

IMS



# HALDB Online Reorganization Guide and Reference

*Version 9*



IMS



# HALDB Online Reorganization Guide and Reference

*Version 9*

**Note**

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 83.

**Quality Partnership Program (QPP) Edition (April 2004) (Softcopy Only)**

This QPP edition replaces or makes obsolete the previous edition, ZES1-2346-00. This edition is available in softcopy format only. The technical changes for this version are summarized under "Summary of Changes" on page xvii.

**© Copyright International Business Machines Corporation 2004. All rights reserved.**

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

# Contents

|   |      |
|---|------|
| <b>Figures</b>  | vii  |
| <b>Tables</b>   | ix   |
| <b>About This Book</b>  | xi   |
| Prerequisite Knowledge  | xi   |
| Terminology of HALDB OLR                                      | xi   |
| How to Read Syntax Diagrams                                   | xii  |
| Syntax Diagram Example  | xiv  |
| How to Send Your Comments                                     | xiv  |
| <b>Summary of Changes</b>                                     | xvii |
| Changes to the Current Edition of This Book for IMS Version 9 | xvii |
| Library Changes for IMS Version 9                             | xvii |
| New and Revised Titles  | xvii |
| Terminology Changes   | xvii |
| Accessibility Enhancements                                    | xvii |
| <b>Chapter 1. Overview of HALDB OLR</b>                       | 1    |
| How HALDB OLR Functions                                       | 1    |
| When HALDB OLR Begins   | 1    |
| When HALDB OLR Initialization Fails                           | 2    |
| After HALDB OLR Initialization Completes                      | 2    |
| After the HALDB OLR Copying Process Completes                 | 4    |
| Naming Conventions for HALDB OLR Data Sets                    | 4    |
| <b>Chapter 2. Planning for HALDB OLR</b>                      | 7    |
| Migration   | 7    |
| Coexistence   | 7    |
| Fallback  | 8    |
| Restrictions  | 8    |
| Performance Considerations for HALDB OLR                      | 9    |
| Global Locking  | 9    |
| Restart, XRF, and RSR with HALDB OLR                          | 10   |
| Logging Effects of Reorganization                             | 11   |
| Reorganization Processing Time                                | 11   |
| Activating Sequential Buffering for OLR                       | 12   |
| Recovery  | 12   |
| Database and Data Set Recovery                                | 12   |
| Remote Site Recovery  | 14   |
| DBRC Considerations   | 15   |
| GENJCL.IC and GENJCL.OIC Considerations                       | 15   |
| NOTIFY.IC and NOTIFY.UIC Considerations                       | 15   |
| DBRC RECON HALDB OLR Information                              | 16   |
| RSR Tracking System   | 17   |
| Recovery Related Records                                      | 18   |
| Online Reorganization Execution                               | 18   |
| Offline Reorganization  | 19   |
| Utilities   | 19   |
| HALDB Index / ILDS Rebuild - DFSPREC0                         | 19   |
| Database Change Accumulation - DFSUCUM0                       | 19   |
| Database Recovery - DFSURDB0                                  | 20   |
| Database Image Copy - DFSUDMP0                                | 20   |

|   |           |
|---|-----------|
| Batch Backout - DFSBBO00 . . . . .  | 20        |
| <b>Chapter 3. The Data Sets of HALDB OLR . . . . .</b>  | <b>21</b> |
| Output Data Set Creation . . . . .  | 21        |
| Required Characteristics for Existing Output Data Sets . . . . .  | 21        |
| Characteristics for Automatically-Created Output Data Sets . . . . .  | 22        |
| Number of Volumes: . . . . .  | 22        |
| Location, SMS-Managed Case: . . . . .   | 22        |
| Location, Non-SMS-Managed, Non-VSAM Data Set Case: . . . . .  | 22        |
| Location, Non-SMS-Managed, VSAM Data Set Case: . . . . .  | 22        |
| Size, Single-Volume Case: . . . . .   | 22        |
| Size, Multiple-Volume Case (SMS-Managed Only): . . . . .  | 23        |
| The Disposition of the Output Data Sets When Initialization Fails . . . . .                                   | 24        |
| The Inactive Data Sets at Successful Completion of the Reorganization<br>Process- Data Set Deletion . . . . . | 24        |
| <b>Chapter 4. Setting Up HALDB OLR . . . . .</b>  | <b>25</b> |
| INITIATE OLREORG Command . . . . .  | 25        |
| /INITIATE OLREORG Command . . . . .   | 25        |
| Format of the INITIATE OLREORG Command . . . . .  | 25        |
| Environments, Keywords, and Parameters of the INITIATE OLREORG<br>Command . . . . .                           | 26        |
| Usage of the INITIATE OLREORG Command . . . . .   | 26        |
| Examples of INITIATE OLREORG Input and Output . . . . .   | 30        |
| INIT.DB. . . . .  | 31        |
| Format of INIT.DB. . . . .  | 31        |
| INIT.DB Command Usage . . . . .   | 32        |
| NOTIFY.REORG . . . . .  | 32        |
| Format of NOTIFY.REORG . . . . .  | 32        |
| NOTIFY.REORG Command Usage . . . . .  | 33        |
| <b>Chapter 5. Monitoring HALDB OLR . . . . .</b>  | <b>35</b> |
| QUERY OLREORG Command . . . . .   | 35        |
| Format of the QUERY OLREORG Command. . . . .  | 35        |
| Environments, Keywords, and Parameters of the QUERY OLREORG<br>Command . . . . .                              | 35        |
| Usage of the QUERY OLREORG Command . . . . .  | 36        |
| Example of QUERY OLREORG Input and Output . . . . .   | 38        |
| /DISPLAY DATABASE Command HALDB OLR Support . . . . .   | 40        |
| Example of /DISPLAY DB Output with HALDB OLR . . . . .  | 40        |
| /DISPLAY DB OLR Command . . . . .   | 40        |
| QUERY DATABASE STATUS(OLR) Command . . . . .  | 42        |
| <b>Chapter 6. Modifying and Tuning HALDB OLR . . . . .</b>  | <b>43</b> |
| TERMINATE OLREORG COMMAND. . . . .  | 43        |
| /TERMINATE OLREORG Command . . . . .  | 45        |
| Format of the TERMINATE OLREORG Command . . . . .   | 45        |
| Environments, Keywords, and Parameters of the TERMINATE OLREORG<br>Command . . . . .                          | 45        |
| Usage of the TERMINATE OLREORG Command . . . . .  | 46        |
| Examples of TERMINATE OLREORG Input and Output . . . . .  | 48        |
| UPDATE OLREORG Command . . . . .  | 49        |
| /UPDATE OLREORG Command . . . . .   | 50        |
| Format of the UPDATE OLREORG Command . . . . .  | 51        |
| Environments, Keywords, and Parameters of the UPDATE OLREORG<br>Command . . . . .                             | 51        |

|   |  |    |
|---|--|----|
|   | Usage of the UPDATE OLREORG Command . . . . .                              | 51 |
|   | Examples of UPDATE OLREORG Input and Output . . . . .                      | 54 |
|   | /START DB Command HALDB OLR Support . . . . .                              | 56 |
|   | /DBRecovery Command HALDB OLR Support . . . . .                            | 56 |
|   | /RECover Command HALDB OLR Support . . . . .                               | 56 |
|   | CHANGE.DB . . . . .  | 56 |
|   | Format of CHANGE.DB. . . . .   | 56 |
|   | CHANGE.DB Command Usage. . . . .   | 57 |
|   | <b>Chapter 7. Messages and Codes of HALDB OLR</b> . . . . .                | 59 |
|   | HALDB OLR DFS Series Messages . . . . .                                    | 59 |
|   | HALDB OLR DSP Series Messages . . . . .                                    | 74 |
|   | HALDB OLR Abends and Codes . . . . .                                       | 76 |
|   | <b>Chapter 8. Log Records of HALDB OLR.</b> . . . . .                      | 77 |
|   | DFSLOG - Database Change Log Records . . . . .                             | 77 |
|   | DFSLOG20 - Database Open and Close (X'20' and X'21') Log Records . . . . . | 77 |
| I | DFSLOG27 - Database Extension (X'27) Log Records . . . . .                 | 77 |
| I | DFSLOG53 - CI/CA and Space Manager (X'53) Log Records . . . . .            | 77 |
|   | Log Records: X'29' OLR Event Log Records . . . . .                         | 77 |
|   | X'2910': Ownership Established. . . . .                                    | 78 |
|   | X'2920': UPDATE OLREORG Command . . . . .                                  | 78 |
|   | X'2930': Output Data Set Information. . . . .                              | 78 |
|   | X'2940': Cursor-Active Status . . . . .                                    | 78 |
|   | X'2950': Cursor Movement . . . . .   | 78 |
|   | X'2970': Cursor-Active Status Reset . . . . .                              | 78 |
|   | X'2990': Ownership Relinquished . . . . .                                  | 78 |
|   | <b>Chapter 9. RECON Listing HALDB OLR Support</b> . . . . .                | 79 |
|   | HALDB MASTER RECORD . . . . .  | 79 |
|   | PARTITION DB RECORD . . . . .  | 79 |
| I | DBDS Record . . . . .  | 80 |
|   | REORG RECORD . . . . .   | 82 |
| I | <b>Notices</b> . . . . .   | 83 |
| I | Trademarks . . . . .   | 85 |
| I | <b>Bibliography</b> . . . . .  | 87 |
| I | IMS Version 9 Library . . . . .  | 87 |
|   | <b>Index</b> . . . . .   | 89 |





---

## Figures

|      |  |    |
|------|--|----|
| 1.   | Relationship between Input and Output Data Sets at a Point During Reorganization . . . . . | 2  |
| 2.   | Normal Processing Flow of a HALDB OLR . . . . .  | 3  |
| 3.   | /DISPLAY DB PVHDJKC POIHDKD output for databases (2 with OLRs in progress) . . . . .       | 40 |
| 4.   | /DISPLAY DB OLR output . . . . .   | 40 |
| 5.   | HALDB OLR Process with a TERMINATE OLREORG . . . . .                                       | 44 |
| I 6. | HALDB OLR Process with an UPDATE OLREORG . . . . .   | 50 |



## Tables

|     |  |     |
|-----|--|-----|
| 1.  | How to Read Syntax Diagrams . . . . .  | xii |
| 2.  | HALDB Data Set Name Examples . . . . .   | 5   |
| 3.  | DSID DCB Numbers for the A-J and X and the M-V and Y data sets . . . . .                 | 16  |
| 4.  | Data Sharing Between the Releases . . . . .  | 17  |
| 5.  | Valid Environments for the INITIATE OLREORG command Keywords, and Parameters . . . . .   | 26  |
| 6.  | Output Fields for INITIATE OLREORG Command . . . . .                                     | 29  |
| 7.  | Return and Reason Code for the INITIATE OLREORG Command . . . . .                        | 29  |
| 8.  | Completion Codes for the INITIATE OLREORG Command . . . . .                              | 29  |
| 9.  | Valid Environments for the QUERY OLREORG command Keywords, and Parameters . . . . .      | 36  |
| 10. | Output Fields of QUERY OLREORG . . . . .   | 37  |
| 11. | Return and Reason Code for the QUERY OLREORG Command . . . . .                           | 38  |
| 12. | Completion Codes for the QUERY OLREORG Command . . . . .                                 | 38  |
| 13. | Valid Environments for the TERMINATE OLREORG Command, Keywords, and Parameters . . . . . | 45  |
| 14. | Output Fields for TERMINATE OLREORG Command . . . . .                                    | 47  |
| 15. | Return and Reason Code for the TERMINATE OLREORG Command . . . . .                       | 47  |
| 16. | Completion Codes for the TERMINATE OLREORG Command . . . . .                             | 48  |
| 17. | Valid Environments for the UPDATE OLREORG Command Keywords and Parameters . . . . .      | 51  |
| 18. | Output Fields for UPDATE OLREORG Command . . . . .                                       | 53  |
| 19. | Return and Reason Code for the UPDATE OLREORG Command . . . . .                          | 54  |
| 20. | Completion Codes for the UPDATE OLREORG Command . . . . .                                | 54  |
| 21. | Reason Codes, Environments, and Errors Detected for Message DFS2991I . . . . .           | 69  |
| 22. | X'29' OLR Event Log Record Subtypes . . . . .  | 78  |



---

## About This Book

This book is available in PDF and BookManager® formats, and also as part of the DB2 Information Management Software Information Center for z/OS Solutions. To get the most current versions of the PDF and BookManager formats, go to the IMS Library page at [www.ibm.com/software/data/ims/library.html](http://www.ibm.com/software/data/ims/library.html). To view the information within the DB2 Information Management Software Information Center for z/OS Solutions, go to [publib.boulder.ibm.com/infocenter/db2zhhelp](http://publib.boulder.ibm.com/infocenter/db2zhhelp).

This information is intended to supply you with IMS HALDB Online Reorganization information.

Licensed information on HALDB Online Reorganization has been excluded from this information and can be found in the *IMS Version 9: Diagnosis Guide and Reference*.

---

## Prerequisite Knowledge

The following prerequisites are needed for you to use HALDB OLR support:

- A working knowledge of database management, HALDBs, and DBRC.
- A working knowledge of the IMS Common Service Layer and an understanding of Operations Manager (OM) and Structured Call Interface (SCI).

### Related Reading:

- For more information on database management and HALDBs, see the *IMS Version 9: Administration Guide: Database Manager*.
- For more information on DBRC, see the *IMS Version 9: DBRC Guide and Reference*.
- For more information on the OM and SCI, see the *IMS Version 9: Common Service Layer Guide and Reference*.

---

## Terminology of HALDB OLR

### **active set of data sets**

In cursor-active status, the paired input and output sets of data sets that the data is being copied from and to respectively. Otherwise, the single set of data sets that comprise the partition.

### **cursor**

The point within a partition that delineates the end of the database records that have been copied.

### **cursor-active status**

Paired sets of data sets are active. OLREORG ACTIVE=YES is displayed in the RECON listing.

### **HALDB OLR**

Non-disruptive reorganization of HALDB PHDAM and PHIDAM partitions. Synonymous with online reorganization and HALDB Online Reorganization, OLR, OLREORG, and HALDB OLREORG.

### **HALDB OLREORG**

See HALDB OLR.

### **HALDB Online Reorganization**

See HALDB OLR.

**inactive set of data sets**

The old input set of data sets that the data was copied from.

**OLR** See HALDB OLR.

**OLR ACTIVE HARD COUNT**

Represents the number of HALDB OLR cursor active calls that are pending in the tracking RECON.

**OLR INACTIVE HARD COUNT**

Represents the number of HALDB OLR cursor inactive calls that are pending in the tracking RECON.

**OLREORG**

See HALDB OLR.

**online reorganization**

See HALDB OLR.

**owned**

Status during which an IMS subsystem has exclusive control over a HALDB OLR.

**SHAREPLEX**

As used in this document, SHAREPLEX is defined as the IMS Data Sharing within the boundaries of one set of DBRC RECONS. As such, this differs from MVS™ SYSPLEX and also from the term, IMSplex or OM API (Operations Manager Application Programming Interface) which is sometimes used in an IMS Shared Queues environment.

**unit of reorganization**

One or more HALDB records (root segment and all of its dependent segments) that are reorganized within one commit boundary.

---

## How to Read Syntax Diagrams

Each syntax diagram in this book begins with a double right arrow and ends with a right and left arrow pair. Lines that begin with a single right arrow are continuation lines. You read a syntax diagram from left to right and from top to bottom, following the direction of the arrows.

