMS system definition parameters for time-controlled operations (TCO). Everything you need for TCO is automatically included when you install IMS.

In order to initialize TCO when IMS starts, include a DFSTCF DD statement in the IMS procedure. You can use the data set name IMS.TCFSLIB, or you can use another data set for your TCO scripts.

IMS defines two logical terminals for TCO: DFSTCF and DFSTCFI. These terminal definitions allow TCO input to appear as though it came from a terminal. You do not need to specify either LTERM name in your IMS stage 1 input. If you do, IMS issues a G979 message telling you that your stage 1 includes duplicate LTERM names.

You can call RACF by specifying TCORACF=Y and call the Command Authorization Exit Routine (DFSCCMD0), if it exists. To sign on a user ID to IMS, you can add a /SIGN ON command at the beginning of the TCO script and a /SIGN OFF command at the end of the script. RACF checks all subsequent commands to determine if the signed-on user is authorized to issue the commands. The commands are then passed to the DFSCCMD0 exit routine.

How to use TCO to complete different tasks

IMS loads and processes time-controlled operations (TCO) scripts and exit routines in the IMS control region. So, you can modify either the script or the TCO exit routine without stopping and restarting IMS or disrupting operations.

TCO operates by processing a script that is an unblocked 80-character member in a partitioned data set. The script contains time schedule requests and, optionally, messages that are passed to the TCO exit routine.

At the time specified in a time schedule request, IMS loads a TCO exit routine and schedules it to perform a task. If the TCO exit routine generates input, IMS passes this input to the TCO device dependent module, which then passes the data to IMS using standard DL/I interfaces. If the data is a command, IMS executes it. If it is a transaction or a message switch, IMS enqueues it for processing.

IMS displays output generated from the input, including commands, on any non-ETO device you specify. Thus, you can log all input and output for automatically-generated commands or other activity.

Automatically starting IMS resources

After IMS starts, TCO can execute any non-restart commands. It can start message regions and telecommunication lines, and automatically display the system status.

Peak load handling

The TCO exit routine can assign extra message regions or terminate message regions that you no longer need based on the time of day or the status of the IMS message queues.

Start non-IMS jobs

TCO can start non-IMS-related jobs at a specific time by issuing the /START REGION *procname* command. For example, you can specify that:

- At 8:00 p.m., TCO notifies users that an application will terminate.
- At 8:15 p.m., TCO stops transactions for the application.
- At 8:16 p.m., TCO starts the non-IMS job.

Schedule low-priority jobs

You might have some types of transactions that are processed only once or twice a day, usually by BMPs. By collecting these transactions in the message queues at 0

priority, TCO can schedule and process the transactions at less busy times. For example, assuming your operators leave for lunch at 12:00 p.m., you can specify that:

- At 12:05 p.m., TCO issues an /ASSIGN command to reassign the transactions' priority (greater than 0), so processing can begin. If the transactions need a BMP, TCO can generate one using the /START REGION command.
- At 12:50 p.m., TCO reassigns the transactions' priority to 0.

By the time the operators return from lunch, processing can continue for high-priority jobs.

Handle multiple time zones

If some of your IMS systems are in one time zone, for example United States Eastern-standard time, and other systems are in another time zone, for example United States Pacific time, the TCO exit routine can automatically start and stop terminals and other resources, taking the time difference into account.

System monitoring

TCO can automatically maintain periodic displays of the pools, message queues, and system status without operator intervention.

Update user status

You can automatically update the status or location of user terminals on a regular basis.

Start large networks in phases

TCO can start a large teleprocessing network in phases to prevent IMS, VTAM, or NCP resource overloads. By issuing a command for a few nodes on each physical line rather than all nodes on a few physical lines, the network starts more quickly because of increased parallel activity.

Starting large networks in phases example

While IMS is starting, TCO can issue an /OPNDST command to start 100 nodes. A minute later, TCO can issue another /OPNDST command to start another 100 nodes. Two minutes after startup, TCO can start another 100 nodes, and so on, until the entire network is started.

Shut down IMS

The TCO exit routine can notify users that IMS will shut down, and, for example, 10 minutes later, it can stop the telecommunication lines and message regions. TCO can even issue the /CHECKPOINT command to shut IMS down.

Time-controlled operations exit routine

The time-controlled operations (TCO) exit routine inserts messages on the IMS message queue for processing. The messages are the commands, transactions, and message switches that you specify in the time schedule requests and message sets that make up a TCO script.

The TCO exit routine passes any data found in columns 56 through 71 of the time schedule request to IMS for processing.

You do not need to write your own exit routine. You can schedule predefined commands, transactions, and message switches at predefined times using DFSTXIT0, the TCO exit routine IMS supplies.

You must bind user-written exit routines with the TCO language interface module (DFSTDLI0) and place them in the IMS.SDFSRESL data set. A user-written exit routine can have any name.

The following example shows a user-written exit routine, MYEXIT, being bound with DFSTDLI0:

```
//XIT      JOB ...
//BIND     EXEC PGM=IEWL
//SYSUT1   DD UNIT=SYSDA,SPACE=(TRK,(20,20))
//SYSPRINT DD SYSOUT=A
//SYSLMOD  DD DSN=IMS.SDFSRESL,DISP=SHR
//INLIB    DD DSN=IMS.OBJ,DISP=SHR
//ADFSLOAD DD DSN=IMS.ADFSLOAD,DISP=SHR
//SYSLIN   DD *
           INCLUDE INLIB(MYEXIT)
           INCLUDE ADFSLOAD(DFSTDLI0)
           NAME     MYEXIT(R)
/*
```

To load and execute a user-written exit routine, a time schedule request in a TCO script that is executing must refer to the routine. The following example shows a time schedule request in a TCO script that executes the user-written exit routine, MYEXIT:

```
*TIME 1200 MYEXIT
```

- Columns 1-5 indicate that this is a time schedule request.
- Columns 7-10 indicate that the request starts at 12:00 p.m.
- Columns 12-19 indicate the name of the user exit routine, in this case MYEXIT.

TCO CNT Edit exit routine

You can use the TCO CNT Edit exit routine DFSTCNT0 to restrict users (LTERMs) from loading TCO scripts.

When IMS loads a TCO script, IMS calls the CNT exit routine to validate the load request. If the routine rejects the load request, IMS sends an error message to the terminal attempting to load the script.

You do not need to write your own exit routine if you do not want to restrict any users from loading TCO scripts. IMS provides a default exit routine, DFSTCNT0, that is always called unless it is replaced with a customized exit routine. The default exit routine allows all users to load TCO scripts.

You can write your own exit routine to authorize certain users (LTERMs) to load TCO scripts. You must name the exit routine DFSTCNT0 and replace the default routine.

IMS uses the DFSTCNT0 exit routine instead of the Message Switching (Input) Edit exit routine when the message switching destination is DFSTCF.

IMS does not use RACF to determine authorization loading for TCO scripts.

Time schedule requests

A time schedule request is a statement in a TCO script that tells IMS to process a particular task at a specified time. A time schedule request also specifies which exit routine is scheduled and what data should be passed to the exit routine.

A time schedule request can specify:

- A time of day, by hour and minute.
- A range of times, for example, every 10 minutes from 10:00 to 16:00.
- IMS startup. Requests are scheduled immediately after IMS restart, or when a new script is loaded.
- A specified delay after IMS startup, for example, 10 minutes after startup.

A time schedule request either contains 16 bytes of data or refers to a variable-length message set that contains one or more single segment or multi-segment messages.

The following figure shows two message sets and time schedule requests for broadcasting a message. The example also shows a scale to indicate the alignment of the fields.

.....1.....2.....3.....4.....5.....6.....7.....8	
/BRO LTERM CTRL	S00000001
THE IMS TIME-CONTROLLED OPERATION IS UP.	00000002
*TIME DFSTXIT0 S ****	00000003
/BRO LTERM CTRL	S00000004
FIVE MINUTES HAVE ELAPSED SINCE STARTUP, AND TCO IS ACTIVE.	00000005
*TIME DFSTXIT0 0005 S ****	00000006

In this example, the first message set is:

```
/BRO LTERM CTRL
THE IMS TIME-CONTROLLED OPERATION IS UP
```

The presence of the S in column 72 indicates the end of a segment and that the next segment is part of the same message. If there are any errors in this message set, the numbers will be added to the error message. The time schedule request indicates that this broadcast should take place as soon as the script is initialized.

The second message set in this example is:

```
/BRO LTERM CTRL
FIVE MINUTES HAVE ELAPSED SINCE STARTUP, AND TCO IS ACTIVE
```

The associated time schedule request indicates that this message should be broadcast 5 minutes after the script is initialized.

Required fields for time schedule requests

The required fields for a time schedule request specify the following: the time you want the task to be performed, and to which exit routine to pass the script.

- Columns 1-5 (ID) IDENTIFICATION
Must contain the characters *TIME. All other requests are treated as message sets.
- Columns 12-19 (NAME) USER EXIT ROUTINE NAME

Must include the one- to eight-character name of the exit routine in the IMS.SDFSRESL data set. If the name has fewer than eight characters, it must be left-justified and padded to the right with blanks. NAME must begin in column 12.

Optional fields for time schedule requests

By specifying one or more optional fields, you can request tasks to be processed during IMS startup or at a specified time following startup. You can also specify for processing to start, stop for an interval, and then restart.

- Columns 7-10 (HHMM) START TIME

The time the task should start. A start time is required unless the request is for startup. Valid values for HH are 00 to 23, and valid values for MM are 00 to 59.

- Columns 21-24 (HHMM) END TIME

The time the task should end. Use this field only if you specify an interval time in columns 26-29. Otherwise, the field is ignored. This time must be a later time than the start time specified in columns 7-10. Valid values for HH are 00 to 23, and valid values for MM are 00 to 59.

If you request multiple dispatching for a task, the end time also stops the multiple dispatching. If you do not specify an end time, multiple dispatching continues until the end of the day (2359).

- Columns 26-29 (HHMM) INTERVAL TIME

This field specifies the amount of time you want to defer a request. The interval time can be any time after the start time (or startup) to a time before this request is to be dispatched. Valid values for HH are 00 to 23, and valid values for MM are 00 to 59.

- Column 31 (Flag) RESIDENT INDICATOR

Specifies whether the exit routine is resident in memory. If you specify any request for an exit routine as resident, all requests use the resident copy. The options for this column are:

Blank The exit routine is resident and paged in and out as required.

D The exit routine is loaded each time a request is generated.

There is an advantage and a disadvantage to specifying the exit routine as resident:

- Advantage: the request is performed more quickly because the system does not have to load the routine each time it is needed to process a time schedule request.
- Disadvantage: the resident copy is always used; IMS does not load a new copy. If you want to load a new copy of the exit routine, you must either change the resident indicator on the time schedule request statements to “D”, or reload the TCO script using DFSTCF LOAD.

- Column 32 (Flag) DISPATCH INDICATOR

Specifies when the request is dispatched (started). The options for this column are:

Blank The request is dispatched each day at the start time specified in columns 7-10. The start time in columns 7-10 is required.

O The request is dispatched the first day the script member is loaded. It is not dispatched again until the TCO script is reinitialized. A start time in columns 7-10 is required.

S The request is dispatched as soon as the script member is initialized,

either when IMS starts or when the script member is loaded. The start time, if specified in columns 7-10, is ignored.

If you specify an interval time in columns 26-29, the request is delayed from startup before the request is scheduled. With the "S" option, the request is processed only once each time IMS is initialized or the script-member is reloaded.

- Columns 33-35

Reserved

- Columns 56-71 MESSAGE FLAG

This field includes 16 bytes of data to be passed to the exit routine, plus the time of day. You can choose any data, or you can leave the field blank. If you set this flag to ****, the previous message set, if any, is passed to the exit routine instead of the 16 bytes found in the time schedule request.

- Columns 73-80 LINE NUMBER

Sequence numbers for the requests. If IMS or the TCO Verification utility detects an error, the number found in columns 73-80 is included in the error message to help you to identify the cause of the error.

Recommendation: Include sequence numbers for all commands and message sets.

Message sets

A message set is made up of one or more messages in a TCO script-member to be passed to an exit routine for processing. A message can consist of one or more message requests.

A message request is an 80-character string that defines data to be passed to the exit routine specified in a time schedule request.

Required fields for message sets

- Columns 1-71 MESSAGE DATA

IMS commands to be executed. Columns 1-5 must not contain "*"TIME" (unless you want to specify a timed schedule request).

- Column 72 CONTINUATION

Allows you to include single-segment or multi-segment messages and allows segments or messages to be longer than 71 characters.

The options for this column are:

Blank Indicates the end of the segment and the end of the message. The exit routine assumes that the next statement begins a new message. At the end of the message, you must include a time schedule request with "****" in columns 56-59.

S Indicates the end of the segment only (that is, that the next statement begins a new segment of the same message).

X Indicates the segment is continued in the next statement. Multiple statements are combined before the segment is passed to the exit routine. Only nine statements can be continued, for a total of 10 statements and 710 characters. You must be careful not to exceed the large message queue buffer size.

Optional field for message sets

- Column 73-80 LINE NUMBER

Sequence numbers for the requests. If IMS or the TCO Verification utility detects an error, the number found in columns 73-80 is included in the error message to help you to identify the cause of the error.

Recommendation: Include sequence numbers for all commands and message sets.