Table 1 describes the conventions that are used in syntax diagrams in this information:

*Table 1. How to Read Syntax Diagrams*









| Convention  | Meaning  |
|---|--|
|  | You must specify values A, B, and C. Required values are shown on the main path of a syntax diagram.       |
|  | You must specify value A, B, or C.   |
|  | You have the option to specify value A. Optional values are shown below the main path of a syntax diagram. |

Table 1. How to Read Syntax Diagrams (continued)


| Convention  | Meaning  |
|---|--|
|    | <p>You have the option to specify A, B, C, or none of these values.</p>  |
|    | <p>You have the option to specify A, B, C, or none of these values. If you don't specify a value, A is the default.</p>  |
|    | <p>You have the option to specify one, more than one, or none of the values A, B, or C. Any required separator for multiple or repeated values (in this example, the comma) is shown on the arrow.</p>                         |
|   | <p>You have the option to specify value A multiple times. The separator in this example is optional.</p>   |
| <p>►►   Name   ◄◄</p> <p><b>Name:</b></p>  | <p>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.</p> |
| <p>Punctuation marks and numbers</p>  | <p>Enter punctuation marks (slashes, commas, periods, parentheses, quotation marks, equal signs) and numbers exactly as shown.</p>   |
| <p>Uppercase values</p>   | <p>Keywords, their allowable synonyms, and reserved parameters appear in uppercase letters for z/OS. Enter these values exactly as shown.</p>  |
| <p>Lowercase values</p>   | <p>Keywords, their allowable synonyms, and reserved parameters appear in lowercase letters for UNIX. Enter these values exactly as shown.</p>  |
| <p>Lowercase values in italic (for example, <i>name</i>)</p>  | <p>Supply your own text or value in place of the <i>name</i> variable.</p>   |
| <p>b</p>  | <p>A b symbol indicates one blank position.</p>  |

Other syntax conventions include the following:

- When you enter commands, separate parameters and keywords by at least one blank if there is no intervening punctuation.
- Footnotes are shown by a number in parentheses, for example, (1).
- Parameters with number values end with the symbol #.
- Parameters that are names end with 'name'.
- Parameters that can be generic end with the symbol \*.

## Syntax Diagram Example

Here is an example syntax diagram that describes the `hello` command.



### Name:



### Greeting:



### Notes:

- 1 You can code up to three names.
- 2 Compose and add your own greeting (for example, how are you?).

According to the syntax diagram, these commands are all valid versions of the `hello` command:

```
hello
hello name
hello name, name
hello name, name, name
hello, your_greeting
hello name, your_greeting
hello name, name, your_greeting
hello name, name, name, your_greeting
```

The space before the `name` value is significant. If you do not code `name`, you must still code the comma before `your_greeting`.

---

## 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 do one of the following:

- Go to the IMS Library page at [www.ibm.com/software/data/ims/library.html](http://www.ibm.com/software/data/ims/library.html) and click the Library Feedback link, where you can enter and submit comments.
- Send your comments by e-mail to [imspubs@us.ibm.com](mailto: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 you are commenting on (for example, a page number in the PDF or a heading in the Information Center).



---

## Summary of Changes

---

### Changes to the Current Edition of This Book for IMS Version 9

This edition is a draft version of this book intended for use during the Quality Partnership Program (QPP). Contents of this book are preliminary and under development.

There have been extensive structural and content changes to this book since the start of the QPP program.

This softcopy book is available only in PDF and BookManager formats.

---

### Library Changes for IMS Version 9

Changes to the IMS Library for IMS Version 9 include the addition of new titles, the change of one title, and a major terminology change. Changes are indicated by a vertical bar (|) to the left of the changed text.

### New and Revised Titles

The following list details the major changes to the IMS Version 9 library:

- *IMS Version 9: HALDB Online Reorganization Guide and Reference*  
The library includes new information: *IMS Version 9: HALDB Online Reorganization Guide and Reference*. This information is available only in PDF and BookManager formats.
- *IMS Version 9: An Introduction to IMS*  
The library includes new information: *IMS Version 9: An Introduction to IMS*.
- The information formerly titled *IMS Version 8: IMS Java User's Guide* is now titled *IMS Version 9: IMS Java Guide and Reference*.

### Terminology Changes

IMS Version 9 introduces new terminology for IMS commands:

#### **type-1 command**

A command, generally preceded by a leading slash character, that can be entered from any valid IMS command source. In IMS Version 8, these commands were called *classic* commands.

#### **type-2 command**

A command that is entered only through the OM API. Type-2 commands are more flexible and can have a broader scope than type-1 commands. In IMS Version 8, these commands were called *IMSpIex* commands or *enhanced* commands.

### Accessibility Enhancements

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products. The major accessibility features in z/OS products, including IMS, enable users to:

- Use assistive technologies such as screen readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

## **User Assistive Technologies**

Assistive technology products, such as screen readers, function with the IMS user interfaces. Consult the documentation of the assistive technology products for specific information when you use assistive technology to access these interfaces.

## **Accessible Information**

Online information for IMS Version 9 is available in BookManager format, which is an accessible format. All BookManager functions can be accessed by using a keyboard or keyboard shortcut keys. BookManager also allows you to use screen readers and other assistive technologies. The BookManager READ/MVS product is included with the z/OS base product, and the BookManager Softcopy Reader (for workstations) is available on the IMS Licensed Product Kit (CD), which you can download from the Web at [www.ibm.com](http://www.ibm.com).

## **Keyboard Navigation of the User Interface**

Users can access IMS user interfaces using TSO/E or ISPF. Refer to the *z/OS V1R1.0 TSO/E Primer*, the *z/OS V1R1.0 TSO/E User's Guide*, and the *z/OS V1R1.0 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.

---

## Chapter 1. Overview of HALDB OLR

High Availability Large Database Online Reorganization (HALDB OLR) provides non-disruptive reorganization of HALDB Partitioned Hierarchical Direct Access Method (PHDAM) and Partitioned Hierarchical Indexed Direct Access Method (PHIDAM) partitions, thus HALDB OLR provides continuous availability of the largest supported databases in the world.

Data processing today requires 24 hour a day, 7 day a week data availability. IMS already leads the industry in capabilities regarding data availability. HALDB OLR goes further and addresses the issue of planned data outage time, that constituted the largest amount of time that data was unavailable. HALDB OLR might also significantly reduce the cost of owning data (for example: reduction in skills and personnel needed to maintain data, and a reduction in storage costs).

This section covers these topics:

- “How HALDB OLR Functions”
- “Naming Conventions for HALDB OLR Data Sets” on page 4

---

### How HALDB OLR Functions

Prior to HALDB OLR support, PHDAM and PHIDAM databases had to be taken offline to be reorganized. The online reorganization of a HALDB PHDAM or PHIDAM partition runs non-disruptively by allowing concurrent IMS updates, including updates by data sharing IMS systems in the DLISAS (Data Language/ separate address space). Concurrent IMS updates are allowed while small amounts of data are moved and reorganized.

This reorganization process does not support data definitional changes. This process only provides reclustered and space distribution advantages.

### When HALDB OLR Begins

Ownership of the online reorganization by an IMS subsystem is established during initialization and is recorded by DBRC in the RECON. When ownership of an online reorganization is established, no other IMS subsystems are allowed to obtain ownership of that online reorganization. Ownership may be released by this IMS prior to the end of initialization or before the completion of the online reorganization; then another IMS may obtain ownership of and complete the online reorganization.

Before online reorganization starts, there is a single *active* set of data sets for the HALDB partition. The data sets in the other (*inactive*) set contain no useful information, and one or more of these data sets need not even exist before the reorganization is started.

In preparation for their use as the output set of data sets, each of the inactive data sets is evaluated. If one of these data sets does not exist, the initialization process creates it automatically. Each preexisting data set is validated to ensure that it meets the requirements that are discussed in “Required Characteristics for Existing Output Data Sets” on page 21.

In the final initialization step, the active set of data sets is treated as the input set, and the inactive set becomes the output set. At the end of this initialization process the online reorganization of the HALDB partition is recorded in the RECON with an *OLREORG cursor-active* status: (OLREORG CURSOR ACTIVE=YES).

## When HALDB OLR Initialization Fails

During initialization, various error conditions, such as an unacceptable preexisting data set or an insufficient amount of DASD space for an automatically-created data set, can cause the initialization to fail.

If any error occurs during or after the data set creation and validation process but before the *OLREORG cursor-active* status of the reorganization has been recorded in the RECON, then any output data sets that were already automatically-created are retained along with any preexisting ones.

## After HALDB OLR Initialization Completes

When the *OLREORG cursor-active* status has been recorded, and until this reorganization completes or until a batch reorganization reload is done, the HALDB partition is comprised of the pair of A-through-J (and X for PHIDAM format only) set of data sets and the M-through-V (and Y for PHIDAM format only) set of data sets. During this time both sets of data sets must be available for the partition to be used.

Database records are copied from the input to the output data sets in units of reorganization. During the reorganization, other database changes might be made to records that have not yet been copied from the input data sets and to records that have already been copied to the output data sets.

Figure 1 illustrates the conceptual relationship between the input and output data sets at a point during the reorganization.

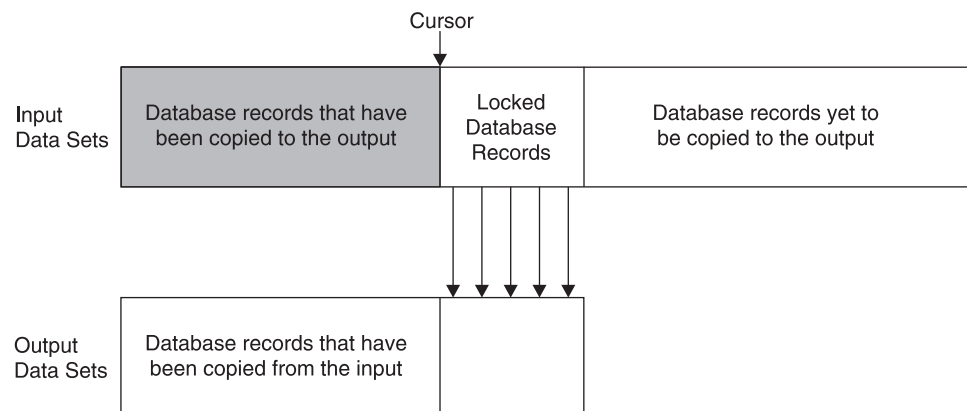


Figure 1. Relationship between Input and Output Data Sets at a Point During Reorganization

Consider the database records for the entire HALDB partition to be ordered from left to right in three categories:

- those that have been copied
- those that are now being copied
- those that are yet to be copied

A cursor marks the point in the partition where the database records have already been copied. As the copying proceeds, this cursor moves forward through the database (from left to right in the diagram).

In the area in the diagram labeled "database records yet to be copied" the database records can be read, replaced, inserted, or deleted. In this case, these changed

records are eventually copied to the output data sets. The records that are currently being copied are locked and cannot be accessed or updated by other applications until they have been copied.

When another application running at the same time as HALDB OLR, attempts to access a database record that is being reorganized online, the following processing occurs internally. If the database record is located before the HALDB OLR cursor, then IMS determines that it is in the input data sets. If the database record is located after the HALDB OLR cursor, then IMS determines that it is in the output data sets.

For possible recovery of either an input or an output data set, all database changes are logged, including the database records copied from the input to the output data sets.

Figure 2 shows the normal processing steps of a successful HALDB OLR with columns depicting flow of control from the User to IMS and the statuses of the sets of data sets as processing events occur.

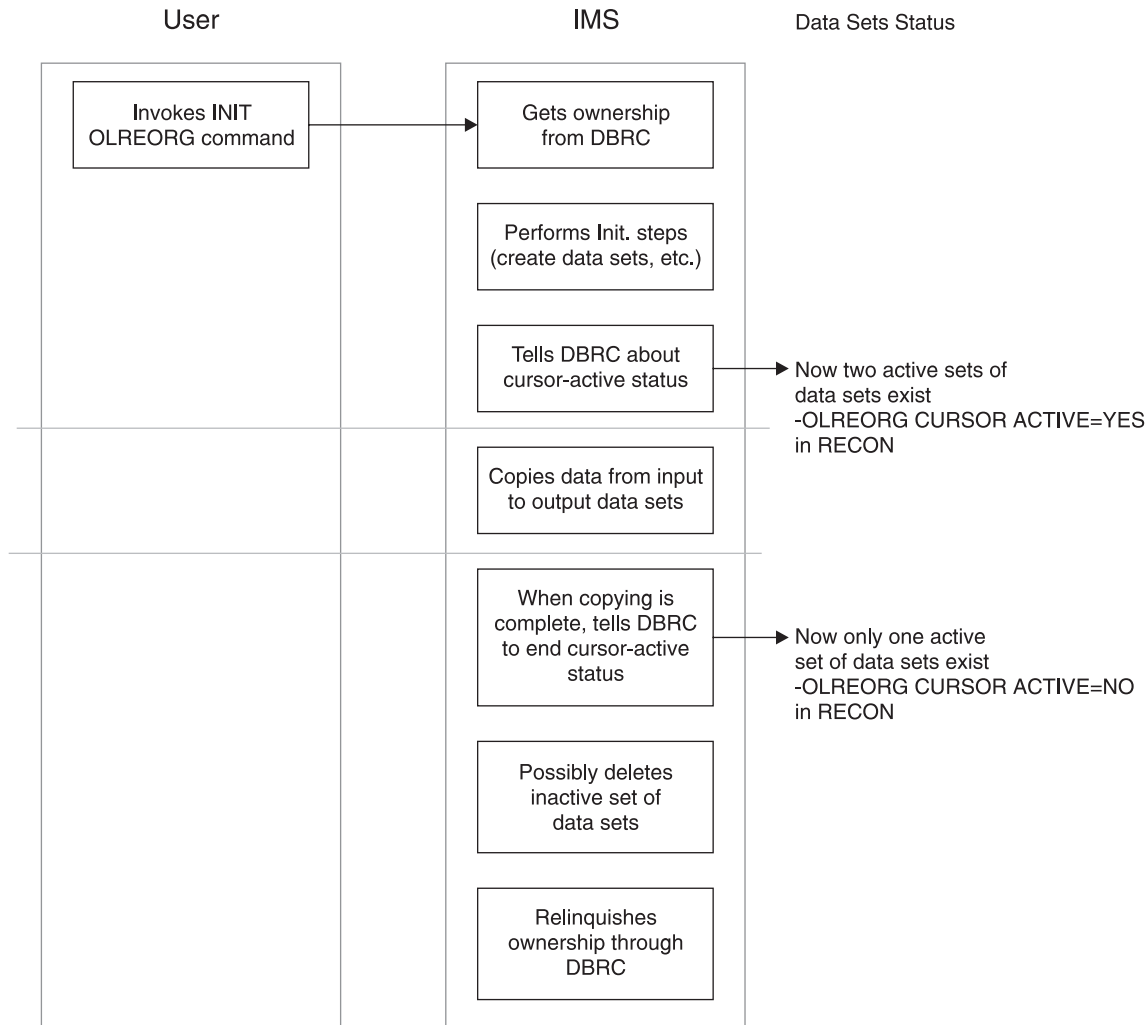


Figure 2. Normal Processing Flow of a HALDB OLR

## After the HALDB OLR Copying Process Completes

At this point in the process, the output data sets become the active set, and the input data sets become the inactive set, possibly to be used as the output in a subsequent reorganization. At this point cursor-active status in the RECON is reset. Unless an initial load or a batch reorganization reload is done, successive online reorganizations for the HALDB partition alternate between the two sets of data sets.

After the OLREORG cursor-active status has been reset, the inactive data sets are deleted if the DEL option is in effect. IMS tells DBRC to relinquish this IMS's ownership of the reorganization.

### **Related Reading:**

- For more detailed information on the DBRC commands that support HALDB OLR, see "INIT.DB" on page 31, "NOTIFY.REORG" on page 32, "CHANGE.DB" on page 56.
- For more detailed information on the RECON records that support HALDB OLR, see Chapter 9, "RECON Listing HALDB OLR Support," on page 79.

---

## Naming Conventions for HALDB OLR Data Sets

HALDB OLR has extended the data definition and data set naming convention established for HALDBs. Multiple data set groups in a HALDB database use the characters A-through-J in the DDNAMEs and data set names of the supported ten data set groups. The primary index for a PHIDAM database uses the character X in these names. This has been expanded by implementing the characters M-through-V and Y for an alternate (or paired) set of data sets.

HALDB data set names have the following format:

***bbbbbbb.dppppp***

where:

***bbbbbbb***

The data set base name of up to 37 characters (as you have defined it with the HALDB Partition Definition Utility or DBRC batch command). This same base name is used for every data set within a HALDB partition.

***d*** An IMS-assigned data set name type character that uniquely identifies a specific data set for a HALDB partition. The single-character values are:

**A, B, ..., or J** where A corresponds to the first (or only) data set group defined in the DBD, B corresponds to the second, and so on.

The data sets in the A-through-J and X set are distinguished from those in the M-through-V and Y set in this way: the initial load or batch reorganization reload of a HALDB partition is always into the A-through-J and X data sets. If the online reorganization has not been run for the HALDB partition since the last initial load or batch reorganization reload, only the A-through-J and X data sets are used.

**M, N, ..., or V** where M corresponds to the first (or only) data set group defined in the DBD, N corresponds to the second, and so on.



The data sets in the M-through-V and Y set are not used until the first time an online reorganization is done for the HALDB partition. The use of the A-through-J and X and M-through-V and Y data sets in the reorganization is discussed in "After HALDB OLR Initialization Completes" on page 2

- L** for the ILDS (indirect list data set). The online reorganization does not make a copy of this data set.
- X** for the primary index of a PHIDAM database. This data set is the index for the A-through-J data sets and is replaced by the Y data set when the online reorganization copies the database records from the A-through-J and X data sets into the M-through-V and Y data sets.
- Y** for the primary index of a PHIDAM database. This data set is the index for the M-through-V data sets and is replaced by the X data set when the online reorganization copies the database records from the M-through-V and Y data sets into the A-through-J and X data sets.

*ppppp* is the 5-digit partition ID assigned by IMS.

Whether you create one of the output data sets, or you let IMS create it automatically, the data set name is identical to that of the corresponding input data set with the exception of the data set name type character. If the A-through-J and X data sets are the active set prior to this reorganization, the data set name for the output data set is the corresponding data set name in the M-through-V and Y set of data sets, and the opposite is also true. Table 2 shows some data set name examples:

Table 2. HALDB Data Set Name Examples

| ACTIVE BEFORE REORGANIZATION | DATA SET GROUP OR INDEX | PARTITION ID | INPUT DATA SET NAME | OUTPUT DATA SET NAME |
|------------------------------|-------------------------|--------------|---------------------|----------------------|
| A-through-J and X            | 1                       | 00003        | DH41.A00003         | DH41.M00003          |
| A-through-J and X            | Index                   | 00065        | ACCT.X00065         | ACCT.Y00065          |
| M-through-V and Y            | 2                       | 00005        | PAY.MST.N00005      | PAY.MST.B00005       |
| M-through-V and Y            | 8                       | 00001        | PAY.EMP.T00001      | PAY.EMP.H00001       |



---

## Chapter 2. Planning for HALDB OLR

These sections describe your planning considerations for HALDB Online Reorganization support:

- “Migration”
- “Coexistence”
- “Fallback” on page 8
- “Restrictions” on page 8
- “Performance Considerations for HALDB OLR” on page 9
- “Recovery” on page 12
- “DBRC Considerations” on page 15
- “Offline Reorganization” on page 19
- “Utilities” on page 19

---

### Migration

These are your migration considerations for HALDB Online Reorganization support.

- When migrating from a prior release, use a `CHANGE.RECON UPGRADE` command to upgrade the RECON to an IMS V9 RECON. At that time, all existing HALDBs are marked to indicate that they not capable of being reorganized online using HALDB Online Reorganization. Prior to running HALDB Online Reorganization for a partition, issue a `CHANGE.DB DBD(HALDB_master) OLRCAP` command on IMS V9 to change the status of the HALDB and all of its partitions to allow HALDB Online Reorganization. You can also issue a `CHANGE.DB ALL OLRCAP` command to enable HALDB OLR for all HALDB databases.
- To allow full data sharing between IMS V8 and IMS V9 with HALDBs undergoing online reorganization, you need to apply the IMS Version 8 OLR Coexistence SPE.
- HALDBs marked as capable of running online reorganization are not known to an IMS V8 without the IMS Version 8 OLR Coexistence SPE, nor are they known to an IMS Version 7.
- IMS Version 8 without the IMS Version 8 OLR Coexistence SPE and IMS Version 7 can access HALDBs marked as not capable of running online reorganization.

---

### Coexistence

These are your coexistence considerations for HALDB OLR support:

- IMS V7 systems and IMS V8 (without the OLR Coexistence SPE applied) can still access the HALDBs defined as not capable of running HALDB OLR. You can create new applications that exploit HALDB OLR capable HALDBs before all IMS systems have been upgraded to IMS V9. IMS V8 with the OLR Coexistence SPE applied is also able to access online reorganization capable HALDBs.
- As part of DBRC RECON upgrade, existing HALDBs are marked to indicate that they are not capable of being reorganized online using HALDB OLR.
- You can use the `CHANGE.DB DBD(HALDB master) OLRCAP` command in IMS V9 to allow HALDB OLR.
- HALDBs marked as capable of running HALDB OLR are not known to an IMS V8 without the IMS Version 8 OLR Coexistence SPE, or to IMS V7.
- To allow full data sharing between IMS V8 and IMS V9 with HALDBs undergoing online reorganization, you must apply the IMS Version 8 OLR Coexistence SPE.

- IMS V9 must always be used to process HALDBs defined as capable of being reorganized online for these utilities:
  - Database Recovery
  - Database Image Copy
  - Database Image Copy 2
  - Database Change Accumulation

These offline reorganization IMS V8 database utilities with the IMS Version 8 OLR Coexistence SPE applied, fail when they encounter a HALDB marked as HALDB OLR capable.

---

## Fallback

These are the fallback considerations for HALDB Online Reorganization support.

**For IMS Version 8 with the OLR Coexistence SPE-** Full data sharing is allowed and no action is necessary. However, if any M-through-V and Y DBDSs are active or the online reorganization cursor is active for any partitions, then you must run the IMS V9 utilities whenever processing those partitions. The exceptions are Batch Backout, Log Recovery, and Log Archive that must be run on the release of IMS that created the partitions.

You will need to run offline reorganization for any partitions with M-through-V and Y DBDSs active or for any partitions with online reorganization cursor active before falling back to use the IMS V8 utilities.

**For IMS Version 8 without the OLR Coexistence SPE and for IMS Version 7-** Databases defined as not capable of being reorganized online are available and no action is required. Databases defined as capable of being reorganized are unavailable until these actions are completed:

- Reorganize all partitions with the M-through-V and Y data sets active or online reorganization cursor-active using the IMS V9 offline reorganization utility. This leaves only the A-through-J and X data sets active.
- Tell DBRC to set databases as incapable of being reorganized online using the IMS V9 command (`CHANGE.DB DBD(HALDB_master) OLRNOCAP`).

Should fallback to a prior version become necessary, you must set DBRC to OLRNOCAP for all HALDBs.

---

## Restrictions

- To activate Online Reorganization using an OM API (Operations Manager Application Programming Interface) command, an IMS Common Service Layer that includes Operations Manager and Structured Call Interface is required. See the *IMS Version 9: Common Service Layer Guide and Reference* for details.
- HALDB Online Reorganization only runs in an LS0=S (local storage option-subordinate) environment. An attempt to initiate online reorganization in an LS0=Y (local storage option- yes) IMS environment is rejected.  
For more information about the LSO specification, see the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.
- Data definitional changes are not supported by the HALDB OLR process. This reorganization process only provides reclustering and space distribution advantages.

- The HALDB OLR commands do not support keyword synonyms. Only one version of each keyword can be specified for a given command.
- These commands are not allowed to process against a HALDB partition on an IMS system while HALDB OLR is running against that partition on that IMS:
  - /START DB or UPDATE DB NAME(name) START(ACCESS)
  - /DBRECOVERY DB or UPDATE DB NAME(name) STOP(ACCESS)
  - /DBDUMP DB or UPDATE DB NAME(name) STOP(UPDATES)
  - /STOP DB or UPDATE DB NAME(name) STOP(SCHD)

If one of these commands is attempted on a HALDB partition that is actively being online reorganized, error message DFS0488I DBN=*database name* RC=58 is displayed and the named DL/I partition is not processed by the command.

For a HALDB master, these commands are not allowed:

- /START DB with ACCESS UP or UPDATE DB NAME(name) START(ACCESS)
- /DBRECOVERY DB or UPDATE DB NAME(name) STOP(ACCESS)
- /DBDUMP DB or UPDATE DB NAME(name) STOP(UPDATES)

#### **Related Reading:**

- For more information about message DFS0488I, see Chapter 7, “Messages and Codes of HALDB OLR,” on page 59.
- For more information about these commands see the *IMS Version 9: Command Reference*.
- HALDB OLR cannot be started for a partition if it is already running for that partition.
- Image copy is not allowed if the HALDB OLR cursor is active.
- Batch backout must be run using a DLI region type to backout in-flight HALDB OLR work.
- HALDB OLR can only be run for HALDBs defined in RECON as online reorganization capable (OLRCAP). For more information on the OLRCAP parameter, see “INIT.DB” on page 31 or “CHANGE.DB” on page 56.
- Existing output data sets must meet certain requirements, see “Required Characteristics for Existing Output Data Sets” on page 21.
- You cannot run HALDB OLR on an XRF alternate system or on an RSR tracking system.
- For information on DBRC command restrictions for HALDB OLR, see “GENJCL.IC and GENJCL.OIC Considerations” on page 15 and “NOTIFY.IC and NOTIFY.UIC Considerations” on page 15.

---

## **Performance Considerations for HALDB OLR**

These sections describe performance effects of global locking, system restart, XRF, and logging.

### **Global Locking**

To maintain partition integrity and recovery at every step of HALDB OLR processing, logging and locking operations are necessary. The type of locking is global locking because every IMS in the IMSplex is affected. Global locking involves initiating lock and unlock requests through the IRLM lock structure in the CF (coupling facility). The performance overhead for each locking operation consists of the synchronous structure access (during which every micro-second is counted as TCB busy) and MVS XES (cross-system extended services) path length.

Each HALDB OLR running at full speed (RATE(100)) incurs a significant number of IRLM lock structure accesses per second. This is a significant part of configuration capacity planning in implementing HALDB OLR. When the CF is over utilized, necessary CF related operations can be delayed and the whole IMSplex performance can be degraded.

**Recommendations:**

- Consider using a second subpool to relieve DB buffer contention for more than 4 concurrent OLRs.
- Use IBM CFSizer<sup>tm</sup> to model the additional CF activities to ensure that your CF configuration is capable of handling the extra load introduced by HALDB OLR.
  - For IRLM 2.1 with PC=N0 specified, each additional 1K of concurrent locks held requires 256 KB of ECSA storage.
  - For IRLM 2.2, the additional storage is utilized from IRLM private above the 64-bit line. No increase in ECSA storage is necessary.
- Review your LOGL latch contention rate, OLDS logging rate, IRLM lock structure access, and DBBP (for OSAM) latch contention.

Here is an example of DFSVSNxx:

```
CFNAMES,CFIRLM=IRLM91,CFVSAM=VSAM91
CFNAMES,CFOSAM=OSAM91
OPTIONS,BGWRT=YES,VSAMFIX=(BFR,IOB)
SBONLINE,MAXSB=8000
POOLID=VALL,FIXBLOCK=YES,FIXDATA=YES,STRINGMX=96
VSRBF=2048,2000
VSRBF=4096,2000
VSRBF=8192,1000
VSRBF=12288,400
IOBF=(2048,400,Y,Y)
IOBF=(4096,2000,Y,Y)
IOBF=(8192,1000,Y,Y)
IOBF=(12288,400,Y,Y)
POOLID=CUSA,FIXBLOCK=YES,FIXDATA=YES,STRINGMX=32
VSRBF=4096,1000,D
VSRBF=4096,400,I
IOBF=(8192,1000,Y,Y,CUOA)
DBD=PCUST12(A,CUOA)
DBD=PCUST12(L,CUSA)
DBD=PCUST12(X,CUSA)
DBD=PCUST14(A,CUOA)
DBD=PCUST14(L,CUSA)
DBD=PCUST14(X,CUSA)
DBD=PCUST16(A,CUOA)
DBD=PCUST16(L,CUSA)
DBD=PCUST16(X,CUSA)
DBD=PCUST18(A,CUOA)
DBD=PCUST18(L,CUSA)
DBD=PCUST18(X,CUSA)
OLDSDEF OLDS=(00,01,02,03,04,05,06,07,08,09),BUFNO=400,MODE=DUAL
WADSDEF WADS=(0,1,2)
```

## Restart, XRF, and RSR with HALDB OLR

All active HALDB OLRs are suspended before a shutdown checkpoint. They are resumed automatically by IMS at the next restart.

All HALDB OLRs that were active (running) when IMS abnormally terminated will have any uncommitted changes dynamically backed out to the last sync point using the same logic that is used to backout online programs. They will be resumed automatically by IMS at the next restart.

Similarly, when an XRF takeover occurs, the HALDB OLRs are resumed on the new active IMS.

When an emergency restart or a warm start occurs, any HALDB OLRs that were active when IMS abnormally terminated or shutdown are automatically resumed on the new active IMS. In the case of RSR, the HALDB OLRs that were active at the time of RSR takeover are automatically resumed at the takeover restart.

After the ILDS and primary index data sets are rebuilt by DFSPREC0, issue the INITIATE OLREORG command to resume the HALDB OLRs.

## Logging Effects of Reorganization

The logging volume that is attributable to online reorganization is important to understand because the reorganization of a partition generates x'50' database change log records for all of the data in the partition. Because of a small number of x'29' log records that support HALDB OLR, which reflect the progress of the reorganization and because of the prefix size of the x'50' log record, the total amount of log data from the reorganization is considerably greater than the amount of data in the partition.

The reorganization of multiple partitions in parallel on the same IMS can generate sufficient log data to impact normal transaction processing. The rate of OLDS log switches and their archiving can certainly be affected.

## Reorganization Processing Time

An online reorganization's impact on the system is affected by the available system resources, by total system utilization (including other online reorganizations), by total logging volume, by log contention, and by the intensity at which this reorganization was set to run. These same factors also affect the speed at which the reorganization runs.

You can use the RATE parameter as a dynamic throttling knob in the INITIATE OLREORG and UPDATE OLREORG commands to control the intensity at which the reorganization runs. This can affect not only the reorganization's speed but also its impact on the rest of the system. The value specified for RATE is the percentage of elapsed time to be devoted to copying records. The remaining time is an intentionally introduced delay in the copying process that minimizes the reorganization's impact on other IMS work and on the whole system.

RATE(100) is the default that allows the online reorganization to run as fast as possible (depending on system resources, system contention, log contention, etc.) with no intentional delays. A RATE value of 50 specifies that 50% of the elapsed time should be spent copying records and the remaining 50% should be spent in a delay. This would cause the reorganization to run approximately twice as long as it would have run with RATE value of 100. A RATE value of 25 would cause the reorganization to take four times as long, and so on. The RATE value can be changed with the UPDATE OLREORG command at any time that adjustments are desired because of the amount of system contention. With higher system contention, the RATE value can be lowered to slow down the online reorganization and minimize its impact on other work. With less system activity, the RATE value can be raised, thus speeding up the reorganization.

## Activating Sequential Buffering for OLR

Online Reorganization of OSAM databases can use Sequential Buffering (SB) to improve performance. SB can be activated by including the SBONLINE statement in DFSVSMxx.

Including the SBONLINE statement loads all SB modules at IMS initialization, and then whenever an Online Reorganization is initiated for an OSAM database, SB is activated and is started immediately. Including the SBONLINE statement eliminates the need to analyze several hundred DLI calls to determine whether SB is suited for this processing, thereby improving performance.

The two forms of the SBONLINE control statement are:

SBONLINE

or

SBONLINE,MAXSB=nnnnn

where nnnnn is the maximum storage (in kilobytes) that can be used for SB buffers.