When IMS reaches the end of the TCO script, or at the end of the day, IMS starts again at the beginning of the script. Time schedule requests for startup are not re-executed, unless you reload the script. Requests with the dispatch indicator set to O (one-time execution) are not rescheduled unless the script is reloaded.

Statements do not need to be in time sequence in the script, nor do startup requests need to be first in the file. If TCO finds multiple startup schedule requests, or more than one time schedule request for the same time, TCO processes them in the order they are listed.

TCO Verification utility

Use the TCO Verification utility, DFSTVER0, to ensure your time-controlled operations (TCO) scripts are error free. You should run the utility before you execute any script online.

The utility detects any script errors that would occur during online execution, except an error caused by insufficient storage.

You must add a TCO script to the script library before you can execute the TCO Verification utility. You can verify more than one script at a time by assigning an input control statement for each script you want to verify.

Related reference:



Time-Controlled Operations Verification utility (DFSTVER0) (System Utilities)

LOAD command and processing a TCO script

You can process a new time-controlled operation (TCO) script without stopping (and restarting) IMS or TCO by sending a message switch to the TCO LTERM, DFSTCF. An operator, a program, or an exit routine can issue the command to load a new script.

When you loaded a new script, and IMS processes it, IMS disregards all previous schedule requests.

Being able to load TCO scripts dynamically allows you to do the following:

- Change or correct TCO scripts without stopping IMS
- Chain multiple scripts, so that they are processed on different days

The format of the load command is:

►►—DFSTCF—LOAD—*membername*—OUTPUT—*lterm*—CONT—*count*—►►

LOAD *membername*

Specifies the load request and the one- to eight-character script member name.

OUTPUT *lterm*

Specifies that you want IMS to send the output from the command to a

specified LTERM, rather than the default LTERM, the IMS MTO. *lterm* specifies the one- to eight-character LTERM name to which IMS sends output.

CONT *count*

Specifies how many segments (continuation lines) a message can have. The default is 9. Use this keyword to change the maximum size of a message. Valid values for *count* can be 1 to 99.

You must include a blank space between each keyword and parameter.

Chapter 10. Type-1 automated operator (AO) application program (GU, GN, CMD, and GCMD calls)

You can write type-1 automated operator (AO) application programs that will retrieve messages from the AO exit routine DFSAOUE0 (GU and GN calls), issue a subset of IMS operator commands (CMD call), or retrieve command responses for the commands that are issued on the CMD call (GCMD call).

Related concepts:

"IMS Automated Operator Interface (AOI)" on page 330

"Restart and recovery considerations" on page 370

Chapter 11, "Type-2 automated operator (AO) application program (GMSG, ICMD, and RCMD calls)," on page 365

"Tools for automated operations" on page 330

About the type-1 AO application program (GU, GN, CMD, and GCMD calls)

The Automated Operator Interface (AOI) consists of several IMS functions that allow installations to better monitor and control IMS activities. These functions can be used separately or in conjunction with user-written programs and the master terminal operator.

For example, one function includes a system definition option and an operator command that gives the master terminal operator more control of the command messages sent to the secondary master terminal.

An AO application can issue a subset of IMS commands normally reserved for the master terminal operator and receive responses to those commands by using the following DL/I calls:

- CMD (to issue commands and receive the first command response segment)
- GCMD (to receive subsequent command response segments)

If there are any response segments, the first is returned to the user's I/O work area. Subsequent segments are retrieved with GCMD. GCMD, similar to a get next (GN), places subsequent command response segments in the user's I/O work area.

The AO application receives control from IMS whenever a message with the correct transaction code has been queued for processing.

The following figure shows an AO application issuing a command with the CMD call (1) and receiving the initial response segment to the command (2). The AO application uses the GCMD call (3) to obtain additional response segments (4), if they exist.

Another AOI function lets an AO application retrieve messages. GU and GN calls retrieve messages. The retrieved message can be from either:

- A terminal from which the transaction was entered
- AO exit routine DFSAOUE0

An AO application using this function typically retrieves messages from the AO exit routine and then takes specific actions based on the content of the message.

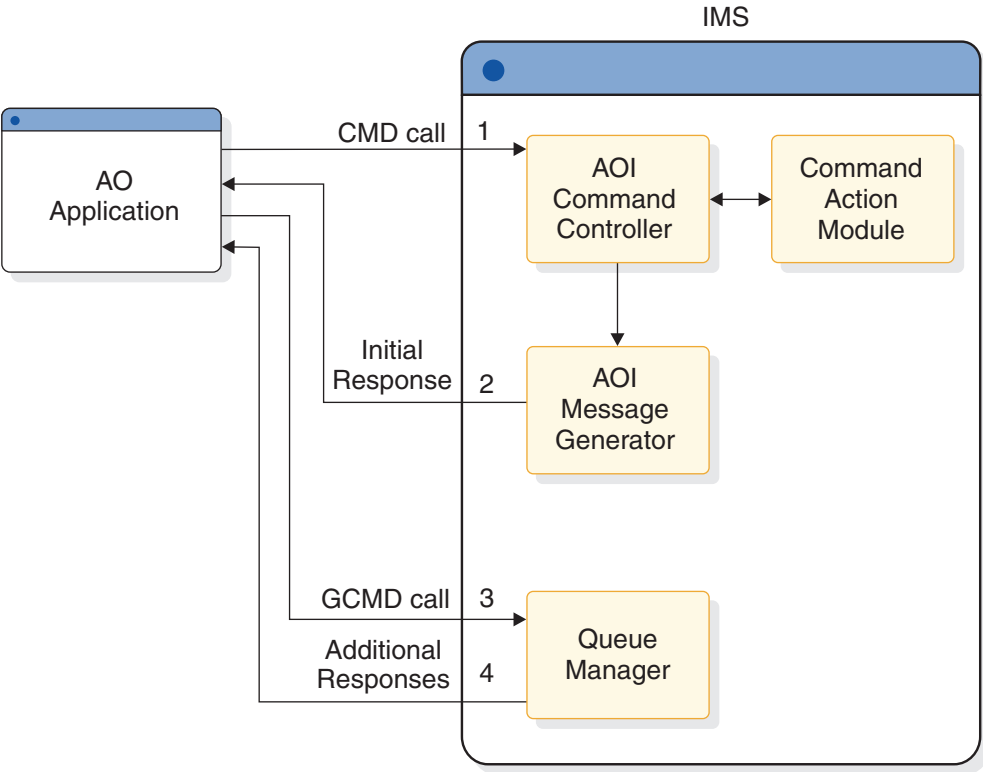


Figure 17. AO application processing

Supported application program environments

The types of automated operator (AO) applications that can issue GU, GN, CMD, and GCMD calls are listed by IMS environment.

Table 52. GU, GN, CMD, and GCMD call support by application region type.

Application region type	IMS environment		
	DBCTL	DB/DC	DCCTL
DRA thread	No	No	N/A
BMP (non-message-driven)	No	No	No
BMP (message-driven)	N/A	Yes	Yes
MPP	N/A	Yes	Yes
IFP	N/A	No	No

AO applications for shared-queues

Some transactions must reside on the same IMS subsystem as DFSAOUE0 to process correctly. If your installation uses shared-queues, define these local transactions as SERIAL to guarantee that they are processed on the local IMS subsystem. A transaction that is not defined as SERIAL can be processed on any IMS subsystem that has that transaction defined.

Commands used with AO applications (CMD)

To receive responses to these commands, use the GCMD call.

Format of commands

The syntax and synonyms for commands are the same whether the command is entered using a CMD call or from a terminal. All messages returned in response to a CMD call or a command entered from the terminal are the same.

The general format of your I/O work area on a CMD call is:

LLZZ/verb KEYWORD1 P1 KEYWORD2 P2, P3. Comments

LL Two-byte field containing the length of the command text, including **LLZZ**

ZZ Two-byte field reserved for IMS

/ Indicates that an IMS command follows

verb The command you issued

KEYWORD_x

Keywords that apply to the command you issued

P_x Parameters for the keywords you specified

. (period)

End of the command

The length of a command is limited by the size of the I/O area; the size is specified in the IOASIZE parameter in the PSBGEN macro during PCB generation. LL is the length of the command text. The size of the I/O area is the length of the actual command text, plus 4 bytes for LLZZ. The minimum size of the I/O work area is 132 bytes. The fifth byte must be a slash (/), and the verb must follow immediately. The /BROADCAST and /LOOPTEST commands must have a period between the command segment and text segment, and must be preceded by an LLZZ field that includes the size of the text. Comments can be added by placing a period (.) after the last parameter.

When issuing the /SSR command, do not code an end-of-message indicator (period). If a period is used, it is considered part of the text.

Responses to commands

The execution of a command with a CMD call might result in a response that comprises one or more segments. The response is similar to that received by a terminal operator upon entering the same command. The maximum size response returned by IMS to the AO application is 132 bytes (including the 4-byte LLZZ field).

DFS messages that are the initial response to the command will be time-stamped even if the NOTIMESTP keyword was coded on the COMM macro.

No response segments

This condition is indicated by a PCB status of blanks, showing that the execution of the command was successful or is in progress and would have resulted in one of the following command responses had the command been entered from a terminal:

DFS058 XXX COMMAND COMPLETED (No EXCEPT phrase)
DFS058 XXX COMMAND IN PROGRESS.

where XXX is the command issued by the automated operator.

DFS058 COMMAND COMPLETED indicates that no exceptional conditions are reported.

When DFS058 COMMAND IN PROGRESS is issued, one or more additional messages are issued to indicate whether the command was successful. These messages are sent to the master terminal rather than to AOI. Commands such as /DBDDUMP, /DBRECOVERY, /START DATABASE, /STOP DATABASE, and /SSR can result in message DFS058, followed by one or more asynchronous messages. Asynchronous messages are not returned as responses to the AO application.

Response segments

Response segments are identical to the response to a command entered by the operator, with preceding LLZZ and carriage control (optional) and trailing carriage controls (optional). A command might have executed successfully even though a non-blank PCB status code is returned. The status code CC indicates that one or more response segments have been produced, which occurs when a command partially executes. For example, the EXCEPT condition of this message might be returned as multiple segments.

Synchronization point processing

When IMS is restarted or resets to an earlier checkpoint, it is possible that a previously-issued command will be rescheduled for execution.

When the automated operator (AO) application regains control after a CMD call, IMS produces an initial response to the command. Command processing might be incomplete because additional IMS processing initiated by the command is still in progress.

Command responses that are not retrieved by a GCMD call are discarded at a synchronization point or by another CMD call.

A *synchronization point* is reached when any of the following events occur:

- The AO application requests the next input message with a GU call, assuming the program is defined as MODE=SINGLE on the TRANSACT macro statement.
- The program makes a CHKP call.
- The program makes a SYNC call.
- The program terminates.

Therefore, when IMS is restarted or resets to an earlier checkpoint, it is possible that a previously-issued command will be rescheduled for execution. An input message, generated by an AO exit routine and destined for an AO application, could be rescheduled although the resources are no longer held by the AO application. PCB status codes indicate more precisely which condition, or combination of conditions, exists.

Format identifiers for the /DISPLAY command

When a /DISPLAY command is issued from a CSLOMCMD or CSLOMI call, a command response segment is always produced. If OPTION=AOPOUTPUT is

specified on the CSLOMCMD or CSLOMI call, the output from the /DISPLAY command also includes a format identifier (FID) on each output segment.

Your application can use the FID to determine what type of line it is processing. The mapping of the /DISPLAY output line with the FID can be found in the DISPLAY macro. A DFS message might be produced with no FID if a syntax error occurs.

The FID is 3 bytes of EBCDIC characters. It appears at the beginning of each line, after the LLZZ and any optional carriage control character. The FID indicates to an AOI program how to map the line of output. The FID can be useful if you are converting existing AOI programs to OM AOI programs. The format of the FID is:

FID character 1

Specific /DISPLAY command class (A-Z) as follows:

FID characters 2-3

Type of output line

The following are the values you can receive in FID characters 1-3 when issuing /DISPLAY using an ICMD.

FID char 1

Defines specific /DISPLAY command:

A	/DISPLAY ACTIVE, /DIS AOITOKEN, /DISPLAY APPC
B	/DISPLAY PSB
C	/DISPLAY CONVERSATION
D	/DISPLAY AREA, /DISPLAY DATABASE, /DISPLAY DESCRIPTOR
E	/DISPLAY MSNAME
F	/DISPLAY LTERM
G	/DISPLAY SHUTDOWN STATUS
H	/DISPLAY ASMT LINE/PTERM/LTERM/NODE/USER, /DISPLAY HSB
I	/DISPLAY ASMT LINK/MSPLINK/SYSID/MSNAME
J	/DISPLAY DBD
K	/DISPLAY LINK
L	/DISPLAY LINE, /DISPLAY LUNAME, /DISPLAY PTERM
M	/DISPLAY MASTER (/RDISPLAY)
N	/DISPLAY NODE
P	/DISPLAY PROGRAM
Q	/DISPLAY Q
R	/DISPLAY RTCODE
S	/DISPLAY STATUS
T	/DISPLAY TRANSACTION, /DISPLAY TIMEOVER, /DISPLAY SYSID TRANSACTION
U	/DISPLAY USER, /DISPLAY TRACE

V /DISPLAY CCTL, /DISPLAY OASN SUBSYS, /DISPLAY SUBSYS
W /DISPLAY OLDS
X Time stamp, /DISPLAY POOL DSECTS
Y /DISPLAY MODIFY
Z /DISPLAY HSSP

FID chars 2-3

Defines type of output line:

00-49 Data line
50-69 Message line
70-89 Heading line
90-98 Reserved
99 Time stamp line

The response to the /DISPLAY POOL command does not include a FID. You can get similar information about a buffer pool using a DL/I STAT call, although the STAT call returns no Fast Path information.