When the MAXSB limit is reached, IMS stops allocating SB buffers to online applications, including OLR, until terminating online programs and OLR release SB buffer space. If you do not specify the MAXSB= keyword, the maximum storage for SB buffers is unlimited. Detailed instructions for coding the SBONLINE control statement are located in the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

---

## Recovery

These sections describe your recovery considerations for HALDB Online Reorganization support.

### Database and Data Set Recovery

There are several differences in the management of database backup and recovery both during and after an online reorganization.

**Restriction:** Only the data sets from the active set or sets of data sets can be recovered. This means that while the cursor-active status is in effect, data sets from either the input set or the output set can be recovered, but while the cursor-active status is not in effect, only data sets from the single active set of data sets can be recovered.

#### No Database Image Copy Allowed During Reorganization

While a partition is in cursor-active status, none of the three Database Image Copy utilities can copy any of the data sets for that HALDB partition. Even if the TERMINATE OLREORG command is used to temporarily stop the copying of database records from the input to the output data sets, an image copy still cannot be made for any data set in the HALDB partition.

#### Log-Only Recovery

After the *cursor active* status for the online reorganization is recorded in the RECON, recovery of any of the output data sets with the Database Recovery Utility (DFSURDB0) can be accomplished solely from database change records (type X'50' log records) until the reorganization has completed and a subsequent image copy has been made. An empty output data set is considered to be the starting point from which database change records can be applied.



If recovery of an output data set is required before the online reorganization has completed, then proceed as follows:

1. Ensure that the reorganization has been temporarily stopped. If the reorganization encountered an abend, this is automatic; otherwise, the TERMINATE OLREORG command may have to be used to temporarily stop a currently running reorganization.
2. Enter the /DBR or the UPDATE DB command for the HALDB partition.
3. Run database change accumulation as necessary. The JCL can be created by the DBRC GENJCL.CA command, or the Database Change Accumulation utility can be run from your JCL. The purge time for change accumulation must be the time of the beginning of the reorganization to represent restoring from the initial empty state of the data set. There are further details on the DBRC GENJCL command and on the utility in the paragraphs that follow.
4. Create the output data set that must be recovered, either through a JCL DD statement or with Access Method Services, as appropriate.
5. Recover the database changes. The JCL can be created by the DBRC GENJCL.RECOV command. Alternatively, the Database Recovery utility can be run from your JCL; in this case, the DD statement for DFSUDUMP must specify DUMMY, which indicates that there is no image copy from which to restore.
6. Execute the Batch Backout utility as backout might be required to backout uncommitted data. Note that the operational aspect of online reorganization in this case is very much like an online application program.  
After all data sets of the HALDB partition have been recovered and possibly backed-out proceed to the next step.
7. Enter the /STA or the UPDATE DB command for the HALDB partition.
8. Resume the reorganization using the INITIATE OLREORG command.

It is also possible to recover an output data set after the reorganization has completed but before an image copy has been made. The procedure to recover an output data set after the reorganization has completed but before an image copy has been made is similar to the procedure to recover an output data set is required before the online reorganization has completed except that, the reorganization need not be terminated or resumed. Although you can recover an output data set using this method, it is *strongly recommended* that you make an image copy of the output data sets as soon as it is practical after the reorganization. Recovering from this image copy is faster than recovering from the database change records logged during the reorganization.

You can also recover an output data set from a point other than the beginning of the reorganization, such as from a full dump of a DASD volume, using existing procedures if the reorganization has been either completed or terminated.

To recover an input data set before the reorganization completes, your only additional consideration is to ensure that the reorganization has been terminated before starting the recovery.

### **Database Change Accumulation Utility:**

When doing database change accumulation for one of the output data sets, specify a purge time equal to the online reorganization start time on the control statement. This is necessary if change accumulation records (or an input log) involving the output data set span the time that an online reorganization was started. This eliminates database change records from before this point; this is analogous to eliminating database change records from prior to the start time of an image copy.

### DBRC GENJCL.CA and GENJCL.RECOV Commands:

Even though no image copy exists for the output data sets, the RECON reflects the beginning of the reorganization as a starting point from which forward recovery of one of these data sets can be done, even after the reorganization is finished. Until an image copy has been made for an output data set, the DBRC GENJCL.CA command treats this starting point as though it were the most recent image copy and causes changes to the output data set to be accumulated from that point. Similarly, GENJCL.RECOV prepares recovery of an output data set from this point even though no physical image copy exists.

### Generic DDNAME for Database Image Copy

Database Image Copy always copies the currently active data set that is recorded in RECON.

Regardless of whether the A-through-J or the M-through-V data sets are active, the JCL and control statements for the three Database Image Copy utilities need not be changed when the active set of data sets switches from the A-through-J to the M-through-V data sets, or the other way around. You do not need to be aware of which set of data sets is the active set and you need not provide this information for these reasons:

- On the utility control card for the Database Image Copy utility, the DDNAME need not refer to the currently active data set. No matter whether the A-through-J or the M-through-V data sets are active, the specified DDNAME is *automatically understood* by the utility to refer to the corresponding active data set.

For example, assume that the data set for the second data set group defined in the DBD is to be copied and that the partition name is PARTNO3. Regardless of which set of data sets is active, a DDNAME of either PARTNO3B or PARTNO3N may be coded on the control card. If the A-through-J data sets are active, then the data set identified by the DDNAME PARTNO3B is copied. If the M-through-V data sets are active, then the data set identified by the DDNAME PARTNO3N is copied.

- In the JCL statements for the Database Image Copy utility, the DD statement referring to the input data set should be omitted.

Based on whether the A-through-J or the M-through-V data sets are active, the appropriate data set is dynamically allocated. Having a DD statement that refers to a specific data set name can cause the job to fail because of a “Data Set Not Found” condition during job step initiation. This would occur if an inactive data set name was coded in the JCL and it does not exist.

## Remote Site Recovery

Remote Site Recovery (RSR) supports HALDB Online Reorganization in much the same fashion as normal database update processing is supported. Unlike offline reorganization, there is no need for the databases to be brought offline when HALDB Online Reorganization is recognized by the RSR Tracking IMS.

As part of RSR tracking, the tracking site RECON is updated with HALDB Online Reorganization related information in both Recovery Level Tracking (RLT) and Database Level Tracking (DLT) modes. If the tracking IMS is running in the DLT mode and the HALDB partition is defined as DBTRACK in the tracking site RECON, the tracking process includes:

- Creating the output data sets for the shadow partition as needed
- Updating both sets of shadow data sets

- Designating the original input data sets as inactive and designating the output data sets as the single active set in order to track the completion of the reorganization
- Deleting the inactive data sets if OPTION(DEL) was in effect when the reorganization completed at the active site

Just as at the active site, errors can occur during the validation or creation of the output data sets at the beginning of the reorganization. The shadow partition is stopped when this happens, and this has no effect on the continuation of the reorganization at the active site. After you have corrected the problem, start the partition to continue tracking the reorganization's updates in OFR mode.

If the output data sets for the reorganization already exist before the tracking of the reorganization begins, ensure that these data sets have same characteristics (block size, record size, control interval size, and so on) as those at the active site.

- If data set characteristics are manually changed at the active site, they must be manually changed at the tracking site for partitions undergoing online reorganizations.
- After RSR takeover, partitions undergoing online reorganization will be stopped. When the primary index and indirect list data sets have been rebuilt after an RSR takeover, the INITIATE OLREORG command must be issued if previously running online reorganizations need to be resumed.

HALDB OLRs that were active at the time of RSR takeover will not be automatically resumed at the takeover restart. When the ILDS and primary index data sets are being rebuilt by DFSPREC0, the INITIATE OLREORG command must be issued to resume the HALDB OLRs.

---

## DBRC Considerations

The following sections describe the things that must be considered for online reorganization in relation to DBRC (Database Recovery Control).

### GENJCL.IC and GENJCL.OIC Considerations

The GENJCL.IC and GENJCL.UIC JCL can be generated in advance and not executed immediately. So, the active set of DBDSs may have changed in the interim time period.

Since dynamic allocation is recommended for the utilities, GENJCL no longer generates DD statements for the HALDB DBDS that is being copied. The SYSIN control statement identifies which DBDS to copy. This can be either the A-through-J or M-through-V set of data sets. The image copy utility then copies the corresponding active set of data sets, A-through-J or M-through-V.

### NOTIFY.IC and NOTIFY.UIC Considerations

These are the HALDB OLR support modifications for the NOTIFY.IC and NOTIFY.UIC commands.

An image copy cannot be taken while OLR cursor is active. The NOTIFY.IC and NOTIFY.UIC commands fail if they attempt to record an image copy while OLR cursor is active.

## DBRC RECON HALDB OLR Information

The DB record for a HALDB partition contains information for HALDB Online Reorganization:

- A field that indicates whether the database is capable of being reorganized online.
- A field that indicates whether online reorganization cursor is active.
- The SSID of the IMS currently performing an online reorganization.
- A field indicating which set or sets of DBDSs is or are active for the partition.
- Counters showing the number of cursor-active status changes and cursor-inactive status changes that are not yet hardened at the RSR tracking site.

The HALDB Master and Partition DB records contain a new field indicating if the HALDB can be reorganized online.

The M-through-V and Y DBDS records are defined in RECON with the exact same attributes as the corresponding A-through-J and X DBDS records. One difference, in addition to DD and data set names, is the DSID. The DSID of the M-through-V and Y data set is the same as the corresponding A-through-J and X data set with its X'80' bit on. For RECON list purposes, the X'80' bit is ignored.

Table 3. DSID DCB Numbers for the A-J and X and the M-V and Y data sets

| DSID / DCB<br>Number | A-J, L, and X Data Sets |       | M-V and Y Data Sets |       |
|----------------------|-------------------------|-------|---------------------|-------|
|                      | PHIDAM                  | PHDAM | PHIDAM              | PHDAM |
| 1                    | A                       | A     | M                   | M     |
| 2                    |                         |       |                     |       |
| 3                    | L                       | L     |                     |       |
| 4                    |                         | B     |                     | N     |
| 5                    | X                       | C     | Y                   | O     |
| 6                    | B                       | D     | N                   | P     |
| 7                    | C                       | E     | O                   | Q     |
| 8                    | D                       | F     | P                   | R     |
| 9                    | E                       | G     | Q                   | S     |
| 10                   | F                       | H     | R                   | T     |
| 11                   | G                       | I     | S                   | U     |
| 12                   | H                       | J     | T                   | V     |
| 13                   | I                       |       | U                   |       |
| 14                   | J                       |       | V                   |       |

The M-through-V and Y DBDS records are added to the indicated change accumulation group. That is, to the same CAGROUP as the A-through-J and X DBDS.

The M-through-V and Y DBDS records are added to each DBDS group where the A-through-J DBDS records are found.

The REORG record for a HALDB partition DBDS contains new information (stop time) for HALDB Online Reorganization. In order to run an offline reorganization of it, a database must have been taken offline successfully with a /DBR command. So,

only the start time of the reorganization is necessary. With online reorganization, the database is online so the duration of the process is important. The stop time of the reorganization is now recorded in the REORG record along with the new indicator showing that this is an online reorganization and whether the online reorganization may be used as input to recovery.

A REORG record is recorded under both the A-through-J and X DBDS and the M-through-V and Y DBDS. The REORG record for the DBDS that was the output data set indicates that it can be used as input for recovery. The REORG record for the DBDS that was the input data set indicates that it cannot be used for recovery after online reorganization has completed.

The online reorganization capability is set at the HALDB Master DB level.

The DB Record header record contains a new field which indicates whether online reorganization is enabled for this partition. Once enabled, sharing of HALDBs with an IMS Version 7, or IMS Version 8 without the OLR Coexistence SPE, is no longer allowed. Full data sharing is allowed for IMS Version 8 with the OLR Coexistence SPE applied. IMS Version 7 or IMS Version 8 without the OLR Coexistence SPE applied are allowed to process databases marked as not capable of being reorganized online. This field is set or reset with the CHANGE.DB DBD(HALDB master) command and is displayed in the LIST.DB output.

Table 4 on page 17 shows whether full data sharing of databases is allowed between these IMS Versions:

- Version 7
- Version 8 (without the OLR Coexistence SPE applied)
- Version 8 (with the OLR Coexistence SPE applied)
- Version 9

The ability to share data is also dependent on whether OLRCAP or OLRNOCAP is specified for a database and this table also depicts those cases.

*Table 4. Data Sharing Between the Releases*

|                 | Version 7 | Version 8<br>Without SPE<br>Applied | Version 8 With<br>SPE Applied | Version 9 |
|-----------------|-----------|-------------------------------------|-------------------------------|-----------|
| <b>OLRCAP</b>   | NO        | NO                                  | YES                           | YES       |
| <b>OLRNOCAP</b> | YES       | YES                                 | YES                           | YES       |

Resetting the field in the DB record header to disable online reorganization fails if online reorganization cursor is active, or if any of the M-through-V and Y data sets are active. Resetting the field in the DB record header to disable online reorganization also fails if any partition for the HALDB is authorized. If resetting the field is permitted, all of the M-through-V and Y data set information in RECON is deleted.

## RSR Tracking System

Ownership of the HALDB Online Reorganization is not recorded in tracking RECONS so the tracking RECONS do not reflect ownership of online reorganizations.

RECONS at the tracking site will always reflect the current status of the covered databases. New counters keep track of outstanding online reorganizations that have been recorded in RECON at the tracking site but not yet applied to the databases.

## Recovery Related Records

The A-through-J and M-through-V data sets are treated as one logical set of DBDSs. The M-through-V data sets are defined in RECON with the same attributes as their corresponding A-through-J data sets.

The recovery information in the two DBDS headers is maintained as if they are one logical DBDS. For example: GENMAX, RECOVPD, and CAGRP are the same. Any changes that are made to the active DBDS are also made to the inactive DBDS.

There is a one-to-one relationship between each partition data set- for example A-to-M, B-to-N, C-to-O, and so on. Either set of data sets (A-J and X or M-V and Y) can be reorganized online at any time depending on which set is active prior to the start of the online reorganization. So data sets A and M are one logical data set during an online reorganization. After an image copy is taken, DBRC removes any extraneous recovery related records (IMAGE, ALLOC, RECOV, and REORG) from RECON. What is extraneous is determined by the GENMAX and the optional RECOVPD settings for the DBDS. GENMAX and RECOVPD apply to the A-through-J or X data sets and the corresponding M-through-V or Y data sets as if the two data sets were one logical data set.

For example, GENMAX=2 for data set A (and M) and data set A has two image copies. When data set M becomes the active data set and an image copy is taken, image copy cleanup results in deleting the oldest image copy for data set A along with any of its extraneous recovery records. Another image copy of data set M results in the rest of data set A's image copy and recovery records being deleted. Note that data set A would become unrecoverable in this case.

## Online Reorganization Execution

When online reorganization is started, the ID of the IMS responsible for the reorganization is assigned to the partition. Once the cursor-active status is recorded in the RECON and copying begins between the A-through-J and M-through-V data sets these things also happen:

- The DB record for the partition indicates 'OLREORG CURSOR ACTIVE =YES'.
- REORG records are written showing the start of the reorganization (the stop time is zero) indicating online.
- The REORG record under the data set that is the output indicates that it can be used as input for recovery.
- The REORG record under the data set that was the input does not indicate that it can be used for recovery. DBRC uses this record to ensure that recovery of this data set cannot occur after online reorganization is complete.
- The M-through-V and Y data sets are defined in RECON if they do not already exist.
- The M-through-V and Y DBDSs are added to the same change accumulation group as the A-through-J and X DBDSs, if one is assigned.
- The M-through-V and Y DBDSs are added to any DBDS group where the A-through-J and X DBDSs are found.

At this point the online reorganization must either complete or an offline reorganization must be run in order to reset the online reorganization settings.

Online reorganization cannot be started if it is already running for this partition. However, another IMS can take over the reorganization if it is terminated; that is, when there is no IMS ID assigned.



When the entire online reorganization process completes successfully the following things happen:

- The REORG records are updated to include the online reorganization stop time.
- The DB record for the partition is changed to indicate 'OLREORG CURSOR ACTIVE = NO'.
- The status of the active DBDSs is switched to the newly reorganized data sets.
- The online reorganization IMS ID is cleared.

**Attention:** Successful completion of an online reorganization does not require that image copies of the partition's DBDSs be taken. However, it is recommended that you take an image copy as soon as possible to reduce recovery time.

**Restriction:** An image copy of a DBDS cannot be taken while online reorganization cursor is active. DBRC causes image copy to fail in this case.

---

## Offline Reorganization

Offline reorganization can only be run if online reorganization cursor is not active or online reorganization was terminated. Successful completion of an offline reorganization results in the active data sets being reset to A-through-J and the online reorganization status, along with its IMS identifier, being reset. The REORG records for the incomplete online reorganization are deleted.

---

## Utilities

These are the HALDB online reorganization considerations for running these IMS utilities: Primary Index and ILDS Rebuild, Change Accumulation, Recovery, Image Copy, and Batch Backout.

### HALDB Index / ILDS Rebuild - DFSPREC0

DFSPREC0 recovers the PINDEX for both input and output if RECOVTYP=INDEX is specified and will do the same if RECOVTYP=ILE is specified when the index is recovered. If RECOVTYP=BOTH is specified then DFSPREC0 recovers both the X and Y PINDEX data sets and the ILDS in a single run. There is no way to specify a particular input or output data set but if OLR was active then all of the necessary data sets are allocated and rebuilt.

**Related Reading:** For more information on DFSPREC0, see the *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

### Database Change Accumulation - DFSUCUM0

DFSUCUM0 accumulates changes for HALDB partition A-through-J data sets and for M-through-V data sets. A DBO control statement is needed to indicate to the utility that it is to accumulate changes or that it is to purge changes before online reorganization started.

CA ([Database] Change Accumulation) header records (type 25 records) may be created for a corresponding A-through-J and M-through-V data set if the online reorganization checkpoint has not yet completed at the time of the CA.

When doing database change accumulation for one of the output data sets, specify a purge time equal to the online reorganization start time on the control card. This is necessary if change accumulation records (or an input log) involving the output data set span the time that an online reorganization was started. This eliminates

database change records from before this point; this is analogous to eliminating database change records from prior to the start time of an image copy.

## Database Recovery - DFSURDB0

**Restriction:** For recovery of a data set using the DFSURDB0 utility IMS expects the data sets to exist and makes no attempt to create them.

For information on the HALDB online reorganization considerations for running the Database Recovery utility, see “Log-Only Recovery” on page 12.

## Database Image Copy - DFSUDMP0

| Database Image Copy always copies the currently active data set that is recorded  
| in RECON.

| Database Image Copy processing works this way for HALDBs: An image copy of a  
| DBDS cannot be taken while HALDB online reorganization cursor is active on that  
| DBDS. DBRC fails the execution of an image copy in this case.

| Database Image Copy inputs remain the same as they were. For each request to  
| copy HALDB partition A-through-J data sets or M-through-V data sets, Database  
| Image Copy issues an error message if HALDB online reorganization is cursor  
| active on the HALDB partition. Database Image Copy also determines if the  
| M-through-V data sets are to be copied instead of the A-through-J data sets.

It is no longer necessary to code a DD statement in the JCL when copying HALDB partition data sets as they will be dynamically allocated.

**Attention:** Successful completion of an online reorganization does not require that image copies of the partition's DBDSs be taken. However, it is recommended that you take an image copy as soon as possible afterward to reduce recovery time.

## Batch Backout - DFSBBO00

| During initialization of HALDB online reorganization, a PSB is dynamically built  
| internally and used for the reorganization. The name of the PSB that is built and  
| scheduled is constructed by prefixing the 7-byte HALDB partition name with a  
| numeric character zero ('0'). This PSB does not exist in PSBLIB or ACBLIB and is  
| explained here for IMS cold start situations. For example, a HALDB partition with  
| the name SSN5603 has a dynamic PSB with the name 0SSN5603 for the  
| reorganization work.

When it is necessary to cold start an IMS system, you would normally run the Log Recovery utility (DFSULTR0) with the option to list the PSBs that require backout. This utility lists these online reorganization dynamic PSBs if there is in-flight reorganization work that requires backout.

Batch Backout (DFSBBO00) needs to be run for each of the listed PSB names. Batch Backout recognizes that the named PSB is from online reorganization because it begins with a numeric character (this is not allowed for defined PSBs).

**Restriction:** Batch Backout must be run using a DLI region type to backout in-flight online reorganization work.



---

## Chapter 3. The Data Sets of HALDB OLR

These sections discuss the details concerning the data sets that are used during a HALDB OLR.

**Restriction:** The data sets used by HALDB OLR must be either PHDAM or PHIDAM type only.

---

### Output Data Set Creation

Creation of the output data sets is considered on a data-set-by-data-set basis. For example: the *M* data set might already exist, while the *N* and *O* data sets need to be created by IMS, and the *P* data set might also already exist. So, in this case, IMS creates the *N* and *O* data sets and uses the existing *M* and *P* data sets.

The DASD space for the output data set may be smaller than, equal to, or larger than the corresponding input data set; whatever is present is used. And because the output data set is completely overwritten, it is of no consequence whether there is already any data in that data set.

---

### Required Characteristics for Existing Output Data Sets

Each pre-existing data set in the output set must meet certain requirements that make it usable as output. If these requirements are not met then one of the messages listed in Chapter 7, "Messages and Codes of HALDB OLR," on page 59 of this document, is displayed and the reorganization does not begin. Here are the major requirements for the different data set types:

**OSAM:** The data set must meet all of these requirements:

- It must be cataloged
- It must be a DASD data set
- It must not be a VSAM data set except for the primary index data set of a PHIDAM database
- It must not be a PDS, PDSE, or a member of a PDS or PDSE

**VSAM:** The data set must have all of these characteristics:

- It must be a VSAM entry-sequenced data set (ESDS) except for the primary index data set of a PHIDAM database
- It must have the REUSE attribute
- It must have a fixed-length record length identical to that of the corresponding input data set
- It must have a control interval size identical to that of the corresponding input data set
- It must have a SHAREOPTIONS attribute value at least as high as that of the corresponding input data set if the database is defined to DBRC with a SHARELVL value of 2 or 3

**PRIMARY INDEX:** The data set must have all of these characteristics:

- It must be a VSAM KSDS (key-sequenced data set)
- It must have the same key offset and length as the corresponding input KSDS
- It must have the other required characteristics listed under "VSAM:"

## Characteristics for Automatically-Created Output Data Sets

These topics discuss the details about size, location, and other characteristics of an automatically-created output data set. If you require further control over any of these characteristics, then create the data set before you issue the INITIATE OLREORG command.

### Number of Volumes:

If a particular input data set is SMS-managed, then the corresponding output data set is created with the same number of volumes.

If the input data set is not SMS-managed, then automatic creation of the corresponding output data set is done only when the input data set resides on a single volume. When a non-SMS-managed input data set resides on multiple volumes, then the corresponding output data set must exist before the HALDB OLR is started. Violation of this restriction causes an error message to be displayed, and the reorganization does not begin.

### Location, SMS-Managed Case:

If a particular input data set is SMS-managed, then the corresponding output data set is also SMS-managed, and the storage class used to create the output data set is that of the input data set. The site's storage administrator must ensure either, that this storage class refers to a storage group with sufficient space to hold the output data set or that the automatic class selection (ACS) routine selects an appropriate storage class for the data set.

### Location, Non-SMS-Managed, Non-VSAM Data Set Case:

Regardless of the type of DASD on which the input data set resides, the corresponding non-VSAM output data set is created as though the following parameter had been specified on its DD statement:

```
UNIT=SYSALLDA
```

The creation of the output data set does not request any specific volume serial number, causing the data set to be created on a storage volume or, if no storage volume is available, on a public volume.

### Location, Non-SMS-Managed, VSAM Data Set Case:

The VSAM output data set is created on the same volume as the corresponding input data set. In many cases this restriction limits the usefulness of automatically creating a VSAM data set that is not SMS-managed.

As already stated automatic output data set creation is not done when the corresponding data set is on multiple-volumes and is not SMS-managed.

### Size, Single-Volume Case:

When the input data set has extents on only one DASD volume, the output data set is also created on a single volume as though the following parameter had been specified on a DD statement:

```
VOLUME=(,,1)
```

The amount of primary space is derived from the space allocation of the input data set.

For a non-VSAM data set, this is the total amount of space in the first five extents on the volume.

For a VSAM data set, this is the primary space allocation used when the input data set was created.

In order to reserve approximately the same amount of space that was reserved for the input data set regardless of the DASD types involved, the output data set's space is requested as a number of OSAM blocks or VSAM records. If the input data set space was not already specified in one of these ways, then based on the input data set's DASD characteristics, the cylinder or track allocation is converted to an equivalent number of blocks or records.

If a secondary space allocation amount is specified for the input data set, then this same secondary amount is used for the output data set. This is converted as necessary from a number of cylinders or tracks to a number of blocks or records.

**Note** that if secondary space for the input data set has been allocated on DASD, then the space for the output data set does not necessarily reflect the amount of space currently allocated or used by the input data set.

The automatically created output data set could have a considerably different amount of available DASD space than was used for the input data set.

For example, for an input data set that has used secondary allocation. The automatic creation reserves the primary space for the output data set, but there may not be enough space on the volume for secondary allocation either during the HALDB OLR or during later database processing.

## Size, Multiple-Volume Case (SMS-Managed Only):

IMS automatically creates multiple-volume output data sets only when the input data set (and therefore the output data set) is SMS-managed. In this case it is highly desirable that the storage class used to create the output data set has the GUARANTEED SPACE attribute. This storage class is either that of the input data set or the one specified by the site's ACS routine.

Using the same considerations as the single-volume case, as already discussed, the requested primary and secondary space is based on the input data set's space allocation on the FIRST volume.

As is the case for a single-volume data set, the automatically created multiple-volume output data set could have a considerably different amount of available DASD space than was used for the input data set. In many cases secondary allocation failures can be avoided if the storage class of the input data set (and thus of the output) has the GUARANTEED SPACE attribute. This causes the primary space allocation to be taken on EACH of the volumes when the output data set is created. Secondary space is then taken later, as needed. However, even this does not guarantee that the output data set gets the same amount of space as the input data set, especially when secondary space was acquired for the input.

Another way to simplify DASD space allocation is to avoid multiple-volume data sets altogether. With the larger DASD that is now available and with the ability to define up to 1001 partitions for a HALDB, the need for multiple-volume data sets is reduced.

---

## The Disposition of the Output Data Sets When Initialization Fails

As discussed in the sections titled, “Required Characteristics for Existing Output Data Sets” on page 21 and “Characteristics for Automatically-Created Output Data Sets” on page 22, any preexisting output data sets must meet certain requirements, and any output data sets that do not exist are created automatically by IMS as part of the initialization of the HALDB OLR. During this process, various error conditions can cause the HALDB OLR initialization to fail. Two error conditions that can occur are: an unacceptable preexisting output data set, or an insufficient amount of DASD space for an automatically created data set.

If any errors occur during the validation, or after the validation and creation of the output data sets, but before the *cursor-active* status of the reorganization has been recorded in the RECON, then any automatically created output data sets are retained along with any preexisting data sets.

---

## The Inactive Data Sets at Successful Completion of the Reorganization Process- Data Set Deletion

As part of the successful completion of the HALDB OLR, all of the input data sets involved in the reorganization become inactive. IMS attempts to delete these data sets automatically regardless of whether they were created by IMS or were pre-existing data sets. This automatic data set deletion can be suppressed by specifying the NODEL option on the INITIATE OLREORG command or on the UPDATE OLREORG command.

This process of data set deletion occurs only after:

1. All the records of the HALDB partition have been copied to the output data sets.
2. The successful completion of the HALDB OLR has been recorded in the RECON, making the output data sets the active ones and the input data sets the inactive ones.
3. All IMS systems, including batch jobs, that are sharing the reorganized partition have been notified that event 2, in this list, has occurred.
4. All of the sharing IMS systems have acknowledged that they have received the notification and that they have closed and deallocated the input data sets (that are now the inactive data sets).

There is the possibility that the acknowledgement from a sharing system could be delayed for a long period of time. This is especially true for a sharing batch job. The IMS system that completed the HALDB OLR attempts to delete the inactive data sets anyway if the acknowledgements are not received within 4.5 minutes. However, in this case, the deletion is likely to fail, but reorganization still completes normally.

---

## Chapter 4. Setting Up HALDB OLR

HALDB Online Reorganization is included in IMS Version 9 and does not require special installation.

There are two types of command formats available in IMS:

- The type-1 command is preceded by a CRC (command recognition character) usually a / (slash). It is issued from terminals, AOI applications, EMCS consoles, and APPC/OTMA programs.
- The type-2 command is much like type-1 except that it has no preceding CRC and it can be issued and processed through the OM (Operations Manager) API (Application Programming Interface) on multiple databases on multiple IMSs.

Establish type-1 command security with RACF® or in the exit routine, DFSCCMD0 that you set up. The CIMS class for is 3 characters followed by blanks. The command verbs INI, TER, and UPD must be defined in the CIMS class or the command security exit routine, DFSCCMD0. Type-1 commands are valid from OM API (Operations Manager Application Programming Interface) and the command security can be set up in OM.

This topic describes how you can do these things:

- How to initiate or resume HALDB OLR.
- How to indicate whether HALDB OLR is allowed to run on a database.
- How to specify whether the reorganization is a HALDB OLR or it is an offline reorganization.

And if it is an online reorganization, how you can also specify whether the HALDB OLR can be used as input for recovery and how you can specify the time that the HALDB OLR completed successfully.

---

### INITIATE OLREORG Command

To initiate or resume HALDB OLR, issue the type-2 INITIATE OLREORG command from the OM API.

**Related Reading:** For details on the OM API see the *IMS Version 9: Common Service Layer Guide and Reference*.

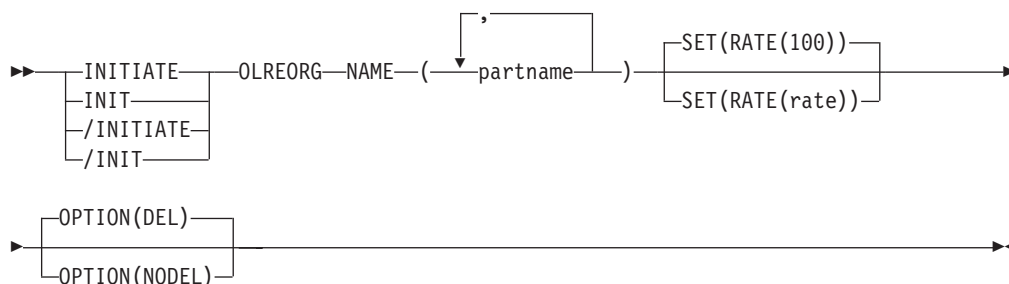
### /INITIATE OLREORG Command

The type-1 command syntax is the same as the type-2 command syntax, but is preceded by a CRC (command recognition character) usually a / (slash).

The type-1 /INITIATE OLREORG command can be issued from terminals, AOI applications, EMCS consoles, and APPC/OTMA programs. When the command is entered as a type-1 command only one NAME parameter can be specified.

### Format of the INITIATE OLREORG Command

The syntax diagram here illustrates the format of the INITIATE OLREORG command.



## Environments, Keywords, and Parameters of the INITIATE OLREORG Command

Table 5 lists the environments, keywords, and parameters of the `INITIATE OLREORG` command.

This command uses syntax that can distinguish between keywords and resource names, so none of these keywords are reserved keywords. You can have resource names that are the same as the HALDB OLR keywords.

**Restriction:** This command does not support keyword synonyms. Only one version of each keyword can be specified for each command issued.