Related reference:

 /DISPLAY commands (Commands)

Getting messages from an AO exit routine or a terminal

Your Automated Operator (AO) application receives control whenever a message with the appropriate transaction code is queued for processing. Your AO application retrieves messages from an AO exit routine or a terminal using the GU call.

Cross-reference of commands and command response messages

When an IMS command executes properly, the program does not receive message DFS058 confirming that the command completed successfully. The commands and their response messages are listed for when an IMS command does not execute properly.

The following table associates IMS commands used in an AO application and the resulting DFS response messages. Refer to this topic when writing your AO application to know what error messages to expect and be able to plan for them in your application.

Most of the messages found in the following table are error messages that the AO application receives when an IMS command executes improperly.

Exception: This is true for all commands except the /DISPLAY, /RDISPLAY, and the /RMxxxxxx commands. When /DISPLAY, /RDISPLAY, or /RMxxxxxx is successfully completed, the AO application receives the data it requested to be displayed rather than message DFS058.

Table 53. IMS commands and DFS response messages

Command	DFS response message
/ACTIVATE	107 127 158 163 181 182 2103

Table 53. IMS commands and DFS response messages (continued)

Command	DFS response message
/ALLOCATE	107 108 121 127 158 163 182 216 1950 1952 1953 2038 3110
/ASSIGN	107 119 124 127 131 138 139 147 150 151 152 157 158 161 163 164 182 201 210 211 212 213 214 232 243 289 794 795 2034 2035 2036 2104 2105 3102 3103 3104 3105 3108 3115 3632 3636 3637 3638
/BROADCAST	099 100 102 113 115 127 147 158 163 182 2105 3633
/CHANGE	107 108 121 127 130 163 181 182 216 696 1951 1953 2038 2105 2231 2232 3458 3619 3630 3632 3633 3639 3805 3811 3812 3817 3818 3819 3826 3827 3832 3833
/CHECKPOINT	107 127 130 140 141 142 163 182 2038 2717
/CLSDST/OPNDST	107 121 127 128 130 140 147 163 181 182 2037 2081 2103 2108 2109 2477 3101 3105 3106 3107 3108 3110 3111 3112 3113 3114 3116 3630 3633 3634 3639 3862
/COMPT	107 127 163 182 2103 2104 2254 2255 2256 2257 2261 3106 3108
/DBDUMP	107 127 130 132 140 141 142 160 163 174 175 182 2038 2717 3316 3319
/DBRECOVERY	107 121 127 130 132 140 141 142 158 160 163 174 175 178 182 2038 2537 2717 3316 3319 3327
/DELETE	102 106 107 113 114 115 116 117 127 128 130 146 150 158 163 2103 2104 2105 3654
/DEQUEUE	107 113 114 115 121 127 128 130 150 151 152 158 163 182 189 216 236 237 238 239 240 1191 1192 1952 1953 2038 2103 2104 2153 2183 2185 3104 3106 3107 3108 3110 3117 3824
All DISPLAYs	074 107 127 130 163 182 1924 2038 2093
/DISPLAY APPC	127 163 182 1953
/DISPLAY ASSIGNMENT SYSID	
MSNAME	147
MSPLINK	147
LINK	147
	147
/DISPLAY AREA	116 132
/DISPLAY ASSIGNMENT NODE	
LTERM	182 2103 3106 3630 3639
LINE-PTE	114
USER	113 115 150 3108
/DISPLAY CCTL	182 216
/DISPLAY CONVERSATION	099 113 115 163 182 216 2103 3108
/DISPLAY DATABASE	116 132 3653
/DISPLAY DBD	182
/DISPLAY DESCRIPTOR	121 127 163 182 1953
/DISPLAY HSB	182 3813
/DISPLAY LINE-PTE	113 115 150 182 216
/DISPLAY LINK	147 2038
/DISPLAY LTERM	114 216 2038
/DISPLAY LUNAME	107 121 127 130 158 163 182 1953 1954 3110
/DISPLAY MASTER	
/DISPLAY MODIFY	163 182 237 3431 3443 3451 3452 3455 3460 3461
/DISPLAY MSNAME	114 216 2038

Table 53. IMS commands and DFS response messages (continued)

Command	DFS response message
/DISPLAY NODE	182 216 2103 3106 3114 3630 3639 3831
/DISPLAY OASN SUBSYS	147 182 216 3601
/DISPLAY OLDS	182 216 2038 3828
/DISPLAY PROGRAM	117
/DISPLAY PSB	147 182
/DISPLAY Q	127 147 156 163 170 216 2038
/DISPLAY RTCODE	147
/DISPLAY SHUTDOWN STATUS	134 198
/DISPLAY SUBSYS	147 182 216 3601
/DISPLAY SYSID TRAN	147
/DISPLAY TRANSACTION	146 216
/DISPLAY USER	3108 3653
/END	113 115 150 152 158 181 182 189 285 2103 2104 3106 3108 3635 3654
/EXCLUSIVE	113 115 150 152 158 181 182 189 241 2103 2104 3106 3108 3635 3654
/EXIT	107 180 181 182 183 184 189 576 577 2103 2104 2105 3106 3108 3635 3655
/FORMAT	070 182 244 3824
/IDLE	107 127 130 134 147 163 181 216 2105 2109 3633
/LOCK	107 114 116 117 127 146 158 163 182 2038 3250
/LOG	
/LOOPTEST	113 115 143 150 152 181 195 196 203
/MONITOR	099 107 113 115 130 140 150 163 182 216
/MSASSIGN	107 127 147 163 164 182 2038 2151 2152 2153 2154 2155 2156 2157 2233 2240 2244 2252 2259 2260 3200 3214
/PSTOP	099 107 113 114 115 117 130 140 146 147 150 156 163 181 182 216 232 2152 2270 2272 2273 2297 3204 3630 3633 3824
/PURGE	099 107 108 113 114 115 117 127 130 140 146 150 156 163 182 216 1953 2038 3630 3633 3824
/QUIESCE	107 130 140 158 163 181 182 216 2103 3106 3630 3632 3633 3824
/RDISPLAY	
/RM	181 182 2038 3322
/RSTART	099 107 113 115 130 140 147 150 163 181 182 216 2103 2152 3104 3106 3108 3204 3630 3633 3639 3824
/SECURE	107 127 163 182 1953
/SMCOPY	165 181 182
/SSR	182 696 2038 3601 3609 3610 3618
/START	099 107 108 113 114 115 117 127 130 132 140 146 150 156 158 163 170 175 176 178 179 182 216 232 696 1953 2010 2020 2021 2022 2023 2024 2025 2038 2103 2105 2475 3104 3106 3108 3110 3315 3316 3319 3601 3604 3609 3630 3633 3639 3798 3799 3805 3806 3807 3808 3809 3816 3817 3824 3829 3843
/STOP	099 107 108 113 114 115 117 127 130 132 140 146 150 156 158 163 170 175 176 178 179 182 216 232 696 1953 2010 2020 2021 2022 2023 2024 2025 2038 2103 2105 2109 2475 3104 3106 3108 3110 3315 3316 3319 3601 3604 3630 3633 3639 3805 3806 3808 3824 3825 3829

Table 53. IMS commands and DFS response messages (continued)

Command	DFS response message
/SWITCH	107 127 140 163 696 2038 3810 3813 3814 3824
/TEST	113 115 143 150 152 158 181 182 189 196 282 283 284 2103 2104 2158 3106 3108 3635 3654
/TRACE	107 108 113 130 147 150 158 163 182 216 237 279 280 281 314 775 1953 2028 2029 2030 2031 2032 2038 2093 2103 2104 2155 2460 3106 3108 3110 3630 3633 3639 3813
/UNLOCK	107 114 116 117 127 146 158 163 182 2038 3250 3813
/VUNLOAD	107 130 140 163 182 216

Sample AO application (UETRANS)

A sample AO application is available from the IMS library (IMS.ADFSSRC, member name UETRANS). The sample shows use of AOI by an AO application and an associated AO exit routine. The AO application is written in PL/I for the Optimizing Compiler.

The sample AO application:

- Accepts only messages originating from the exit routine (indicated by a CG PCB status code).
- Upon receipt of a DFS994 message, looks for the character string EMERGENCY. If found, it issues the command /RSTART LINE ALL. Otherwise, it issues the following commands:

```
/RSTART LINE ALL
/START DATABASE ALL
/START TRAN ALL
```
- Upon receipt of a terminal error message, it issues the command /RDISPLAY to determine on which line and PTERM is the master terminal. If the error occurred on that line, it issues a /DISPLAY LINE ALL command to determine the first available started line, and reassigns the master terminal to that line and PTERM.
- Upon receipt of an /ASSIGN command, if the response is a DFS119 message, reissues the command until it is successful. If the response is a DFS138 message, it restarts the line with an /RSTART command and reissues the /ASSIGN command. It reissues the /ASSIGN command a maximum of 20 times; if the command is not successfully completed after the 20th attempt, an appropriate message is issued.

The associated AO exit routine is assumed to:

- Pass all DFS994 (checkpoint identifier) messages to the AO application.
- Pass all terminal error messages to the AO application. The line number of the terminal in error is placed in the last 20 bytes of the buffer. Message numbers passed are DFS001, DFS025, DFS026, DFS027, DFS029, DFS072, DFS250, DFS251, DFS253, DFS260, DFS269, DFS973, DFS998.
- Pass all /ASSIGN commands that receive a DFS119 or DFS138 message to the AO application. These messages indicate that either the LTERM is busy or the line is not started and cannot be reassigned.
- Ignore all other messages.

Chapter 11. Type-2 automated operator (AO) application program (GMSG, ICMD, and RCMD calls)

You can write type-2 automated operator (AO) application programs that will retrieve messages from AO exit routine DFSAOE00 (GMSG call), issue a subset of IMS operator commands (ICMD call), or retrieve command responses for commands issued on the ICMD call (RCMD calls).

Related concepts:

“IMS Automated Operator Interface (AOI)” on page 330

Chapter 10, “Type-1 automated operator (AO) application program (GU, GN, CMD, and GCMD calls),” on page 355

“Tools for automated operations” on page 330

About the type-2 AO application program (GMSG, ICMD, and RCMD calls)

This automated operator (AO) application can help you monitor and control IMS activities. You can use it separately or with AO exit routine DFSAOE00.

This AO application can:

- Issue DL/I calls to retrieve messages from DFSAOE00 (GMSG call).
- Issue a subset of IMS operator commands (ICMD call).
- Retrieve responses to those commands from DFSAOE00 (RCMD call).

Retrieving messages (GMSG call)

When AO exit routine DFSAOE00 is used to route messages to an AO application, the IMS message queues are not used. Instead, applications that cannot access the message queues or that do not want to use the message queues retrieve the messages from an AO exit routine.

The AO application retrieves messages using the GMSG call. GMSG associates an AOI token name with a message. The application passes an 8-byte AOI token to IMS, and IMS returns to the application a message associated with the AOI token. By using different AOI tokens, DFSAOE00 can direct messages to different AO applications.

The GMSG call, specifying an 8-byte token, retrieves the first segment of the message associated with the token. Subsequent GMSG calls, with no token specified, retrieve the second through n segments of a multi-segment message. If an application needs all segments of a multi-segment message, it must issue GMSG until all segments are retrieved before issuing another GMSG call with a token specified. IMS discards all remaining message segments from the previous GMSG call when a new GMSG call specifying an AOI token is issued from the same AO application.

An application can specify a wait on the GMSG call. If there are no messages for the specified AOI token, the application is put into a wait state until a message is available. The decision to wait is specified by the application, unlike a WFI (wait for input) transaction where the wait is specified in the transaction definition.

Defining the AOI token

You define the AOI token value. You also decide how the AO application gets the AOI token value. Some possible methods for doing these follow:

- The token could be a hard-coded value understood by both DFSAOE00 and the AO application.
- A naming convention could be established that relates the name of an AOI token to an AO application. For example, an AO application with the name ABCAOI02 might use AOI token AOI02nnn.
- The AO application could issue an ICMD '/LOG data ' call with the AOI token in the data. The application would then issue a GMSG call with WAITAOI specified for the AOI token named in the LOG call.

DFSAOE00 will see the ICMD and can parse the data in the LOG command to get the AOI token name and any other information in the data. DFSAOE00 inserts messages to the AOI token name; the messages are then returned to your AO application when the GMSG call is issued.

- DFSAOE00 could create a dynamic AOI token when it processes an IMS message. Then DFSAOE00 could issue an SVC 34 to start a BMP. The START command can override the APARM parameter on the EXEC statement. The APARM parameter is a way of passing information to the AO application. To retrieve the value specified on the APARM parameter, the AO application can issue the DL/I INQY call.
- DFSAOE00 could store an AOI token in an area that both DFSAOE00 and the AO application have access to. The AO application, when it starts, could get the AOI token from this area.