Table 5. Valid Environments for the `INITIATE OLREORG` command Keywords, and Parameters

| Command / Keywords | DB/DC | DBCTL | DCCTL |
|--------------------|-------|-------|-------|
| OLREORG            | X     | X     |       |
| NAME               | X     | X     |       |
| SET                | X     | X     |       |
| OPTION             | X     | X     |       |

## Usage of the INITIATE OLREORG Command

The `INITIATE OLREORG` command causes HALDB OLR to begin or resume on the partition or partitions that are specified by *partname*.

The `INITIATE OLREORG` command might be routed to one or more of the IMSs in the IMSplex depending on the `ROUTE` specifications. The default for the command is `ROUTE=ALL` which means that the command is routed to all active DB/DC and DBCTL IMSs in the IMSplex. The first IMS to successfully process the `INITIATE OLREORG` command owns the online reorganization and the command fails on all of the other IMS subsystems where the command has been routed.

Part of the HALDB OLR processing involves:

- The creation of the output data sets at the beginning of the online reorganization
- The optional deletion of the inactive data sets when the online reorganization completes

If the partition is tracked at an RSR tracker site, the `OPTION` value (`DEL` or `NODEL`) that is in effect at the completion of the HALDB OLR also determines whether the inactive data sets for the shadow partition are deleted at the completion of the tracking of the reorganization.

When HALDB OLR completes, a message indicating that the reorganization has completed is sent to the System Console.

The output for this command is defined in XML and is available to automation programs that communicate with OM if the command is issued as a type-2 command.

**Related Reading:** Refer to the *IMS Version 9: Common Service Layer Guide and Reference* for sample XML for IMS commands.

**NAME()** For the type-2 version of this command, you can specify one or more HALDB partitions to HALDB OLR.

For the type-1 version of the command, you can specify only one NAME keyword for each of these commands that you issue.

**PARTNAME**

The name of a HALDB partition. Only PHDAM or PHIDAM HALDB partition names can be specified. Also, the use of wildcard character (\*), for example DBPART\* or just \*, is NOT allowed.

**SET()** Allows you to specify the RATE option as described here.

**RATE()**

Allows you to specify the RATE at which the HALDB OLR is to be run.

**rate** You can specify a rate value of 1 to 100. An online reorganization's impact on the system is affected by the available system resources, by total system utilization (including other online reorganizations), by total logging volume, by log contention, and by the intensity at which this reorganization was requested to run. These same factors also affect the speed at which the reorganization runs.

A rate value of 50 specifies that 50% of the elapsed time be spent copying records and the remaining 50% be spent in a delay. This causes the reorganization to run approximately twice as long as it would have run with a rate value of 100.

You can use the RATE parameter to control the intensity at which the reorganization runs. This can affect both the reorganization's speed and its impact on the rest of the system. The value you specify for rate is the percentage of elapsed time to be devoted to copying records. The remaining time is to be an intentionally introduced delay in the copying process that minimizes the reorganization's impact on other IMS work and on the whole system.

RATE(100) is the default.

**OPTION()** Allows the specification of the DEL/NODEL options as described here. If the partition is tracked at an RSR tracker site, the OPTION value (DEL or NODEL) in effect at the completion of the online



reorganization also determines whether the inactive data sets for the shadow partition are deleted at the completion of the tracking of the reorganization.

**DEL** Indicates the deletion of the inactive data sets is to be attempted when the online reorganization completes. The attempted deletion occurs regardless of who created the data sets or when the data sets were created. DEL is the default.

**NODEL**

Indicates that the deletion of the inactive data sets is *not* to be attempted at online reorganization completion.

**Attention:** If the HALDB OLR is stopped prior to completion, the OPTION(NODEL) if specified is not retained or remembered by IMS. OPTION(NODEL) must be specified on the INITIATE OLREORG command that is issued to resume the stopped online reorganization. OPTION(NODEL) can also be specified on the UPDATE OLREORG command.

**Command Responses for /INITIATE OLREORG**

When the /INITIATE OLREORG command is issued as a type-1 command, the command response is returned in a message format.

When the command completes successfully, a DFS0725I message with a completion code of 0 is returned to the system console and master terminal. If the command results in an error, a non-zero completion code or an error message is returned to the master terminal and system console.

If the command results in an error, an error message is returned to the entering terminal. One or more of the following messages might also be returned:

```
DFS0725I INITIATE|UPDATE|TERMINATE OLREORG COMMAND FOR DB dbnamexx COMPLETE.
CC= nn
```

where: dbnamexx is the HALDB partition name entered on the command  
nn is the completion code

```
DFS058 - INITIATE COMMAND IN PROGRESS
DFS107 - REQUIRED KEYWORD NOT PRESENT
DFS110 - COMMAND KEYWORD OLREORG INVALID FOR yyyyyyy (tracker, XRF alt)
DFS128 - TOO FEW KEYWORDS
DFS136 - COMMAND xxxxxxxxxxxx INVALID FOR yyyyyyy (tracker, XRF alt)
DFS165 - COMMAND VERB IS INVALID
DFS2026 - NO WKAP STORAGE
DFS2038 - INSUFFICIENT STORAGE AVAILABLE TO PROCESS COMMAND
DFS2262 - TOO MANY DATABASE PARAMETERS
DFS3630 - GENERIC PARAMETER NOT ALLOWED
```

For more complete information about these messages, see and the *IMS Version 9: Messages and Codes, Volume 2*.

**Command Responses for INITIATE OLREORG from OM API**

The following sections describe the responses from the OM API for the INITIATE OLREORG command.

***INITIATE OLREORG Output Fields:***

Table 6 on page 29 shows the INITIATE OLREORG output fields. The columns in the table are as follows:



|                     |   |
|---------------------|---|
| <b>Short Label</b>  | Contains the short label generated in the XML output.<br><br><b>Note:</b> This field does not apply to the /INITIATE or to the /INIT command. |
| <b>Show Keyword</b> | Identifies keyword on the command that caused the field to be generated.<br>N/A appears for output fields that are always returned.           |
| <b>Meaning</b>      | Provides a brief description of the output field.   |

Table 6. Output Fields for INITIATE OLREORG Command

| Short Label | Show Keyword | Meaning   |
|-------------|--------------|---|
| PART        | N/A          | Partition name.   |
| MBR         | N/A          | The IMS that processed the command and built the command response line. |
| CC          | N/A          | Completion code.  |

**Return, Reason, and Completion Codes for the INITIATE OLREORG Command:**

The OM return and reason codes that might be returned as a result of the INITIATE OLREORG command are standard for all commands entered through the OM API. See the *IMS Version 9: Common Service Layer Guide and Reference* for a list of the OM codes and the code meanings.

Table 7 includes the return and reason codes and a brief explanation of the INITIATE OLREORG command. Table 8 includes an explanation of the completion codes. Errors unique to the processing of INITIATE OLREORG command are returned as completion codes. A completion code is returned for each action against a HALDB partition.

Table 7. Return and Reason Code for the INITIATE OLREORG Command

| Return Code | Reason Code | Meaning  |
|-------------|-------------|--|
| X'00000000' | X'00000000' | The INITIATE OLREORG command completed successfully. |
| X'00000010' | X'0000400C' | Command issued on an XRF alternate.                  |
|             | X'00004014' | Command issued on an RSR tracker.                    |
|             | X'00004030' | Command Invalid for LSO=Y.                           |
| X'00000014' | X'00005000' | A GETMAIN error occurred.                            |

Table 8. Completion Codes for the INITIATE OLREORG Command

| Completion Code | Meaning  |
|-----------------|--|
| C' 0'           | The INITIATE OLREORG command completed successfully for the partition. |
| C' 10'          | Resource name invalid.   |
| C' 14'          | Resource not a partition name.   |
| C' 1C'          | Resource is a partitioned secondary index.                             |
| C' 20'          | HALDB OLR is already in progress.                                      |

## Examples of INITIATE OLREORG Input and Output

This section provides type-2 and type-1 input and output examples and explanations for the INITIATE OLREORG command.

### **EXAMPLE 1:**

#### **Entry ET:**

```
/INITIATE OLREORG NAME(HALDBPRT) SET(RATE(50))
```

#### **Response ET:**

```
DFS058I INITIATE COMMAND IN PROGRESS
```

#### **System Console/MTO:**

```
DFS0725I INITIATE OLREORG COMMAND FOR DB HALDBPRT COMPLETE. CC= 0
```

#### **Explanation:**

Issue the INITIATE OLREORG command to start HALDB OLR for partition HALDBPRT at the entering terminal. The DFS058I command response is returned to the entering terminal followed by the DFS0725I message to the system console and MTO.

**Related Reading:** For more information about messages, see the *IMS Version 9: Messages and Codes, Volume 2*.

### **EXAMPLE 2:**

#### **Entry ET:**

```
/INIT OLREORG NAME(PDHDOKA) SET(RATE(5))
```

#### **Response ET:**

```
DFS058I INITIATE COMMAND IN PROGRESS
```

#### **Response MTO/System console:**

```
DFS2970I - OLR STARTED FOR NAME=PDHDOKA  
DFS0725I INITIATE OLREORG COMMAND FOR DB PDHDOKA COMPLETE. CC= 0
```

**Explanation:** The INIT OLREORG command for partition PDHDOKA is issued from a terminal, resulting in the message, DFS058I INITIATE COMMAND IN PROGRESS being sent to that entering terminal.

When OLR is started successfully, the DFS0725I and the DFS2970 messages are sent to the system console and to the MTO.

### **EXAMPLE 3:**

#### **TSO SPOC Input:**

```
INIT OLREORG NAME(PDHDOKA,PDHDOKB) SET(RATE(5))
```

#### **TSO SPOC Output:**

| Partition | MbrName | CC |
|-----------|---------|----|
| PDHDOKA   | IMSA    | 0  |
| PDHDOKA   | IMS1    | C3 |
| PDHDOKB   | IMSA    | 0  |
| PDHDOKB   | IMS1    | C3 |

**OM API Input:**

```
CMD (INIT OLREORG NAME(PDHDOKA,PDHDOKB) SET(RATE(5)))
```

**OM API Output:**

```
<imsout>
<ctl>
<omname>OM10M </omname>
<omvsn>1.2.0</omvsn>
<xmlvsn>1 </xmlvsn>
<stime>2003.168 21:19:06.827362</stime>
<stotime>2003.168 21:19:07.036292</stotime>
<staseq>B99629BFEC66256E</staseq>
<stoseq>B99629C01F684F44</stoseq>
<rqsttkn1>USRT005 10141906</rqsttkn1>
<rc>00000000</rc>
<rsn>00000000</rsn>
</ctl>
<cmd>
<master>IMSA </master>
<userid>USRT005 </userid>
<verb>INIT</verb>
<kwd>OLREORG </kwd>
<input>INIT OLREORG NAME(PDHDOKA,PDHDOKB) SET(RATE(5)) </input>
</cmd>
<cmdrsphdr>
<hdr s1b1="PART" l1b1="Partition" scope="LCL" sort="A" key="1"
  scroll="NO" len="9" dtype="CHAR" align="left" />
<hdr s1b1="MBR" l1b1="MbrName" scope="LCL" sort="N" key="0" scroll="NO"
  len="8" dtype="CHAR" align="left" />
<hdr s1b1="CC" l1b1="CC" scope="LCL" sort="N" key="0" scroll="YES"
  len="4" dtype="INT" align="right" />
</cmdrsphdr>
<cmdrspdata>
<rsp> PART(PDHDOKB ) MBR(IMSA ) CC( 0) </rsp>
<rsp> PART(PDHDOKA ) MBR(IMS1 ) CC( C3) </rsp>
<rsp> PART(PDHDOKB ) MBR(IMS1 ) CC( C3) </rsp>
</cmdrspdata>
</imsout>
```

**Explanation:** The INIT OLREORG command is issued from TSO SPOC for partitions PDHDOKA and PDHDOKB. The command is routed to IMSA and IMS1. The command is successful at IMSA for both partitions. The command is not successful for both partitions at IMS1.

---

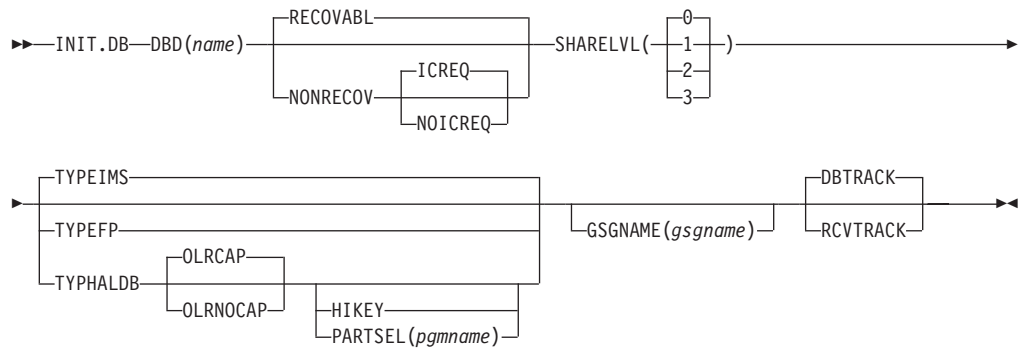
## INIT.DB

Use this command to register a database with DBRC and specify the level of database sharing and whether HALDB OLR is allowed to run on the database.

A database must be registered with DBRC before you can initialize a new DBDS, HALDB partition, or DEDB area with the INIT.DBDS command or INIT.PART command.

These are the HALDB OLR support keywords and parameters for the INIT.DB command.

## Format of INIT.DB



## INIT.DB Command Usage

**Additional parameters valid for HALDB master(type=HALDB) only:**

### OLRCAP | OLRNOCAP

Mutually exclusive, optional parameters used to indicate that HALDB OLR is allowed to run.

**OLRCAP** specifies in the master database record and in the partition database record that HALDB OLR is allowed to run. This is the default.

**OLRNOCAP** specifies in the master database record and in the partition database record that HALDB OLR is not allowed to run.

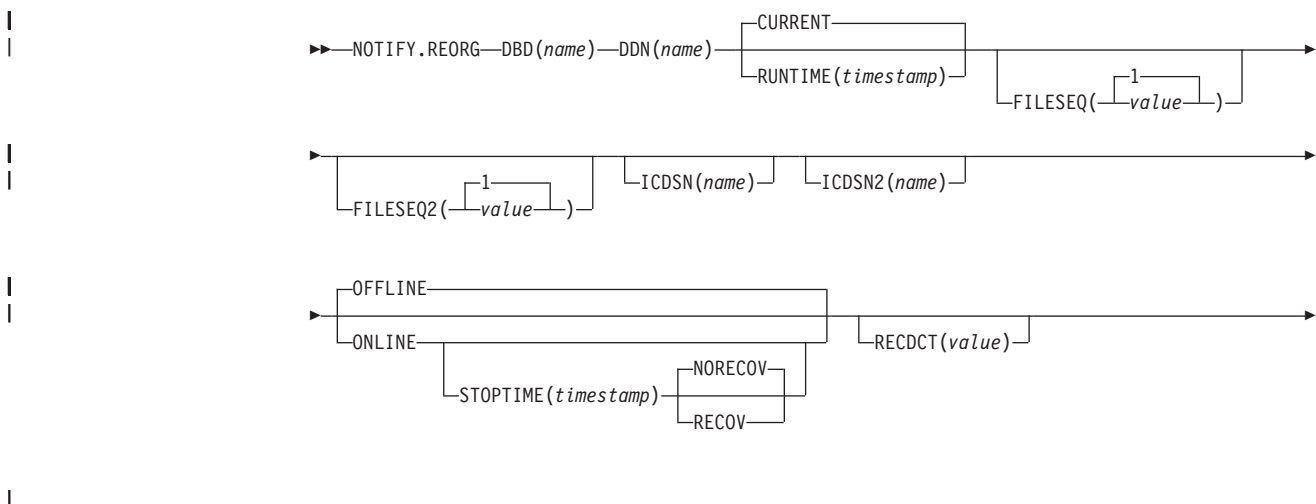
## NOTIFY.REORG

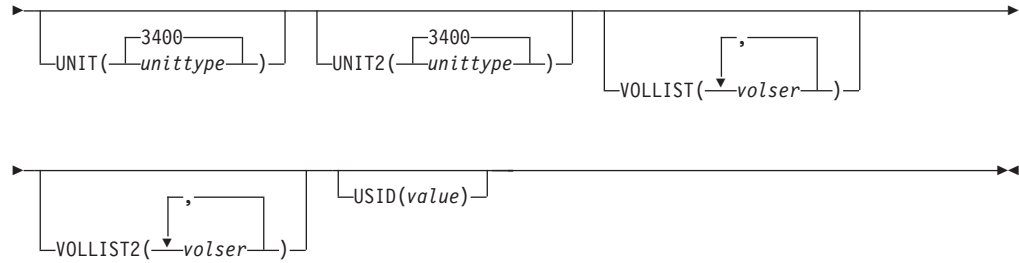
These are the HALDB OLR support keywords and parameters for the NOTIFY.REORG command.