Related reference:

 DL/I calls for IMS TM system services (Application Programming APIs)

Issuing commands and retrieving command responses (ICMD and RCMD calls)

ICMD issues an IMS command and retrieves the first command response segment if one exists. RCMD retrieves subsequent command response segments. When the response to an ICMD call results in a multi-segment message, you might want to retrieve all message segments.

If you need all message segments, you must issue RCMDs until all segments are retrieved before issuing another ICMD call. IMS discards all remaining message segments from the previous ICMD call when a new ICMD is issued from the same AO application.

Restriction: All messages queued to an AO application by DFSAOE00 remain in the IMS subsystem to which they are queued. Only AO applications in the local subsystem can issue a DL/I GMSG call to access the queued messages. Because these messages are not written to a share-queue structure, applications on other IMS subsystems cannot access them.

Supported application program environments

The types of automated operator (AO) applications that can issue GMSG, ICMD, and RCMD are listed. In addition, these calls are supported from a CPI-C driven program.

Table 54. GMSG, ICMD, and RCMD call support by application region type.

Application region type	IMS environment		
	DBCTL	DB/DC	DCCTL
DRA thread	Yes	Yes	N/A
BMP (non-message-driven)	Yes	Yes	Yes
BMP (message-driven)	N/A	Yes	Yes
MPP	N/A	Yes	Yes
IFP	N/A	Yes	Yes

AO application security

Security consists of controlling which applications can issue the Issue Command (ICMD) call and which commands they can specify using ICMD.

ICMD security is implemented using RACF (or equivalent), the Command Authorization Exit Routine (DFSCCMD0), or both. To specify one or both methods of securing ICMD, use the AOIS parameter in the DBC, DCC, and IMS procedures.

You need to use the following RACF commands to implement security for AO applications:

ADDGROUP
ADDUSER
PERMIT
RDEFINE

Related concepts:

 Security for AO application programs (System Administration)

Related reference:

 Parameter descriptions for IMS procedures (System Definition)

Commands used with AO applications (ICMD)

The syntax and synonyms for commands are the same whether the command is entered using an ICMD call or from a terminal. All messages returned in response to an ICMD call or a command entered from the terminal are the same. To receive responses to these commands, use the RCMD call.

Format of commands

The general format of your I/O work area on an ICMD call is:

LLZZ/verb KEYWORD1 P1 KEYWORD2 P2, P3. Comments

LL Two-byte field containing the length of the command text, including **LLZZ**.

ZZ Two-byte field reserved for IMS.

/ or CRC

Indicates an IMS command follows. CRC (command recognition character) rather than a slash (/) is used in the DBCTL environment.

verb The command you issued.

KEYWORD_x

Keywords that apply to the command you issued.

P_x Parameters for the keywords you specified.

. (period)

End of the command.

The length of a command is limited by the size of the I/O area; the size is specified in the IOASIZE parameter in the PSBGEN macro during PCB generation. LL is the length of the command text. The size of the I/O area is the length of the actual command text, plus 4 bytes for LLZZ. The minimum size of the I/O work area is 132 bytes.

The fifth byte must be a slash (/) (or CRC for DBCTL), and the verb must follow immediately. The /BROADCAST and /LOOPTEST commands must have a period between the command segment and text segment, and must be preceded by an LLZZ field that includes the size of the text. Comments can be added by placing a period (.) after the last parameter.

When issuing the /SSR command, do not code an end-of-message indicator (period). If a period is used, it is considered part of the text.

Responses to commands

The execution of a command with an ICMD call might result in a response that comprises one or more segments. The response is similar to that received by a terminal operator upon entering the same command. The maximum size of the response returned by IMS to the AO application is 132 bytes (including the 4-byte LLZZ field). DFS messages that are the initial response to the command are time-stamped even if the NOTIMESTP keyword is coded on the COMM macro.

No response segments

This condition is indicated when AIB return and reason codes are zero and the length of the message returned is zero. These indicators tell you execution of the command was successful or is in progress and would have resulted in one of the following command responses had the command been entered from a terminal:

```
DFS058 XXX COMMAND COMPLETED (No EXCEPT phrase)
DFS058 XXX COMMAND IN PROGRESS.
```

where XXX is the command issued by the automated operator.

When DFS058 COMMAND COMPLETED is issued, no exceptional conditions are reported.

When DFS058 COMMAND IN PROGRESS is issued, one or more additional messages are issued to indicate whether the command was successful. These messages are sent to the master terminal rather than to AOI. Commands such as /DBDDUMP, /DBRECOVERY, /START DATABASE, /STOP DATABASE, and /SSR can result in message DFS058, followed by one or more asynchronous messages. Asynchronous messages are not returned as responses to the AO application.

Response segments

Response segments are identical to the response to a command entered by the operator, with preceding LLZZ and carriage control (optional) and trailing carriage controls (optional). Command responses not retrieved by an RCMD call are discarded on the next ICMD call.

Format identifiers for the /DISPLAY command

Output from the /DISPLAY command includes a format identifier (FID) on each output segment. /DISPLAY commands generally return multiple lines with different formats to the AO application. Your AO application uses the FID to determine what type of line it is processing. A DFS message might be produced with no FID when there is a syntax error.

When a /DISPLAY command is entered with an ICMD call, a command response segment is always produced. Normally, the response is a multi-segment message. Each segment contains EBCDIC and line control characters. Each segment varies in length with a maximum of 132 but usually less than 80 characters.

The FID is 3 bytes of EBCDIC characters. It appears at the beginning of each line, after the LLZZ and any optional carriage control character. The format of the FID is:

FID character 1

Specific /DISPLAY command class (A-Z)

FID characters 2-3

Type of output line

The following are the values you can receive in FID characters 1-3 when issuing /DISPLAY using an ICMD.

FID char 1

Defines specific /DISPLAY command:

A	/DISPLAY ACTIVE, /DIS AOITOKEN, /DISPLAY APPC
B	/DISPLAY PSB
C	/DISPLAY CONVERSATION
D	/DISPLAY AREA, /DISPLAY DATABASE, /DISPLAY DESCRIPTOR
E	/DISPLAY MSNAME
F	/DISPLAY LTERM
G	/DISPLAY SHUTDOWN STATUS
H	/DISPLAY ASMT LINE/PTERM/LTERM/NODE/USER, /DISPLAY HSB
I	/DISPLAY ASMT LINK/MSPLINK/SYSID/MSNAME
J	/DISPLAY DBD
K	/DISPLAY LINK
L	/DISPLAY LINE, /DISPLAY LUNAME, /DISPLAY PTERM
M	/DISPLAY MASTER (/RDISPLAY)
N	/DISPLAY NODE

P	/DISPLAY PROGRAM
Q	/DISPLAY Q
R	/DISPLAY RTCODE
S	/DISPLAY STATUS
T	/DISPLAY TRANSACTION, /DISPLAY TIMEOVER, /DISPLAY SYSID TRANSACTION
U	/DISPLAY USER, /DISPLAY TRACE
V	/DISPLAY CCTL, /DISPLAY OASN SUBSYS, /DISPLAY SUBSYS
W	/DISPLAY OLDS
X	Time stamp, /DISPLAY POOL DSECTS
Y	/DISPLAY MODIFY
Z	/DISPLAY HSSP

FID chars 2-3

Defines type of output line:

00-49	Data line
50-69	Message line
70-89	Heading line
90-98	Reserved
99	Time stamp line

The response to the /DISPLAY POOL command does not include a FID. You can get similar information about a buffer pool using a DL/I STAT call, although the STAT call returns no Fast Path information.

Restart and recovery considerations

Messages directed to an automated operator (AO) application cannot be recovered at restart because the automated operator interface (AOI) does not use the IMS message queues in this environment.

When a recoverable command is issued by the ICMD call, log record X'02' is generated when the command completes. If a failure occurs, the command is not reprocessed at restart time. If a failure occurs after a command is completed but before the command response is returned, the command response is lost. Your AO application could issue the appropriate /DISPLAY command, if applicable, to determine if the ICMD call executed successfully (as is required with asynchronous commands).

Related concepts:

Chapter 10, "Type-1 automated operator (AO) application program (GU, GN, CMD, and GCMD calls)," on page 355

"AOI and the IMS environments" on page 335

Sample AO application (DFSAOPGM)

A sample AO application is available from the IMS library IMS.ADFSSMPL, member name DFSAOPGM.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS," without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies, and have been used at least once in this information:

- Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
- 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.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

| For more information about the use of various technologies, including cookies, for
| these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and
| IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the
| section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM
| Software Products and Software-as-a-Service Privacy Statement" at
| <http://www.ibm.com/software/info/product-privacy>.

Bibliography

This bibliography lists all of the publications in the IMS Version 12 library, supplemental publications, publication collections, and accessibility titles cited in the IMS Version 12 library.

For information about the locally installable version of the Information Management Software for z/OS Solutions Information Center, see <http://pic.dhe.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.dzic.doc/installabledzic.htm>.

IMS Version 12 library

Title	Acronym	Order number
<i>IMS Version 12 Application Programming</i>	APG	SC19-3007
<i>IMS Version 12 Application Programming APIs</i>	APR	SC19-3008
<i>IMS Version 12 Commands, Volume 1: IMS Commands A-M</i>	CR1	SC19-3009
<i>IMS Version 12 Commands, Volume 2: IMS Commands N-V</i>	CR2	SC19-3010
<i>IMS Version 12 Commands, Volume 3: IMS Component and z/OS Commands</i>	CR3	SC19-3011
<i>IMS Version 12 Communications and Connections</i>	CCG	SC19-3012
<i>IMS Version 12 Database Administration</i>	DAG	SC19-3013
<i>IMS Version 12 Database Utilities</i>	DUR	SC19-3014
<i>IMS Version 12 Diagnosis</i>	DGR	GC19-3015
<i>IMS Version 12 Exit Routines</i>	ERR	SC19-3016
<i>IMS Version 12 Installation</i>	INS	GC19-3017
<i>IMS Version 12 Licensed Program Specifications</i>	LPS	GC19-3024
<i>IMS Messages and Codes, Volume 1: DFS Messages</i>	MC1	GC18-9712
<i>IMS Messages and Codes, Volume 2: Non-DFS Messages</i>	MC2	GC18-9713
<i>IMS Messages and Codes, Volume 3: IMS Abend Codes</i>	MC3	GC18-9714
<i>IMS Messages and Codes, Volume 4: IMS Component Codes</i>	MC4	GC18-9715
<i>IMS Version 12 Operations and Automation</i>	OAG	SC19-3018
<i>IMS Version 12 Release Planning</i>	RPG	GC19-3019
<i>IMS Version 12 System Administration</i>	SAG	SC19-3020
<i>IMS Version 12 System Definition</i>	SDG	GC19-3021
<i>IMS Version 12 System Programming APIs</i>	SPR	SC19-3022
<i>IMS Version 12 System Utilities</i>	SUR	SC19-3023

Supplementary publications

Title	Order number
<i>Program Directory for Information Management System Transaction and Database Servers V12.0</i>	GI10-8843
<i>Program Directory for Information Management System Transaction and Database Servers V12.0 Database Value Unit Edition</i>	GI10-8943
<i>IRLM Messages and Codes</i>	GC19-2666

Publication collections

Title	Format	Order number
IMS Version 12 Product Kit	CD	SK5T-7394

Accessibility titles cited in the IMS Version 12 library

Title	Order number
<i>z/OS TSO/E Primer</i>	SA22-7787
<i>z/OS TSO/E User's Guide</i>	SA22-7794
<i>z/OS ISPF User's Guide Volume 1</i>	SC34-4822

Index

Special characters

/ACTIVATE command 24
/ASSIGN command 24, 25
 ETO 35
 terminals 24
/BROADCAST command 296
/CHANGE command 36
/CHECKPOINT command 114
 DUMPQ 116
 FREEZE 116
 PURGE 116
 QUIESCE 119
 SNAPQ 202
 using 116
/CLSDST command 119
/COMPT command 24
/DBDUMP command 25
/DBRECOVERY command
 effects 169
 HALDB 169
 stopping database access 25
/DEQUEUE command 24, 35
 dequeuing messages, responses, and
 transactions 280
/DISPLAY command 369
 CMD-generated 359
 CONV 36
 HALDB 169
 MSC 276
 terminals 24
 XRF 207
/DISPLAY TRACKING command 177
/ERESTART command
 cold start 95
 emergency restart 77, 97
 emergency restart and dynamic
 resource definition 79
 message queue recovery 227
/FORMAT command 300
/IAM command 299, 300
/IDLE command 24
/MSASSIGN command 35
/MSVERIFY command 35
/NRESTART command
 cold start 95
 message queue recovery 227
 warm start 77, 96
/RSTART LINK command 269, 271
/SET command 296
/START command 169
 APPC 90
 DC 88
 HALDB 169
 IMS 62
 REGION 81
/STOP command
 DC 113
/STOP REGION command 114
/STOP SUBSYS command 46
/TRACE command 103, 111
 IMS Monitor 111

/TRACE command (*continued*)
 options 253
 VTAM I/O Timeout facility 24

Numerics

3270 Information Display System 299
3270 keyboard 302
3270 master terminal support 307
3270 MFS bypass option 303
3270 terminal
 message-advance function 301
 page-advance function 301
3270 terminal components 300
3270 terminal copy function 311
3270 terminal function keys 312
3275
 connecting to IMS 299
 disconnecting from IMS 300
3284 Model 3 printer 313
3290 information panel 325
33/35 teletypewriter 317, 318