Use a NOTIFY.REORG command to add a record to RECON about the reorganization of the database to which an identified DBDS belongs.

The information in the reorganization record is used by DBRC to determine which image copy data sets, change accumulation data sets, and log data sets are valid as input to a subsequent recovery of the identified DBDS.

## Format of NOTIFY.REORG





## NOTIFY.REORG Command Usage

### **ONLINE | OFFLINE**

Mutually exclusive, optional parameters used to specify whether the reorganization is a HALDB OLR or an offline reorganization.

**ONLINE** specifies that the reorganization is a HALDB OLR

**OFFLINE** specifies that the DBDS was reorganized offline (default)

### **RECOV | NORECOV**

Mutually exclusive, optional parameters used to specify whether the online reorganization may be used as input for recovery. This parameter is ignored unless specified with the ONLINE parameter.

**RECOV** specifies that the HALDB OLR may be used as input to recovery

**NORECOV** specifies that the HALDB OLR may not be used as input to recovery

### **STOPTIME(timestamp)**

Optional parameter used to specify the time when a HALDB OLR successfully completed. This parameter can only be specified with the ONLINE keyword.



## Chapter 5. Monitoring HALDB OLR

This section describes how IMS commands allow you to:

- Display status information about cursor active HALDB OLRs in an OM API environment.
- Monitor the status of the specified databases or partitions on the IMS where the command is issued (including those with HALDB OLRs in cursor-active status).
- Display all HALDB partitions that have HALDB OLRs in cursor-active status on the IMS locally.
- Display a list of all databases with HALDB OLR in cursor-active status.

### QUERY OLREORG Command

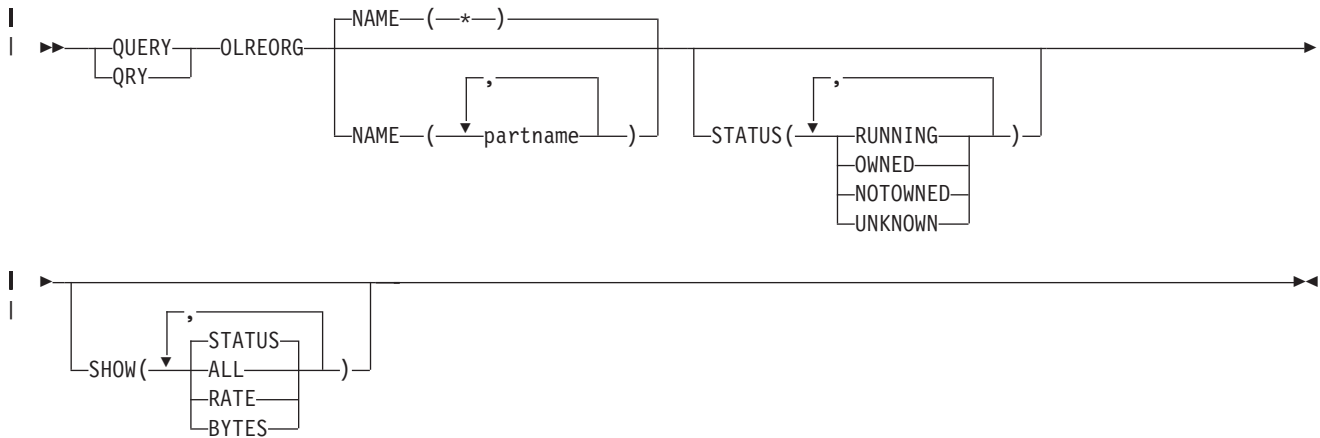
The QUERY OLREORG command returns status information about active HALDB OLRs in an OM API environment.

The /DISPLAY DB OLR command displays the status of the HALDB OLR in progress locally.

**Related Reading:** For more information on the /DISPLAY DB OLR command see the section entitled, “/DISPLAY DATABASE Command HALDB OLR Support” on page 40.

### Format of the QUERY OLREORG Command

This syntax diagram illustrates the format of the QUERY OLREORG command.



### Environments, Keywords, and Parameters of the QUERY OLREORG Command

Table 9 on page 36 lists the environments, keywords, and parameters of the QUERY OLREORG command.

This command uses syntax that can distinguish between keywords and resource names, so none of these keywords are reserved keywords. You can have resource names that are the same as the HALDB OLR keywords.

Table 9. Valid Environments for the QUERY OLREORG command Keywords, and Parameters

| Command / Keywords | DB/DC | DBCTL | DCCTL |
|--------------------|-------|-------|-------|
| OLREORG            | X     | X     |       |
| NAME               | X     | X     |       |
| STATUS             | X     | X     |       |
| SHOW               | X     | X     |       |

## Usage of the QUERY OLREORG Command

The QUERY OLREORG command is valid only from OM API. Responses from each IMS that the command was routed to are consolidated by OM. If particular names are specified on the NAME parameter, response lines are returned for each partname that was specified. For NAME(\*), response lines are returned only for the HALDB partitions that have online reorganizations in progress at each IMS.

A non-zero return code and non-zero reason code is returned when the command is routed to an XRF (Extended Restart Facility) alternate system and when the command is routed to an RSR (Remote Site Recovery) tracking system.

**Restriction:** This command does not support keyword synonyms. Only one version of each keyword can be specified for each command issued.

The output for the QUERY OLREORG command is defined in XML and is available to automation programs which communicate with OM.

**Related Reading:** See the *IMS Version 9: Common Service Layer Guide and Reference* for sample IMS command XML.

|                 |  |
|-----------------|--|
| <b>NAME()</b>   | Displays the online reorganization status for the specified HALDB partitions.  |
| <b>PARTNAME</b> | The name of a HALDB PHDAM or PHIDAM partition. The wildcard character (*) is NOT allowed except as NAME(*) for all defined HALDB partitions. NAME(*) is the default.   |
| <b>SHOW()</b>   | Specifies the output fields to return. If SHOW is not specified only the partname is returned.   |
| <b>ALL</b>      | Returns STATUS, RATE, and BYTES on each response line.   |
| <b>BYTES</b>    | Returns the number of bytes that have been moved to the output data set while this HALDB OLR was running continuously.   |
|                 | If this HALDB OLR was terminated and resumed, the value of BYTES reflects only the number of bytes moved to the output data set since the subsequent INITIATE OLREORG. The number of bytes moved before the TERMINATE OLREORG command was entered is not included in the output of a QUERY OLREORG command issued after the resumption of the HALDB OLR. See "Example of a /DISPLAY DB OLR Output with a Resume" on page 41. |



|  |                 |                 |  |
|--|-----------------|-----------------|--|
|  |                 | <b>RATE</b>     | Returns the rate from 1 to 100, at which the HALDB OLR is running.   |
|  |                 |                 |  |
|  |                 | <b>STATUS</b>   | Online reorganization status.  |
|  | <b>STATUS()</b> |                 | Displays online reorganizations that possess at least one of the specified statuses. If STATUS is not specified then any online reorganization with a status of either RUNNING, OWNED, or NOTOWNED is returned.  |
|  |                 | <b>NOTOWNED</b> | Specifies that the output is for the online reorganizations that have been temporarily stopped by the TERMINATE OLREORG command and therefore, are not owned by any IMS.   |
|  |                 | <b>OWNED</b>    | Specifies that the output is for HALDB OLRs that are owned by other IMS'. Those owned by this IMS display a STATUS of RUNNING. The OLRs that are running on the IMS where the command is being processed display RUNNING. OLRs running on other IMS subsystems displayed a STATUS of OWNED.  |
|  |                 | <b>RUNNING</b>  | Specifies the output is for the HALDB OLRs that are owned by each IMS for the specified <i>partname</i> or <i>partnames</i> . This option can be used to determine which IMS has an online reorganization running for which <i>partname</i> .  |
|  |                 | <b>UNKNOWN</b>  | Specifies that the output is for those <i>partnames</i> on each IMS where status of the HALDB OLR cannot be determined. This inability to determine the status can be due to situations such as: the HALDB master being taken offline by a /DBR DB command, or the IMS not being authorized to the named <i>partname</i> because of an IRLM failure. |

**QUERY OLREORG Output Fields**

Table 10 shows the QUERY OLREORG output fields. The columns in the table are as follows:

**Short Label**

Contains the short label generated in the XML output.

**Show Keyword**

Identifies the keyword on the command that caused the field to be generated.

N/A appears for output fields that are always returned.

**Meaning**

Provides a brief description of the output field.

*Table 10. Output Fields of QUERY OLREORG*

| Short Label | Show Keyword | Meaning                                    |
|-------------|--------------|--|
| PART        | N/A          | Partition name.                            |
| MBR         | N/A          | The IMS from which the command was issued. |
| CC          | N/A          | Completion code.                           |
| LSTT        | STATUS       | Status of HALDB OLR.                       |

Table 10. Output Fields of QUERY OLREORG (continued)

| Short Label | Show Keyword | Meaning   |
|-------------|--------------|---|
| RATE        | RATE         | The speed at which HALDB OLR runs. A value of 1 to 100 percent.                                 |
| BYTES       | BYTES        | Number of bytes moved to the output data set while the HALDB OLR has been running continuously. |

### Return, Reason, and Completion Codes for the QUERY OLREORG Command

The OM return and reason codes that might be returned as a result of the QUERY OLREORG command are standard for all commands entered through the OM API. See *IMS Version 9: Common Service Layer Guide and Reference* for a list of the OM codes and the code meanings.

Table 7 on page 29 includes the return and reason codes and a brief explanation of the QUERY OLREORG command. Table 12 includes an explanation of the completion codes. Errors unique to the processing of QUERY OLREORG command are returned as completion codes. A completion code is returned for each action against a HALDB partition.

Table 11. Return and Reason Code for the QUERY OLREORG Command

| Return Code | Reason Code | Meaning   |
|-------------|-------------|---|
| X'00000000' | X'00000000' | The QUERY OLREORG command completed successfully. |
| X'00000008' | X'00002004' | Invalid command keyword.                          |
| X'00000010' | X'00004030' | Command Invalid for LSO=Y.                        |
|             | X'00004014' | Command issued on an RSR tracker.                 |
|             | X'0000400C' | Command issued on an XRF alternate.               |
| X'00000014' | X'00005000' | A GETMAIN error occurred.                         |

Table 12. Completion Codes for the QUERY OLREORG Command

| Completion Code | Meaning   |
|-----------------|---|
| C' 0'           | The QUERY OLREORG command completed successfully for the partition. |
| C' 10'          | Resource name invalid.  |
| C' 14'          | Resource not a partition name.                                      |
| C' 1C'          | Resource is a partitioned secondary index.                          |
| C' 24'          | No HALDB OLR is in progress.  |
| C' CB'          | Partition not in specified status.                                  |

## Example of QUERY OLREORG Input and Output

### Example 1

This section provides type-2 input and output examples and explanations for the QUERY OLREORG command.

#### TSO SPOC Input:

```
QRY OLREORG NAME(*) SHOW(ALL)
```

**TSO SPOC Output:**

| Partition | MbrName | CC | LCLSTAT | RATE | BYTES  |
|-----------|---------|----|---------|------|--------|
| PDHDOKA   | IMSA    | 0  | RUNNING | 5    | 115260 |
| PDHDOKB   | IMSA    | 0  | RUNNING | 5    | 634    |

**OM API Input:**

CMD (QRY OLREORG NAME(\*) SHOW(ALL))

**OM API Output:**

```

<imsout>
<ctl>
<omname>OM10M </omname>
<omvsn>1.2.0</omvsn>
<xmlvsn>1 </xmlvsn>
<statime>2003.168 21:05:23.374319</statime>
<stotime>2003.168 21:05:23.376148</stotime>
<staseq>B99626AE9E0EF969</staseq>
<stoseq>B99626AE9E8140A4</stoseq>
<rqsttkn1>USRT005 10140523</rqsttkn1>
<rc>00000000</rc>
<rsn>00000000</rsn>
</ctl>
<cmd>
<master>IMSA </master>
<userid>USRT005 </userid>
<verb>QRY </verb>
<kwd>OLREORG </kwd>
<input>QRY OLREORG NAME(*) SHOW(ALL) </input>
</cmd>
<cmdrsphdr>
<hdr s1b1="PART" l1b1="Partition" scope="LCL" sort="A" key="1"
  scroll="NO" len="9" dtype="CHAR" align="left" />
<hdr s1b1="MBR" l1b1="MbrName" scope="LCL" sort="N" key="0" scroll="NO"
  len="8" dtype="CHAR" align="left" />
<hdr s1b1="CC" l1b1="CC" scope="LCL" sort="N" key="0" scroll="YES"
  len="4" dtype="INT" align="right" />
<hdr s1b1="LSTT" l1b1="LCLSTAT" scope="LCL" sort="NO" key="0"
  scroll="YES" len="*" dtype="CHAR" align="left" />
<hdr s1b1="RATE" l1b1="RATE" scope="LCL" sort="N" key="0" scroll="YES"
  len="4" dtype="INT" align="right" />
<hdr s1b1="BYTES" l1b1="BYTES-MOVED" scope="LCL" sort="NO" key="0"
  scroll="YES" len="20" dtype="INT" align="right" />
</cmdrsphdr>
<cmdrspdata>
<rsp> PART(PDHDOKA ) MBR(IMSA ) CC( 0) LSTT(RUNNING ) RATE( 5)
  BYTES( 115260) </rsp>
<rsp> PART(PDHDOKB ) MBR(IMSA ) CC( 0) LSTT(RUNNING ) RATE( 5)
  BYTES( 634) </rsp>
</cmdrspdata>
</imsout>

```

**Explanation:** The QUERY OLREORG command is routed to IMSA. The command is issued to obtain the information on all of the OLRs that are in progress at IMSA. The output that is returned contains the following information:

- The partition name
- The IMS ID
- The status of OLRs in progress
- The rate of OLR
- The number of bytes moved

## **/DISPLAY DATABASE Command HALDB OLR Support**

The /DISPLAY DB command output shows the status of the specified databases or partitions locally. Use the OLR attribute to show all HALDB partitions that have online reorganizations in progress locally.

### **Example of /DISPLAY DB Output with HALDB OLR**

Here is an example of output that is returned to the originating terminal from a /DISPLAY DB PVHDJJC POIHKD command.

| DATABASE | TYPE   | TOTAL UNUSED | TOTAL UNUSED | ACC | CONDITIONS |
|----------|--------|--------------|--------------|-----|------------|
| DBVHDJ05 | PHDAM  |              |              | UP  |            |
| PVHDJ5C  | PART   |              |              | UP  | OLR        |
| DBOHIDK5 | PHIDAM |              |              | UP  |            |
| POIHKD   | PART   |              |              | UP  | OLR        |

*Figure 3. /DISPLAY DB PVHDJJC POIHKD output for databases (2 with OLRs in progress)*

## **/DISPLAY DB OLR Command**

The /DISPLAY DB OLR command returns a list and status of all databases that have HALDB OLR in progress locally to the originating terminal. The /DISPLAY DB OLR information is returned with these FIDs:

- Header line with FID D77
- Database information with FID D18

### **RESTRICTIONS:**

The attribute OLR cannot be specified on the /DISPLAY DB command with any other attributes such as STOPPED or ALLOCS, or others. If additional attributes are specified, the command is rejected and message DFS0135 is displayed.