A

abends
 U0970 77
abnormal termination 280
ACBLIB
 failure 134
accessibility
 features ix
 keyboard shortcuts ix
ACTIVATE (/ACTIVATE) command 24
active subsystem
 how commands affect the alternate
 subsystem 214
allocation
 WADS, spare 44
ALTID
 XRF 201
AO (automated operator)
 application 357, 368
 /DISPLAY command 369
 allowed commands 357
 CMD call 357
 use 355
 command response
 no response segment 357
 size 357
 cross-reference of commands to
 responses 360
 definition 355
 described 355, 365
 DISPLAY command
 command 359
 environments 356, 367
 functions 355
 GMSG call 365
 GU call status codes 358

AO (automated operator) application
 (*continued*)
 ICMD call 366
 issuing commands 366, 367
 RCMD call 366
 recovery 370
 responses to commands 357, 368
 restart 370
 retrieving command responses 366
 retrieving messages 365
 sample application
 DFSAOPGM 371
 UETRANS 363
 security 367
 segments
 no response 357
 size 357
 synchronization processing
 occurrences 358
 rescheduling 358
 with shared queues 357
AO (Automated Operator) exit
 routine 360
AOI (Automated Operator Interface)
 application program 335, 338
 command responses 333
 commands 333
 comparison of types 343
 DL/I calls 335
 exit routines 335, 338
 functions 331
 messages 332, 333
 overview 330
 type-1 335
 type-2 338
 types, comparison 343
 uses 332
APG (automatic priority group) 81
API
 REXX SPOC 343
APPC keyword 90
APPC/IMS
 stop 114, 120
APPC/MVS
 connecting to 90
application program
 create 27
 failures 139
 recording activity for 240
 recovery 139
 restart after abend 332
archiving SLDSs 166
areas
 effect of commands on 15
 shadowing 153
 stopping 178
ARM (Automatic Restart Manager)
 restart RM 65
 SCI 63
ASSIGN (/ASSIGN) command 24, 25
 ETO 35

- ASSIGN (/ASSIGN) command
 - (continued)
 - terminals 24
- Attn key 323
- attributes
 - updating
 - dynamically 29
- audit log 10
- audit trail
 - create 332
 - operations manager (OM) 243
- auditing operational control 250
- automated operations
 - advantages 329
 - introduction 330
 - overview 329
 - TCO 344
 - what to automate 330
- Automated Operator (AO) exit
 - routine 360
- Automated Operator Interface (AOI)
 - messages 332
 - uses 332
- automatic priority group (APG) 81
- Automatic Restart Manager (ARM) 64, 100
 - policy 84
 - restart RM 65
 - SCI 63
 - usage 84
- Availability Manager (AVM)
 - starting 200
- availability requirements 244
- AVM (Availability Manager)
 - starting 200

B

- backout
 - dynamic
 - XRF 220
 - XRF 221
- backward recovery 140
- Base Primitive Environment. 169
- batch jobs
 - restarting 102
 - setting up 332
- Batch SPOC utility 10
- batch window
 - setting up 332
- block size
 - OLDS 41
- BMP (batch message processing program)
 - deferring backout of 99
 - restarting 101, 139
 - starting 82
- break key
 - disconnecting from IMS 318
- BUILDQ keyword 227

C

- calls
 - GMSG 365
 - ICMD 366
 - RCMD 366

- CANCEL command 126
- Cancel key 323
- card reader 319
- carriage return (CR) 327
- catch-up processing
 - DL/I database tracking 177
 - log router 177
 - performance, tape 177
 - RSR 177
- CCTL (coordinator controller)
 - failures 144, 145
 - thread 145
 - thread looping 145
 - reconnecting 102
 - starting 83
- CEMT command 90, 120
- change
 - online 162, 171
- CHANGE command 36
- changing modes
 - OLDS 41, 42
 - WADS 43, 44
- CHECKPOINT (/CHECKPOINT)
 - command
 - DUMPQ 116
 - FREEZE 116
 - PURGE 116
 - SNAPQ 202
 - using 116
- CHECKPOINT command
 - QUIESCE 119
- checkpoints
 - system action 114
 - XRF 202
- CICS
 - connecting to 90
 - ISC session
 - connecting to 90
 - terminating 120
- class-1 terminals
 - XRF 205
- CLDST (/CLSDST) command 119
- cleanup failure 85
- clear key
 - disconnecting from IMS 318
- clock
 - setting 71
- CMD call
 - use 355
- CMD calls 335
 - format 357
 - status codes 357
- cold start 84
 - dynamic resource definition and cold
 - start 95
 - general considerations 267
 - MSC 267
 - performing 73, 74
 - with IMSRSC repository 76
 - with RDDS 75
 - preparing 173
 - resource structures 87
 - structures 86
 - when to perform 95
- command format
 - CMD call 357

- command responses
 - TSO SPOC 5
 - type-1 and type-2 5
- command shortcuts
 - entering commands 7
- command shortcuts panel
 - TSO SPOC 6
- command status
 - TSO SPOC 6
- commands
 - /ASSIGN 24, 25, 35
 - /BROADCAST 296
 - /CHANGE 36
 - /CHECKPOINT 114, 119, 227
 - FREEZE|DUMPQ|PURGE 114
 - /CHECKPOINT FREEZE 222
 - /DBDUMP 25
 - /DBRECOVERY 25
 - /DEQUEUE 24
 - /DISPLAY 35, 36, 295, 369
 - /DISPLAY TRACKING 177
 - /ERE 218
 - /ERESTART 227, 344
 - /EXIT 36, 295
 - /FORMAT 293, 300
 - /HOLD 295
 - /MODIFY 35
 - start FDBR 151
 - /MSVERIFY 35
 - /NRESTART 227, 344
 - /RELEASE 295
 - /RMxxxxxx 226
 - /RSTART LINE 295
 - /SECURE 35
 - /SET 293, 296
 - /START LINE 295
 - /STOP 36, 113
 - /STOP BACKUP 222
 - /STOP REGION 114
 - /STOP SUBSYS 46
 - /STOP WADS 44
 - /SWITCH SYSTEM ACTIVE
 - [FORCE] 215
 - /SWITCH SYSTEM FORCE 215
 - /TRACE 103, 111
- AO application 367
- AOI 333
- automated operator 357
- CANCEL 126
- cancel remote takeover 186
- CEMT 120
- CHANGE (/CHANGE) 36
- controlling MSC 275
- create recovery points 224, 225
- DBRC 226
- DEQUEUE 35
- DFSSPOC 3
- DUMP 127
- effect on
 - area 15
 - database 15
 - logical link 15
 - logical link path 15
 - logical terminal 15
 - MSLINK 15
 - MSNAME 15
 - MSPLINK 15

- commands (*continued*)
 - effect on (*continued*)
 - node 15
 - physical link 15
 - physical terminal 15
 - program 15
 - resources 15
 - subsystem 15
 - telecommunication line 15
 - transaction 15
 - transaction class 15
 - user 15
 - emergency restart 97
 - FDBR, using 149
 - for end users 288, 296
 - for multiple resources 12
 - HALT NET 113
 - issuing
 - Control Center 1
 - IMS Control Center 1
 - TSO SPOC 1
 - MODIFY 36, 122, 124
 - FDBR status 150
 - terminate FDBR 151
 - response
 - AOI 333
 - RSTART 89
 - SETXCF 85
 - shadowing 157
 - SHUTDOWN CSLPLEX 124
 - special operating modes,
 - controlling 296
 - START 62, 63, 89
 - FDBR tracking 150
 - start alternate subsystem 200
 - START DC 88
 - starting AVM 200
 - STOP 89, 122
 - syntax
 - z/OS 124
 - to terminate TMS 176
 - TRACE CT 112
 - Transport Manager subsystem (TMS) 160
 - type-2 47
 - unauthorize subsystems 155
 - unplanned takeover 183
 - UPDATE 25, 103, 113
 - UPDATE PGM 29
 - UPDATE TRAN 31
 - UPDATE TRANDESC 31
 - XRC tracking 192
 - XRF 206
 - Z NET 113
 - commands and responses
 - message cross-reference 360
 - Common Service Layer (CSL)
 - shutting down 123
 - communicating with IMS 293
 - communication terminals
 - SLU-1 devices 321
 - component
 - shutting down 93
 - COMPT (/COMPT) command 24
 - connecting to IMS 292

- connections
 - IMS Connect
 - checking status, overview 103
 - monitoring 103
 - status by client 104
 - status by port 104
 - status for IMSplexes 105
 - status for remote IMS Connect
 - instances 105
 - status for SCI 105
 - console configuration 196
 - Control Center
 - issuing
 - type-1 commands 1
 - type-2 commands 1
 - overview 1
 - control region
 - activity
 - recording 238
 - automatic restart 70
 - failures 132
 - restarting 61
 - starting 61, 70
 - stop 114
 - control responsibility 232
 - controlling
 - IMS 1
 - conversation
 - status 36
 - conversational transactions
 - controlling 251
 - MTO 251
 - remote terminal operator 251
 - stopped 280
 - conversational transactionscode 295
 - conversations
 - accepting and starting 164
 - isolated log send 193
 - log router 193
 - terminate 36
 - coverage
 - initiate 155
 - CPI Communications driven program
 - failure 142
 - recovery 142
 - CPI Communications driven programs
 - recovery 143
 - CQS (Common Queue Server)
 - failure 85
 - failures 144
 - log recovery 144
 - resource structure cold start 87
 - restarting
 - cold start 84
 - description 83
 - warm start 83
 - restarting structures
 - allocation 85
 - shutting down 122
 - starting 81
 - starting manually 81
 - structure cold start 86
 - structure warm start 85
 - CQS0034A
 - CQS structure rebuild 87
 - create recovery points 224, 225

- CSL (Common Service Layer)
 - ODBM
 - shutting down 124
 - OM
 - shutting down 125
 - RM
 - shutting down 125
 - SCI
 - shutting down 126
 - shutting down 122, 123
 - starting 62
- CSL OM
 - restart 64
- CSLRM000
 - started procedure 64
- CSLZSHUT 122
- CTTRACE 112

D

- data communications
 - APPC/IMS 114
 - stop 114
- data entry database 82
- Data Entry Database (DEDB)
 - updates 162
- data group
 - starting 178
- data propagation
 - XRF 199
- data protection
 - IRLM 147
- data sets
 - ACBLIBx 175
 - dumps
 - available 128
 - failure 134
 - failures 133
 - FORMATx 175
 - image copy
 - install 157
 - MODBLKSx 175
 - RECON 46
 - removing spare 45
 - RECONspare 45
 - recovery utility 136
 - system
 - XRF 213
 - XRF 205, 213
 - initializing 199
- data sharing
 - FDBR (Fast Database Recovery) 148
 - recovery 146
 - coordinating with DB2 for
 - z/OS 191
 - RSR 163
- data store connections
 - status
 - IMS commands 106
- data-sharing environments
 - XRF 199
- database
 - accessing 46
 - coverage
 - initiate 155
 - create dynamically 27
 - effect of commands on 15

- database (*continued*)
 - failures 140
 - integrity
 - re-establishing 132
 - quiesce
 - recovery points 223
 - recovery 140, 228
 - RSR 158
 - RSR tracking 154
 - shadow
 - restoring 174
 - shadowing 153, 154, 157
 - starting 178
 - stopping 178
 - stopping access to 25
 - tracking 165
 - update set 159
- database recovery control (DBRC)
 - tracking 158
- Database Recovery Control (DBRC)
 - /RMxxxxxx commands 226
- database resource adapter (DRA)
 - storage 129
- database tracking
 - Fast Path 175
- DB Monitor
 - activating 110
 - data, logging 110
 - description 109
 - performance gathering 109
 - reactivating 110
 - stop 110
- DB2
 - XRC tracking 191
- DB2 for z/OS
 - coordinating recovery with 191
 - tracking
 - routing delays 192
 - XRC tracking
 - during takeover 194
- DBCTL (Database Control)
 - failures 145
 - recovery 145
 - service 223
 - shutdown 114
 - XRF 222
- DBDUMP (/DBDUMP) command 25
- DBRC (database recovery control)
 - address space, starting 70
 - database integrity 147
 - remote takeover 187
 - tracking 158
- DBRC (Database Recovery Control)
 - /RMxxxxxx commands 226
 - commands, issuing 226
- DBRECOVERY (/DBRECOVERY)
 - command 169
 - effects 169
 - stopping database access 25
- DC keyword 88
- dead letter queue 35
- DEADB (data entry database)
 - utility regions, starting 82
- defining groups
 - TSO SPOC 7
- dependent regions
 - adjusting processing load 24
- dependent regions (*continued*)
 - failures 138
 - modifying 24
 - starting 81
 - stop 114
- DEQUEUE (/DEQUEUE) command 24, 35
 - dequeuing messages, responses, and transactions 280
- descriptors
 - updating
 - dynamically 29
- devices
 - SLU-2 324
- devices supported 299
- DFS3871I 200
- DFSAOE00 338
- DFSAOE0 335
- DFSCCMD0 (IMS Command
 - Authorization exit routine) 338
- DFSERA10 (File Select and Formatting
 - utility) 108
- DFSILTA0 (Log Transaction Analysis
 - utility) 109
- DFSSPOC
 - command
 - syntax 3
- DFSTCF load command 352
- DFSTVER0 (TCO Verification
 - utility) 352
- DFSULTR0 (Log Recovery Utility) 136
- diagnostics
 - XRF 222
- disk data sets
 - allocate 321
- DISPLAY (/DISPLAY) command
 - CONV 36
 - HALDB 169
 - terminals 24
 - XRF 207
- display area 309
- DISPLAY TRACKING command 177
- displaying
 - current application name 211
 - number of primary and backup sessions 210
 - session status 210
 - terminal status 209
- DL/I
 - batch
 - jobs 163
 - remote takeover 187
 - RSR 172
 - catch-up processing 177
 - tracking 154
- DL/I database tracking
 - behavior, termination 172
 - remote takeover 187
- DRA (database resource adapter)
 - storage 129
- DRD
 - create application program 27
 - create database 27
 - create transactions 27
 - creating Fast Path routing code 27
 - deleting runtime resources 33
 - DRD and RSR 175
- DRD (*continued*)
 - querying runtime descriptor
 - definitions 34
 - querying runtime resource
 - definitions 34
 - update runtime resource
 - definitions 28
- DUMP command 127
- DUMPQ keyword 116
- dumps
 - DUMP command 127
 - formatting 127, 256
 - keeping data sets available 128
 - message queues 227
 - MODIFY command 127
 - SADMP 127
 - stand-alone 127
- dynamic application program creation
 - commands 27
- dynamic resource definition
 - cold start 95
 - emergency restart 79
- dynamic resource definition (DRD) 25
 - create application program 27
 - create database 27
 - create transactions 27
 - creating Fast Path routing code 27
 - deleting runtime resources 33
 - querying runtime descriptor
 - definitions 34
 - querying runtime resource
 - definitions 34
 - shutdown 118
 - update runtime resource
 - definitions 28

E

- EEQE (extended error queue
 - element) 140
- emergency restart 77
 - /ERE command 227
- commands 97
- dynamic resource definition
 - (DRD) 79
- failures 132
- tracking subsystem 173
- end of volume (EOV)
 - magnetic tape 320
- end user
 - tasks
 - writing 286
- end-of-transmission (EOT) key
 - disconnecting from IMS 318
- end-user
 - developing operating procedures
 - for 285
 - establishing operating instructions
 - for 287
 - problem reporting 291
- entering single segment input 323
- environments
 - AO application 356, 367
- EOM key 323
- EOT key
 - disconnecting from IMS 318

- EOV
 - magnetic tape 320
- ERESTART (/ERESTART) command
 - cold start 95
 - emergency restart 77, 97
 - emergency restart and dynamic resource definition 79
 - message queue recovery 227
- errors
 - I/O
 - IMS Monitor 112
 - log 135
 - recovery 136
 - OLDS read 135
 - OLDS write 135
 - responding to 298
 - SLDS read 135
 - WADS read 135
 - WADS write 135
- ESTAE (Extended Specified Task Abnormal Exit) routines 131
- ETR (External Time Reference) device 71
- exit routines
 - AOI 335, 338
 - Automated Operator (AO) 360
 - DFSTCNT0 348
 - ESTAE (Extended Specified Task Abnormal Exit), processing 77
 - Extended Specified Task Abnormal Exit (ESTAE), processing 77
- extended error queue element (EEQE) 140
- Extended Recovery Facility (XRF)
 - data sets
 - initializing 199
 - data-sharing environment 199
 - determining parameters 199
 - planning 196
 - session management
 - terminals 201
 - starting 199
 - terminals 218
 - tracking 205
 - USERVAR 202
- Extended Specified Task Abnormal Exit (ESTAE) routines 131
- extended terminal operation (ETO)
 - user assignments 35
- external subsystems
 - failure 46

F

- failures
 - ACBLIB 134
 - application program 139
 - CCTL 144, 145
 - thread 145
 - communication 142
 - control region 132
 - CQS 144
 - data sets 133
 - message queue 133
 - database 140
 - DBCTL 145
 - dependent regions 138
 - emergency restart 132
 - failures (*continued*)
 - IRLM 146
 - links 176
 - network 141
 - re-establish database integrity 132
 - RECON 134
 - region controller 139
 - session 142
 - system 142
 - system data sets 134
 - Fast Database Recovery
 - overview 148
 - Fast Path
 - regions
 - restarting 139
 - response mode 294
 - Fast Path database tracking
 - behavior, termination 172
 - remote takeover 187
 - Fast Path routing code
 - create dynamically 27
 - FDBR (Fast Database Recovery)
 - commands 149
 - initiate 151
 - MODIFY
 - stop region functions 151
 - overview 148
 - process 149
 - reestablish tracking after recovery 149
 - regions 159
 - status 150
 - stopping region 151
 - surveillance 149
 - tracking 150
 - stopping 151
 - using 149
 - FID (format identifier) 359, 369
 - File Select and Formatting Print utility (DFSERA10) 108
 - flowcharts
 - operating procedures 256
 - format
 - AO commands 357, 367
 - format identifier (FID) 369
 - DISPLAY command 359
 - formatting
 - MODIFY command 127
 - offline 127
 - offline dumps 256
 - screen 304
 - forms
 - incident report 240
 - log activity 241
 - utilities
 - execution 238
 - forward recovery 140
 - FPUTILstart DEDB utility region 82
 - FREEZE keyword 116
 - full function databases
 - response mode 294
 - function keys
 - 3270 terminal 312

G

- gaps
 - monitor status 177
- global command status 91
- global online change
 - supported environments 13
- global resource
 - status 99
- GMSG call 365
- groups
 - defining
 - TSO SPOC 7
- GU call status codes 358

H

- HALDB (High Availability Large Database) 169
- HALDB Online Reorganization (OLR)
 - restrictions 47
 - RSR, limitations 157
- HALT NET command 113
- high-speed online output terminal 320
- HSSP (High-Speed Sequential Processing)
 - RSR 162

I

- I/O prevention
 - XRF takeover 215
- I/O toleration
 - ending 216
- ICMD call 338
- ICMD calls
 - AO (automated operator)
 - application 366
- IDLE (/IDLE) command 24
- IFP (IMS Fast Path)
 - regions 82
 - restart 139
- image copy
 - recovery, using 146
 - RSR 153
- image copy data set
 - install 157
- IMS
 - shutting down 113
 - subsystem
 - connecting 46
- IMS Application menu 2
- IMS Command Authorization exit routine (DFSCCMD0) 338
- IMS Connect
 - connections
 - checking status, overview 103
 - monitoring 103
 - status 106
 - status by client 104
 - status by port 104
 - status for IMSplexes 105
 - status for remote IMS Connect instances 105
 - status for SCI 105
 - IMS-to-IMS TCP/IP communications
 - cleaning up an MSC logical link 265

- IMS Connect (*continued*)
 - IMS-to-IMS TCP/IP connections
 - MSC, viewing connection information 257
 - MSC, viewing link information 259
 - OTMA, viewing connection information 260
 - starting communication with an IMSplex 265
 - stopping communication with an IMSplex 264
 - viewing connection information 257
 - MSC
 - checking status, overview 103
 - cleaning up an MSC logical link 265
 - monitoring 103
 - viewing IMS-to-IMS TCP/IP connection information 257
 - viewing link information 259
 - ODBM
 - checking status, overview 103
 - monitoring 103
 - OTMA
 - connection status 106
 - viewing IMS-to-IMS TCP/IP connection information 260
 - viewing IMS-to-IMS TCP/IP connection information 257
 - XCF
 - connection status 106
- IMS Control Center
 - overview 1
- IMS Database Recovery Facility
 - compared to DFSURDB0 140
 - remote terminal commands 288
- IMS logger 161
- IMS Monitor
 - activating 111
 - description 111
 - I/O errors 112
 - log 111
 - performance gathering 109
- IMS Queue Control Facility 227
- IMS-to-IMS TCP/IP communications
 - cleaning up an MSC logical link in IMS Connect 265
- IMS Connect
 - cleaning up an MSC logical link in IMS Connect 265
 - stopping an MSC logical link 265
 - viewing connection information 257
- MSC
 - viewing connection information 257
 - viewing link information in IMS Connect 259
- operations 257
- OTMA
 - viewing connection information 260
- restarting a connection from IMS Connect 262

- IMS-to-IMS TCP/IP communications (*continued*)
 - starting IMS Connect IMSplex communications 265
 - stopping a connection from IMS Connect 261
 - stopping a send client socket connection 263, 266
 - stopping IMS Connect IMSplex communications 264
 - TCP/IP connection status 105
- IMSID
 - XRF 201
- IMSplex
 - CSL
 - starting 62
 - CSLZSHUT 122, 123
 - IMS Connect
 - connection status 105
 - modifying resources 13
 - RSR tracking 157
 - shutting down 122
 - starting 61
 - starting CSL address spaces 62
- IMSRSC repository
 - open 66
 - recovering data sets 38, 39, 40
 - start 67
 - status 65
 - stop 67
 - update process 36
- incident report forms
 - guidance 240
- initialization
 - data sets
 - XRF 199
- input message
 - entering multiple physical pages 305
- internal resource lock manager (IRLM)
 - data protection 147
 - failures 146
 - monitoring 112
 - tracing 112
- IRLM (internal resource lock manager)
 - data protection 147
 - failures 146
 - monitoring 112
 - recovery 146
 - starting 80
 - stop 121
 - tracing 112
- ISC (Intersystem Communication)
 - session, terminating 120
 - sessions
 - connecting to 90
 - users, assigning 35
- isolated log transport processing 176

- J**
 - Java dependent regions
 - starting 82
 - JCL
 - standard
 - procedures 281
 - JMP region 82

K

- keyboard shortcuts ix
- keys
 - Attn 323
 - Cancel 323

L

- legal notices
 - notices 371
 - trademarks 373
- linefeed (LF) 327
- link paths
 - control 278
 - MSC 278
- local time
 - setting 73
 - conditions 71
- locking
 - terminal keyboard 303
- locks
 - block-level 147
- log data
 - active
 - receiving 164
 - storing 164
- Log Recovery utility (DFSKRSR0)
 - SLDS 138
- Log Recovery utility (DFSULTR0) 136
- log router
 - behavior, termination 172
 - conversations 164
 - data sharing 166
 - database change records 165
 - remote takeover 187
 - reprocessing log records 175
 - restart 174
 - RSR 164
- logging
 - system logger
 - z/OS 243
- logical link
 - effect of commands on 15
- logical link (MSLINKs)
 - commands for controlling 15
- logical link path
 - verifying consistency 35
- logical link paths
 - effect of commands on 15
- logical link paths (MSNAMEs)
 - commands for controlling 15
- logical paging 301
- logical terminal
 - effect of commands on 15
- logs
 - activity, recording 241
 - controlling the characteristics of 41
 - CQS, recovery 144
 - errors 135
 - gaps 166, 177
 - IMS Monitor 111
 - position 174
 - records
 - printing 108
 - reports 109
 - records gaps 176

- logs (*continued*)
 - recovery 136
 - RSR 161
 - RSR tracking 158
 - system utilities 108
 - transport processing, isolated 176
 - truncation 183
- looping
 - thread
 - CCTL failure 145
- lost messages in MSC 267
- LTERM
 - static user assignment 35

M

- magnetic tape
 - replacing 320
- maintenance
 - XRF 211
- maintenance procedures 281
- master terminal
 - location 196
- master terminal devices
 - 3270 Information Display System 327
 - SLU-1 device 327
 - SLU-2 device 327
 - System console 327
- master-terminal format
 - operation 308
- MAXPROCS keyword 166
- MAXSLDS keyword 166
- message
 - AOI 333
 - CQS0009W 85
 - CQS0020I 83
 - CQS0031A 83
 - CQS0032A 83, 116
 - CQS0034A 87
 - CQS0102E 85
 - CQS0300I 122
 - customize 332
 - DFS2216A 110
 - DFS3872I 200
 - DFS3873I 200
 - DFS3906 256
 - DFS3906I 128
 - DFS3907 256
 - DFS3907A 128
 - DFS3907I 128
 - DFS395A 139
 - DFS681I 139
 - DFS682I 139
 - IEA793A 128
 - IEA911E 128
 - IST634I 210
 - IST635I 209, 210
 - IST809I 210
 - queues
 - using Queue Control Facility 227
- message area 309
- Message Control/Error exit routine (DFSCMUX0) 280
- Message Format Service (MFS) 309, 323
 - entering a password 301
- message formats
 - WTO 318

- message formats (*continued*)
 - WTO 318
- message processing regions
 - starting 82
- message queues
 - dumping 227
 - recovery 133, 227
- message recovery 279
- message sets
 - TCO 351
- messages
 - DFS0617I 77
 - DFS3127I 77
 - send to other operators 296
 - WTO 318
 - WTO 318
- MFS (Message Format Service) 323
- milestone processing 167
- MINSLDS keyword 166
- MNPS
 - Multinode Persistent Session 96, 97
- modified data tags
 - resetting 302
- MODIFY command 127
- modifying resources online 13
- monitoring
 - /TRACE command, using 111
 - database activity
 - FDBR 149
 - IMS 103, 253
 - IRLM (internal resource lock manager) 112
 - XRF 207
- MPP
 - starting 82
- MSASSIGN (/MSASSIGN) command 35
- MSC
 - IMS Connect
 - SCI connection status 105
- MSC (Multiple Systems Coupling)
 - /DEQUEUE command 280
 - /DISPLAY command 276
 - /RSTART LINK command 271
 - assignment 35
 - commands to control MSC 275
 - dequeuing messages, responses, and transactions 280
 - initializing 267
 - link paths 278
 - message resynchronization 143
 - physical links
 - switching TCP/IP and VTAM types 272
 - QUERY command 276
 - restarting a logical link 269, 271
 - stopping
 - terminating 268
 - TCP/IP links
 - operations, overview 273
 - restarting after errors 274
- MSLINKs
 - effect of commands on 15
- MSNAMEs
 - effect of commands on 15
 - verifying consistency 35
- MSPLINKs
 - effect of commands on 15

- MSVERIFY (/MSVERIFY) command 35
- MTO (master terminal operator) 247
 - /RSTART LINK command 267
 - abilities and responsibilities 246
 - auditing operational control 250
 - conversational transactions 251
 - description 61
 - monitoring IMS 253
 - network operations, planning for 249
 - operations, controlling 246
 - procedures
 - content 234
 - resources controlled by 249
 - tasks 151
 - termination of transmission 268
 - UPDATE MSLINK 267
- multi-segment input
 - enter 323
- Multinode Persistent Session 96, 97
- Multinode Persistent Sessions (MNPS)
 - session management
 - terminals 201
- Multiple Systems Coupling (MSC)
 - /DISPLAY command 276
 - /RSTART LINK command 271
 - assignment 35
 - commands to control MSC 275
 - initializing 267
 - link paths 278
 - QUERY command 276
 - restarting a logical link 269, 271
 - stop 268
 - TCP/IP links
 - operations, overview 273
 - restarting after errors 274
 - terminating 268

N

- NCP (Network Control Program) 195
- network
 - failures 141
 - fixing availability 245
 - ID, deleting 36
 - prerequisites to starting 88
 - shutdown 120
 - XRF 208, 218
- network considerations 218
- Network Control Program (NCP)
 - planning 195
- network terminal option (NTO)
 - devices 325
 - disconnecting 327
- Network Terminal Option (NTO)
 - device 326
- node
 - effect of commands on 15
- NRESTART (/NRESTART) command
 - cold start 95
 - message queue recovery 227
 - warm start 77, 96
- NTO devices
 - connecting 325
 - disconnecting 327
- NTO logical devices 325

O

- ODBA (Open Database Access)
 - reconnecting 102
- ODBM
 - IMS Connect
 - SCI connection status 105
- ODBM (Open Database Manager)
 - restarting 64
 - shutting down 124
 - starting 63
- offline dump formatting 127, 256
- OLDS
 - block size
 - changing 41
 - changing number of buffers 42
 - changing space allocation 43
 - mode
 - changing 41, 42
- OLDS (online log data set)
 - archiving 228
 - buffer, changing 41
 - changing characteristics of 41
 - location, changing 41
 - mode, changing 41
 - read error 135
 - recovery 137
 - write error 135
- OLR (HALDB Online Reorganization)
 - restrictions 47
- OM
 - audit trail 243
 - starting 64
- OM (Operations Manager)
 - shutting down 122
 - starting 62
- OM audit log display 10
- OM audit trail 9
- online change
 - ACBLIBx 175
 - FORMATx 175
 - introduction to 13
 - MODBLKSx 175
 - RSR 162, 171
 - shadowing 162
- online forward recovery
 - RSR 178
 - tracking site 178
- online log data set (OLDS)
 - archiving 228
 - buffer, changing 41
 - changing characteristics of 41
 - location, changing 41
 - mode, changing 41
 - read error 135
 - recovery 137
 - write error 135
- online reorganization (OLR)
 - HALDB
 - restrictions 47
- open
 - IMSRSC repository 66
- Open Database Manager (ODBM)
 - restarting 64
 - shutting down 124
 - starting 63
- operating IMS
 - recovering
 - archiving OLDS 228
 - message queues 227
 - using DBRC commands 226
 - using utilities for 228
- operating procedures
 - designing
 - using graphics 256
 - using text 256
- operations
 - automated 329, 330
 - Tivoli NetView for z/OS 330
 - controlling with the MTO 246
 - end-user
 - establishing instructions for 287
 - IMS-to-IMS TCP/IP connections
 - cleaning up an MSC logical link in IMS Connect 265
 - MSC, viewing connection information 257
 - MSC, viewing link information in IMS Connect 259
 - OTMA, viewing connection information 260
 - restarting a connection from IMS Connect 262
 - starting IMS Connect IMSplex communications 265
 - stopping a connection from IMS Connect 261
 - stopping a send client socket connection 263, 266
 - stopping an MSC logical link in IMS Connect 265
 - stopping IMS Connect IMSplex communications 264
 - viewing IMS Connect connection information 257
- MSC
 - TCP/IP link operations, overview 273
 - TCP/IP link restart after errors 274
- MTO, auditing 250
- personnel 229
- procedures
 - designing 256
 - developing 229, 285
 - documenting 232
 - maintenance 281
 - multisystem 267
 - planning 233
 - record keeping 237
 - testing 281
- remote terminal operator
 - problem reporting 291
- responsibilities, establishing 246
- send messages 296
- setting up operator actions 247, 248
- specifying resources 249
- strategy 236
- TCP/IP connections
 - cleaning up an MSC logical link in IMS Connect 265
 - MSC, viewing connection information 257

operations (continued)

- TCP/IP connections (continued)
 - MSC, viewing link information in IMS Connect 259
- OTMA, viewing connection information 260
- restarting an IMS-to-IMS connection from IMS Connect 262
- starting IMS Connect IMSplex communications 265
- stopping a send client socket connection 263, 266
- stopping an IMS-to-IMS connection from IMS Connect 261
- stopping an MSC logical link in IMS Connect 265
- stopping IMS Connect IMSplex communications 264
- viewing IMS Connect connection information 257
- TCP/IP connections between IMS systems 257
- operations interactions 232
- operations manager (OM)
 - audit trail 243
- Operations Manager (OM)
 - shutting down 125
 - starting 64
- operator
 - control
 - levels 236
- operator actions
 - setting up 248
- operators
 - prerequisite knowledge vii
- options 253
- OTMA
 - stop 114
 - transaction instance block (TIB)
 - setting 36
- output messages
 - interrupting 323
- output terminal
 - high-speed printer 320

P

- password
 - entering through MFS 301
- performance
 - gathering data 109
- PF keys 247
- PF keys for MTO 247, 310
- physical link
 - effect of commands on 15
 - failure 269, 271
- physical link (MSPLINKs)
 - commands for controlling 15
- physical links
 - switching TCP/IP and VTAM types 272
- physical terminal
 - effect of commands on 15
- planned remote takeover
 - overview 180

- planning
 - availability of service 244
- preinitialized DBCTL environment 223
- preparing for a cold start 173
- PRILOG records 165
- printers
 - local 320
 - SLU-1 devices 321
- procedures
 - content 234
 - operations
 - designing 256
 - developing 229, 285
 - documenting 232
 - maintenance 281
 - multisystem 267
 - planning 233
 - record keeping 237
 - testing 281
 - setting up standard JCL 281
- processing
 - command 340
 - system message 340
- processing load
 - adjusting 24
- production cycle
 - defining 244
- program
 - effect of commands on 15
- programmable remote systems
 - communication with IMS 321
- programmed symbols (PS) 303
 - administering 313
 - definition 313
 - designing applications 315
 - determining if loaded 313
 - example of loading 314
 - how to load 314, 315
 - solving load problems 314
- programs
 - updating
 - dynamically 29
- protecting
 - 3270 terminal screen 302
- PS (programmed symbols)
 - administering 313
 - definition 313
 - designing applications 315
 - determining if loaded 313
 - example of loading 314
 - how to load 314, 315
 - solving load problems 314
- PURGE keyword 116

Q

- QCF (Queue Control Facility) 133
- QUERY command
 - HALDB 169
 - MSC 276
- Queue Control Facility (QCF)
 - recovery 133
- queues
 - shared
 - remote takeover 188
 - suspend 96
- QUIESCE keyword 119

R

- RACF (Resource Access Control Facility)
 - data space
 - reinitialization 35
- RCMD call 338
- RCMD calls 366
- RDS (restart data set)
 - recovery 138
- RECON data sets
 - adding or spare 45
 - allocation records 159
 - changing the characteristics of 45
 - DBRC 147
 - defining 46
 - recovering 134
 - removing spare 45
 - replacing 45
 - replacing active 45
 - restore 134
 - spare
 - adding 45
 - unusable 134
- records
 - DB Monitor, data 110
 - PRILOG 165
- recovery
 - 3270 errors 316
 - ACBLIB failures 134
 - backward 140
 - CCTL failures 144
 - considerations 279
 - control region failures 132
 - coordinating with DB2 for z/OS 191
 - CPI Communications driven
 - programs 143
 - CPI Communications failures 142
 - CQS log 144
 - data set failures 133
 - data sharing 146
 - database 140
 - making recoverable 228
 - database failure 140
 - emergency restart failures 132
 - executing related functions 226
 - FDBR (Fast Database Recovery) 148
 - FDBR, using 149
 - forward 140
 - in a DBCTL environment 145
 - introduction 131
 - IRLM 146
 - log errors 136
 - LU 6.2 failures 142
 - message queues 133, 227
 - messages for AO application 370
 - methods, comparing 140
 - multiple systems 279
 - network failures 141
 - OLDS 137
 - RECON 134
 - reestablish FDBR 149
 - region controller failures 139
 - remote terminal commands 288
 - SLDS 138
 - structure failures 144
 - structures, for restart 87
 - user access problems 151
 - utilities 228

- recovery (*continued*)
 - VTAM message
 - resynchronization 143
 - z/OS failures 131
- recovery points
 - create 224, 225
 - database quiesce 223
- region
 - assignment 24
 - CCTL 83
 - class 24
 - DEDB online utility 82
 - failures 139
 - Fast Path message-driven 82
 - Java dependent 82
 - looping 139
- REGION keyword 81
- Remote Site Recovery (RSR) 152
 - databases 158
 - logger 161
 - starting new tracking subsystem 190
 - takeover 190
 - XRC tracking
 - shutdown processing during 193
- remote systems, programmable
 - communication with IMS 321
- remote terminal
 - failure 141
- remote terminal operator
 - conversational transactions 251
- replacing magnetic tape 320
- replacing RECON data sets 45
- report
 - Statistical Analysis utility 109
- Repository Server
 - start 68
 - stop 68, 69
 - subordinate
 - start 69
- resource
 - changing online 13
 - cleanup failure 85
 - commands, remote terminal
 - operator 288
 - modifying 12
 - structure failures 144
- Resource Manager (RM)
 - shutting down 125
 - starting 64
- resource structure
 - cold start 87
- response mode
 - transactions 294
- responses to commands 368
- restart
 - active site 173
 - after shutdown 94
 - AO application 370
 - automatic 70, 94, 100
 - backout of BMPs, deferring 99
 - batch jobs 102
 - BMP regions 139
 - component 93
 - DL/I database tracking 174
 - emergency 77, 97
 - failures 132

- restart (*continued*)
 - emergency restart and dynamic resource definition 79
 - IMS Fast Path regions 139
 - IMSplex
 - RM 99
 - log router 174
 - SLDS as input 79
 - structure recovery 87
 - tracking site 173
 - warm start 96
 - without remote takeover 173
 - z/OS Automatic Restart Manager 84
- restart commands 93
- restart data set (RDS)
 - recovery 138
- restart RM
 - automatic restart manager 65
- restarting BMPs 101
- restarting CQS
 - cold start 84
 - description 83
 - structure initialization 85
 - warm start 83
- restarting IMS 61
- restarting ODBM 64
- restarting SCI
 - ARM 63
- REXX SPOC API
 - overview 343
- RM (Resource Manager)
 - shutting down 122
 - starting 62
- routing codes
 - updating
 - dynamically 30
- RSR (Remote Site Recovery) 152
 - allocating DBDSs and areas 159
 - begin
 - shadowing 153
 - catch-up processing 177
 - database tracking 167
 - database update set 159
 - databases 158
 - discontinue 191
 - DL/I batch 172
 - DL/I batch jobs 163
 - DRD and RSR 175
 - enabling XRC tracking 192
 - gap processing 193
 - HALDB 169
 - HSSP 162
 - IMSplex 157
 - log gaps 176
 - log router 164
 - logger 161
 - new active subsystem, starting 188
 - normal operations
 - active site 160
 - online change 162, 171
 - online forward recovery 178
 - operating 160
 - operations
 - tracking site 164
 - restart without remote takeover 173
 - takeover 180
 - canceling 186

- RSR (Remote Site Recovery) (*continued*)
 - takeover (*continued*)
 - message queues, sharing 188
 - new active subsystem, starting 188
 - performing 182, 185
 - planned 181
 - starting new tracking subsystem 190
 - tracking site 187
 - unplanned 183
 - termination
 - abnormal 175
 - active site 172
 - tracking site 172
 - transport manager 172
 - Transport Manager subsystem 176
 - TMS 160
 - tracing
 - database setup 154
 - tracking
 - preparations 152
 - RSR complex 157
 - sample JCL 155
 - tracking log 158
 - XRC tracking
 - during takeover 194
 - restart considerations 193
 - resume 192
 - routing delays 192
 - shutdown processing during 193
 - stop 192
- RSTART (/RSTART) LINK
 - command 271
- runtime
 - descriptor definition
 - create 26
 - resource definition
 - create 26
- runtime descriptor definitions 25
 - querying dynamically 34
- runtime resource definitions
 - querying dynamically 34
 - update dynamically 28
- runtime resources
 - deleting dynamically 33

S

- SADMP 127
- sample AO application
 - DFSAPGM 371
 - UETRANS 363
- SCI (Structured Call Interface)
 - restarting 63
 - shutting down 122
 - starting 62, 63
- screen format
 - display 300
- secondary logical unit (SLU-1)
 - telephone data set 322
- security
 - AO applications 367
 - options
 - specifying at startup 79

- security (*continued*)
 - RACF (Resource Access Control Facility) 335, 338
 - requirements, fulfilling 292
- service availability
 - planning 244
- session initiation
 - SLU-2 devices 324
- sessions
 - immediate termination 119
 - orderly termination 119
 - recovery 142
 - terminating 113, 119
 - XRF 221
- SET CLOCK command 73
- shutdown
 - accelerating 116
 - control region 114
 - dynamic resource definition (DRD) 118
 - forcing 126
 - immediate 116
 - IMS 113
 - /CHECKPOINT command 116
 - network 120
 - orderly 116
 - restart 94
 - VTAM 113
- shutdown commands 93
- SHUTDOWN CSLPLEX command 124
- shutdown processing
 - XRC tracking 193
- shutting down
 - IMSplex 122
- shutting down CQS 122
- shutting down the CSL 123
- single point of control (SPOC)
 - introduction 1
- single-segment
 - entering 323
- SLDS (system log data set)
 - as input to restart 79
 - changing the characteristics of 45
 - read error 135
 - recovery 138
- SLU-1 communication terminal
 - AUTO switch 323
- SLU-1 devices 322
 - communication terminals 321
 - disconnecting 324
 - printers 321
- SLU-2 devices
 - connect to IMS 324
- SNAPQ keyword 202
- space allocation
 - changing
 - OLDS 43
- SPOC (single point of control)
 - introduction 1
 - setting up 2
 - starting 2
- stand-alone dump 127
- standby
 - DBCTL 223
- start
 - BMP regions 82
 - CCTL 83

- start (*continued*)
 - cold 73, 74, 95
 - IMSRSC repository 76
 - RDDS 75
 - cold start and dynamic resource definition 95
 - control region 70
 - CQS subsystem 81
 - DBRC address space 70
 - DEDB utility regions 82
 - dependent regions 81
 - FDBR 151
 - IMS 91
 - IMS network, prerequisites 88
 - IMSRSC repository 67
 - IRLM subsystem 80
 - Java dependent regions 82
 - MPP regions 82
 - MSC 267
 - Repository Server 68
 - subordinate 69
 - Transaction Manager 88
 - warm 77, 96
 - XRF 199, 200
- START (/START) command 169
 - APPC 90
 - DC 88
 - HALDB 169
 - REGION 81
- START command
 - CQS 81
 - IMS 62, 70
 - IRLM 80
- starting IMS 61
- startup
 - security options, specifying 79
 - XRF 200
- status
 - IMSRSC repository 65
- status information
 - tracking 167
- status line 308
- stop
 - APPC/IMS 114, 120
 - control region 114
 - dependent regions 114
 - IMSRSC repository 67
 - IRLM subsystem 121
 - MSC 268
 - OTMA 114
 - Repository Server 68, 69
 - Transaction Manager 113
- STOP (/STOP) command
 - DC 113
- Stop Job key
 - disconnecting SLU-1 devices 324
- structure
 - cold start 86
 - empty 85
 - initialization 85
 - recovery for restart 87
 - restarting CQS 85
 - warm start 85
- Structured Call Interface (SCI)
 - shutting down 126
 - starting 63

- subsystems
 - connecting or disconnecting 46
 - effect of commands on 15
 - identifying 201
 - unauthorize commands 155
- support groups
 - operations 232
 - systems operations 232
 - user liaison 232
- surveillance
 - change 212
 - start 212
 - stop 212
 - XRF 212
- syntax diagram
 - how to read viii
- Sys Req key
 - disconnecting SLU-1 devices 324
- system
 - console
 - z/OS 318
 - failure 142
- system data set
 - reallocating during restart 77
- system definitions
 - XRF 203
- system failure 142
 - restarting BMP 101
- system log data set (SLDS)
 - as input to restart 79
 - changing the characteristics of 45
 - read error 135
 - recovery 138
- system programmers
 - prerequisite knowledge vii

T

- takeover
 - active subsystem
 - easing workload 220
 - cancel 186
 - discontinue RSR 191
 - performing 182
 - planned 181, 214
 - remote 194
 - RSR 180
 - sharing message queues 188
 - starting new tracking subsystem 190
 - unplanned 183
 - unplanned, performing 185
 - XRF 214
- tasks
 - commands for 47
- TCO (time-controlled operations)
 - CNT Edit exit routine 348
 - components 345
 - description 344
 - exit routine 347
 - installing 345
 - large networks, starting 346
 - message sets 351
 - request 351
 - multiple time zones 346
 - peak loads 346
 - scheduling low-priority jobs 346
 - scheduling non-IMS jobs 346

- TCO (time-controlled operations) (*continued*)
 - script loading 352
 - shutting down IMS 346
 - starting IMS resources 346
 - system, monitoring 346
 - time schedule request
 - optional fields 349
 - required fields 349
 - time schedule requests 349
 - user status, updating 346
 - uses 344
 - using 346
 - verification utility 352
- TCO Verification utility (DFSTVER0) 352
- TCP/IP
 - connections
 - checking status, overview 103
 - monitoring 103
 - status by client 104
 - status by port 104
 - status for IMSplexes 105
 - status for remote IMS Connect instances 105
 - status for SCI 105
 - IMS-to-IMS TCP/IP connections
 - cleaning up an MSC logical link in IMS Connect 265
 - MSC, viewing connection information 257
 - MSC, viewing link information in IMS Connect 259
 - operations 257
 - OTMA, viewing connection information 260
 - restarting a connection from IMS Connect 262
 - starting IMS Connect IMSplex communications 265
 - stopping a connection from IMS Connect 261
 - stopping a send client socket connection 263, 266
 - stopping an MSC logical link in IMS Connect 265
 - stopping IMS Connect IMSplex communications 264
 - viewing IMS Connect connection information 257
 - MSC
 - link operations, overview 273
 - link restart after errors 274
 - telecommunication lines
 - effect of commands on 15
 - modifying 24
 - telephone data set
 - secondary logical unit (SLU-1) 322
 - teletypewriter 317, 318
 - terminal display size 300
 - terminal keyboard
 - locking and unlocking 303
 - terminal operator logical paging 301
 - terminal, startup 292
 - terminals
 - administering 24
 - assigning 24

- terminals (*continued*)
 - location
 - master 196
 - modifying 24
 - operating, end-user procedures 285
 - restore operation 141
 - session management
 - XRF 201
 - supervise 332
 - support
 - XRF 218
 - XRF 218
- terminating sessions 119
- termination
 - Fast Path database tracking 175
 - IMS
 - break key 318
 - EOT key 318
 - RSR 175
 - Transport Manager subsystem (TMS) 176
- testing procedures 281
- thread looping
 - CCTL 145
- TIB
 - OTMA 36
- time
 - setting 73
 - conditions 71
- time stamps 71
- time-controlled operations (TCO)
 - script loading 352
- TIMEZONE keyword 73
- Tivoli NetView for z/OS
 - automated operations 330
- TMS (Transport Manager subsystem)
 - commands 160
 - RSR 160
- TOD (time-of-day) clock 71
- TRACE (/TRACE) command 103
 - IMS Monitor 111
 - options 253
 - VTAM I/O Timeout facility 24
- TRACE CT command 112
- trace options 253
- tracing
 - CTRACE, using 112
- tracking
 - data sharing 163
 - database, RSR 167
 - DBRC 158
 - FDBR 150
 - log router 164
 - routing delays 192
 - RSR
 - preparation 152
 - start 150
 - subsystem 175
 - subsystem log 158
 - XRC
 - enable 192
 - gap processing 193
 - restart considerations 193
 - resume 192
 - shutdown processing during 193
 - stop 192
 - unplanned remote takeover 194

- tracking (*continued*)
 - XRF 205
- tracking site
 - image copy data set 157
 - restart 173
 - RSR 164
 - RSR, termination 172
- tracking subsystem
 - gaps 177
 - milestone processing 167
 - starting 190
- trademarks 373
- transaction class
 - effect of commands on 15
- transaction code 293
 - format of 293
- transaction instance block (TIB)
 - setting 36
- Transaction Manager
 - APPC/MVS, connecting to 90
 - CICS 90
 - ISC, connecting to 90
 - starting 88
 - stop 113
 - VTAM, connecting to 88
- transactions
 - assignment 25
 - controlling 25
 - conversational 280, 295
 - create dynamically 27
 - effect of commands on 15
 - entering 293
 - modifying 25
 - priorities 25
 - response mode 294
 - stopped 280
- Transport Manager subsystem (TMS)
 - commands 160
 - RSR 160
 - termination 176
- Transport Manager Subsystem (TMS)
 - component failure 173
 - discontinue 191
 - IMS logger 161
 - terminating 172
- TSO SPOC 8, 9
 - command responses 5
 - command short cuts 7
 - command shortcuts 6
 - command status 6
 - defining groups 7
 - interface 1
 - issuing
 - commands 1, 8
 - type-1 commands 1
 - type-2 commands 1
 - overview 1
 - screen example 1
 - setting up 2
 - starting 2, 5
- TTY device 327
- type-1 AOI
 - GCMD calls 335
- type-2 AOI 338
- type-2 commands 47
- types of AOI, comparison 343

U

- unlocking
 - terminal keyboard 303
- unplanned remote takeover
 - overview 180
- UPDATE command
 - HALDB 169
- update process
 - IMSRSC repository 36
- updating resources 25
- user
 - access problems 151
 - ISC, LTERM assignment 35
- user-input area 310
- users
 - effect of commands on 15
- USERVAR
 - initializing 202
 - Management Enhancement 202
- utilities
 - DFSISTS0 (Statistical Analysis utility) 109
 - execution, recording 238
 - Log Transaction Analysis utility (DFSILTA0) 109
 - recovery 228
 - Statistical Analysis utility (DFSISTS0) 109

V

- VTAM (Virtual Telecommunications Access Method)
 - connecting to 88
 - message resynchronization 143
 - MNPS 96, 97
 - shutdown 113
 - terminating a session 113, 119
 - VTAM I/O Timeout facility 24
 - XRF 221

W

- WADS (write-ahead data set)
 - adding or removing spare 43
 - changing
 - allocation 44
 - location 44
 - space 44
 - characteristics
 - changing 43
 - error 135
 - location
 - changing 43
 - mode
 - changing 43
 - modes
 - changing 43, 44
 - recovery 138
 - spare 44
- warm start 83
 - introduction 96
 - performing 77
 - processing 96
- warm starting structures 85
- warning message area 309

- write-ahead data set (WADS)
 - adding or removing spare 43
 - characteristics
 - changing 43
 - error 135
 - location
 - changing 43
 - mode
 - changing 43
 - recovery 138
 - removing spare 44
- write-to-operator (WTO)
 - messages 318
- write-to-operator-with-reply (WTOR)
 - messages 318
- WTO
 - messages 318

X

- XRC tracking
 - DB2 for z/OS 191
 - enable 192
 - gap processing 193
 - restart considerations 193
 - resuming 192
 - routing delays 192
 - shutdown processing 193
 - stop 192
 - takeover, planned 194
 - unplanned remote takeover 194
- XRF (Extended Recovery Facility)
 - availability manager 200
 - backout 220, 221
 - change surveillance 212
 - checkpoint 202
 - class-1 terminals 205
 - commands 206, 207
 - CPC 201
 - data propagation 199
 - data set 205
 - data sets
 - initializing 199
 - data-sharing environment 199
 - DBCTL 222
 - DBCTL service 223
 - determining parameters 199
 - diagnostics 222
 - end users 218
 - I/O toleration
 - end 216
 - IMS identifier 201
 - IMSID, ALTID 201
 - maintenance 211
 - monitoring 207
 - network status 208
 - new alternative, establishing 222
 - online change 198
 - operations 198
 - system status, explanation 207
 - overview 195
 - planned
 - rules 214
 - planning 196
 - post-takeover 220
 - session management
 - terminals 201

- XRF (Extended Recovery Facility)
 - (*continued*)
 - start surveillance 212
 - starting 199
 - startup procedures 200
 - stop surveillance 212
 - subsystems
 - starting 200
 - surveillance 212
 - synchronization 202
 - system data sets
 - updating 213
 - system definitions 203
 - takeover 214
 - conditions 214
 - failure 218
 - flow 215
 - I/O prevention 215
 - I/O toleration, ending 216
 - manual control 217
 - post-takeover 220
 - terminal support 218
 - TSO 219
 - unplanned 214
 - terminal sessions, managing 202
 - terminals 218, 221
 - terminating 222
 - tracking 205, 213
 - control region, dumping 213
 - suspending operations 213
 - USERVAR 202, 219
 - VTAM 221

Z

- Z NET command 113
- z/OS
 - Automatic Restart Manger 84
 - command syntax 124
 - failures 131



Product Number: 5635-A03
5655-DSQ

Printed in USA

SC19-3018-02



Spine information:

IMS Version 12

Operations and Automation