A /DIS DB OLR command issued on the RSR tracker or the XRF alternate results in the display of this version of message DFS132I - ALL PARAMETERS ARE INVALID.

Returns the number of bytes that have been moved to the output data set while this HALDB OLR was running continuously.

If this HALDB OLR was terminated and resumed, the value of BYTES reflects only the number of bytes moved to the output data set since the subsequent INITIATE OLREORG. The number of bytes moved before the TERMINATE OLREORG command was entered is not included in the output of a QUERY OLREORG command issued after the resumption of the HALDB OLR.

### **Example of /DISPLAY DB OLR Output**

Here is an example of output returned to the originating terminal from a /DISPLAY DB OLR command.

| DATABASE | PART    | RATE | BYTES        | STATUS  |
|----------|---------|------|--------------|---------|
| DBOHIDK5 | POHIDK5 | 1    | 123456789999 | RUNNING |
| DBVHDJ05 | PVHDJ5C | 50   | 0            | RUNNING |

*Figure 4. /DISPLAY DB OLR output*

### **Entry ET:**

```
| /DIS DB OLR
```

| **Response ET Output:**

```
| DATABASE PART RATE BYTES STATUS
| DBHDOK01 PDHDOKA 5 39348 RUNNING
| *03168/143937*
```

| **Explanation:** The /DIS DB OLR command is issued to obtain the information on all OLRs that are in progress. The following information is returned:

- | • Master database name
- | • Partition name
- | • OLR rate
- | • Number of bytes moved
- | • Status of the OLR
- | • Date/Time stamp

| **Example of a /DISPLAY DB OLR Output with a Resume**

| Here is an example of output returned to the originating terminal first from a /DISPLAY DB OLR command entered after an /INIT OLREORG command, then after a /TERM OLREORG and another /INIT OLREORG command.

| **Entry ET:**

```
| /INIT OLREORG NAME(CUSTMA1)
| /DIS DB OLR
```

| **Response ET Output:**

```
| DFS4444I DISPLAY FROM ID=IMS1 808
| DATABASE PART RATE BYTES STATUS
| PDCUSTMA CUSTMA1 30 36448420 RUNNING
| *03233/110628*
```

| **Entry ET:**

```
| /TERM OLREORG NAME(CUSTMA1)
| /INIT OLREORG NAME(CUSTMA1)
| /DIS DB OLR
```

| **Response ET Output:**

```
| DFS4444I DISPLAY FROM ID=IMS1 855
| DATABASE PART RATE BYTES STATUS
| PDCUSTMA CUSTMA1 100 20662176 RUNNING
| *03233/120706*
| DFS2974I - OLR COMPLETED FOR NAME=CUSTMA1 IMS1
```

| **Explanation:** The /INIT OLR is issued to start the HALDB OLR, then the /DIS DB OLR command is issued to obtain the information on all OLRs that are in progress. The following information is returned:

- | • Master database name
- | • Partition name
- | • OLR rate
- | • Number of bytes moved while this HALDB OLR has been running continuously uninterrupted.

- Status of the OLR
- Date/Time stamp

Then a /TERM OLR command, another /INIT OLR NAME(CUSTMA1), and a subsequent /DIS DB OLR command is issued showing that the number of bytes moved (in italics) to the output data set does not include the original amount from the first /INIT OLR NAME(CUSTMA1) command.

To get the total number of bytes moved for partition NAME(CUSTMA1), you would have to add the two BYTES values.

---

## QUERY DATABASE STATUS(OLR) Command

The QUERY DB command output displays HALDB OLR status. The QUERY DB STATUS(OLR) command returns a list of all databases with HALDB OLR in progress. No rate information is returned. The rate information is returned on the QUERY OLREORG command output.

**Related Reading:** For more information regarding the QUERY DB command, see the *IMS Version 9: Command Reference*.

---

## Chapter 6. Modifying and Tuning HALDB OLR

This section describes how to make these modifications and tuning adjustments:

- Stop HALDB OLRs.
- Change the intensity and impact of HALDB OLRs.
- Specifying whether to delete inactive HALDB OLR data sets when HALDB OLR processing completes.
- Modify the status of a HALDB OLR in the RECON data set.

This section also describes the effects to a HALDB master database that has been taken offline with a /DBRecovery DB command.

---

### TERMINATE OLREORG COMMAND

The TERMINATE OLREORG command is used to stop one or more HALDB OLRs in an OM API environment.

**Restriction:** This command does not support keyword synonyms. Only one version of each keyword can be specified for each command issued.

Figure 5 on page 44 shows the processing steps of a HALDB OLR where you invoke TERMINATE OLREORG to temporarily stop a HALDB OLR during copying and then invoke INITIATE OLREORG to resume the online reorganization. This figure contains columns depicting your (as the user) and IMS control flow and the statuses of the sets of data sets as processing occurs.

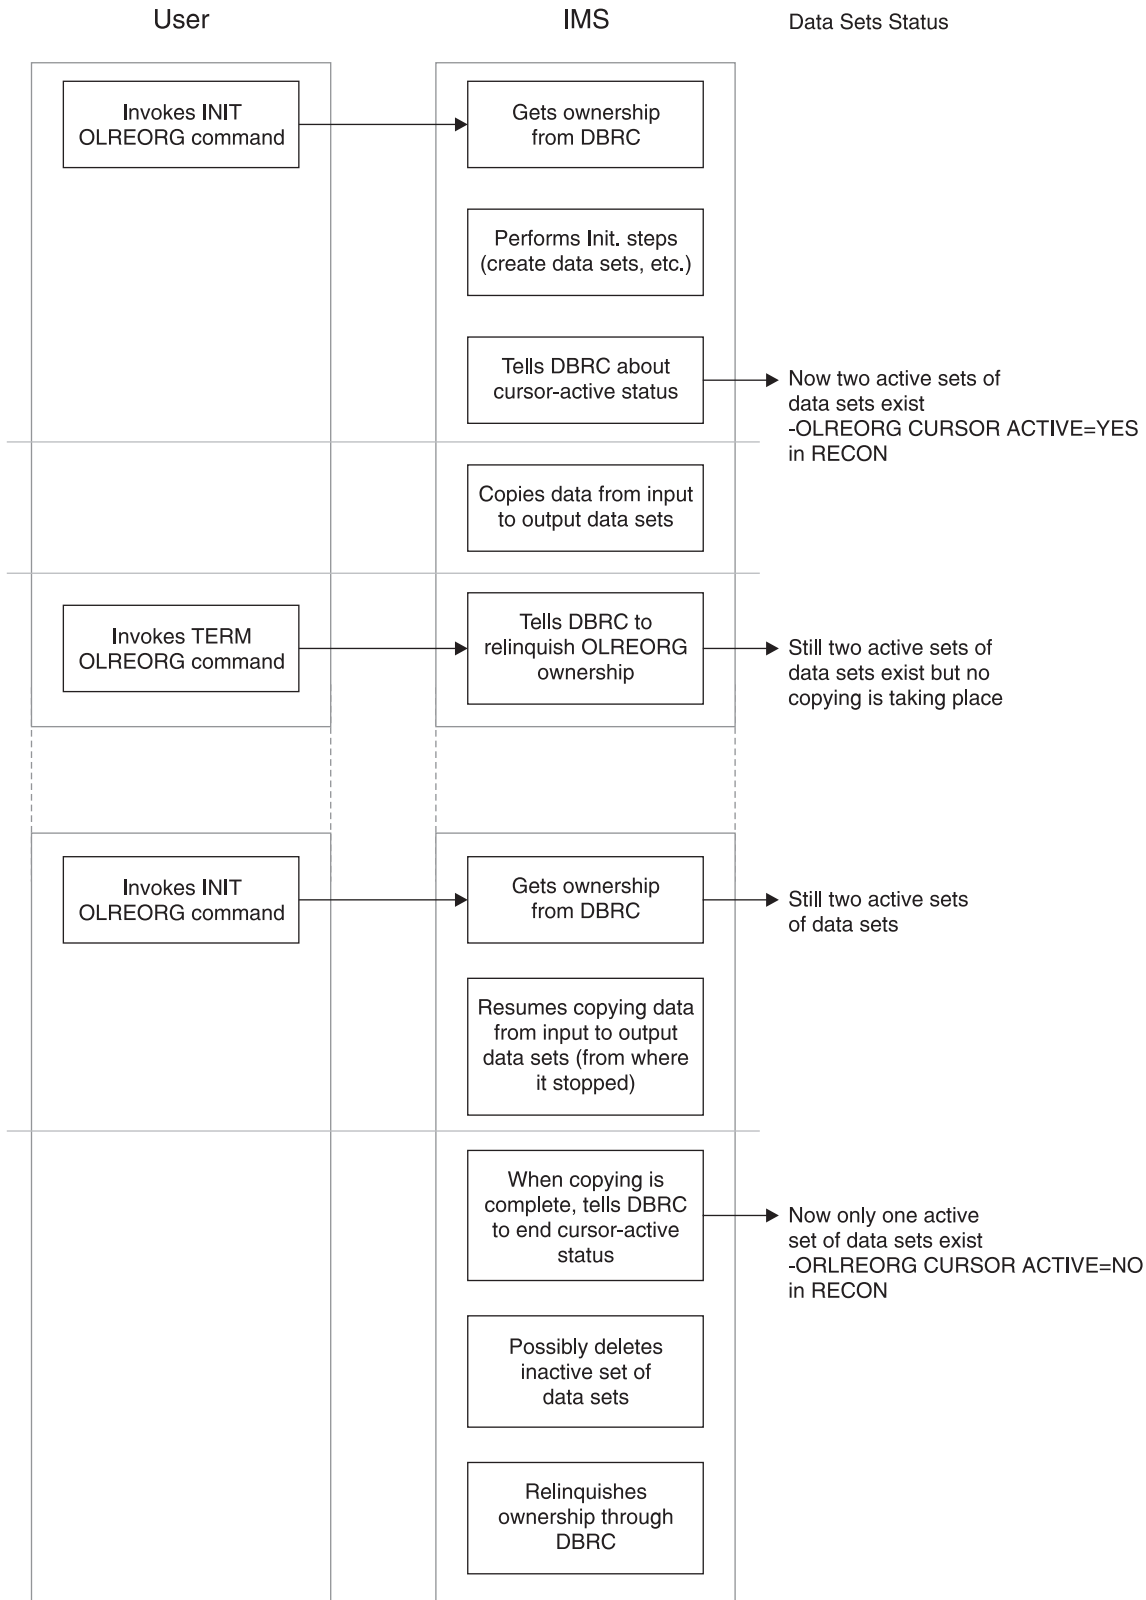


Figure 5. HALDB OLR Process with a TERMINATE OLREORG



## /TERMINATE OLREORG Command

The type-1 command syntax is the same as the type-2 command syntax issued from OM API but is preceded by a CRC (command recognition character) usually a / (slash).

The type-1 format /TERMINATE OLREORG command can be issued from terminals, AOI applications, EMCS consoles, and APPC/OTMA programs. When the command is entered as a type-1 command only one *NAME* parameter can be specified.


When this command is entered as a type-1 command, the command response is returned as a DFS0725I pre-edit message.

### Related Reading:

- See Chapter 7, "Messages and Codes of HALDB OLR," on page 59 for more information about message DFS0725I.

## Format of the TERMINATE OLREORG Command

The syntax diagram here illustrates the format of the TERMINATE OLREORG command.



## Environments, Keywords, and Parameters of the TERMINATE OLREORG Command

Table 13 lists the environments, keywords, and parameters of the TERMINATE OLREORG command.

This command uses syntax that can distinguish between keywords and resource names, so none of these keywords are reserved keywords. You can have resource names that are the same as the HALDB OLR keywords.

**Restriction:** This command does not support keyword synonyms. Only one version of each keyword can be specified for each command issued.

Table 13. Valid Environments for the TERMINATE OLREORG Command, Keywords, and Parameters

| Command / Keywords | DB/DC | DBCTL | DCCTL |
|--------------------|-------|-------|-------|
| OLREORG            | X     | X     |       |
| NAME               | X     | X     |       |
| OPTION             | X     | X     |       |

## Usage of the TERMINATE OLREORG Command

The TERMINATE OLREORG command causes the HALDB OLR to be stopped for the specified HALDB partitions. After a HALDB OLR has been terminated it is no longer active nor does it have an owning IMS. The partitions remain in OLREORG cursor-active status until the online reorganization is resumed with an INITIATE OLREORG command and completes, or an offline reorganization is run.

The output for this command is defined in XML and is available to automation programs which communicate with OM if the command is issued from an OM API environment and not as a type-1 command.

**Related Reading:** See the *IMS Version 9: Common Service Layer Guide and Reference* for sample XML for IMS commands.

- NAME()** For the type-2 version of this command, you can specify one or more HALDB partitions for which HALDB OLR is to be stopped.
- For the type-1 version of the command, you can specify only one NAME keyword for each of these commands that you issue.
- PARTNAME** The name of a HALDB partition. Only PHDAM or PHIDAM HALDB partition names can be specified. Also, the use of the wildcard character (\*), for example DBPART\*, is *not* allowed.
- \* All defined HALDB partitions.
- OPTION()** Allows you to specify the FORCE or ABORT options as described here.
- ABORT** Causes the HALDB OLR to be stopped immediately, possibly with a completion code of Abend U0474, without waiting for the current unit of reorganization to complete. Backout may be required depending on the state of the online reorganization at termination.
- FORCE** Indicates that the HALDB OLR for the named partname is to be stopped when the next record boundary is encountered. All of the moved data up to that point is committed to DASD and no backout is required.

**Attention:** If the HALDB OLR is stopped prior to completion, the OPTION(NODEL) is *not* retained and must be specified on the INITIATE OLREORG command that is issued to resume the stopped online reorganization or on the UPDATE OLREORG command.

### Command Responses for /TERMINATE OLREORG

When the /TERMINATE OLREORG command is entered as a type-1 command, the command response is returned in a message format.

When the command completes successfully, message DFS0725I with a completion code of 0 is returned to the system console and master terminal. If the command results in an error, a non-zero completion code or an error message is returned to the master terminal and system console.

If the command results in an error, this error message is returned to the entering terminal:

```
DFS0725I INITIATE|UPDATE|TERMINATE OLREORG COMMAND FOR DB dbnamexx COMPLETE.
CC= nn
      where: dbnamexx is the HALDB partition name entered on the command
            nn       is the completion code
```

For more complete information about this message, see “HALDB OLR DFS Series Messages” on page 59.

### Command Responses for TERMINATE OLREORG from OM API

The following sections describe the responses from the OM API for the TERMINATE OLREORG command.

**TERMINATE OLREORG Output Fields:** Table 14 shows the TERMINATE OLREORG output fields. The columns in the table are as follows:

**Short Label**

Contains the short label generated in the XML output.

**Note:** This field does not apply to the /TERMINATE or to the /TERM command.

**Show Keyword**

Identifies keyword on the command that caused the field to be generated.

N/A appears for output fields that are always returned.

**Meaning**

Provides a brief description of the output field.

Table 14. Output Fields for TERMINATE OLREORG Command

| Short Label | Show Keyword | Meaning                                       |
|-------------|--------------|---|
| PART        | N/A          | Partition name.                               |
| MBR         | N/A          | The IMS that built the command response line. |
| CC          | N/A          | Completion code.                              |

### Return, Reason, and Completion Codes for the TERMINATE OLREORG Command:

The OM return and reason codes that might be returned as a result of the TERMINATE OLREORG command are standard for all commands entered through the OM API. See the *IMS Version 9: Common Service Layer Guide and Reference* for a list of the OM codes and the code meanings.

Table 15 includes the return and reason codes and a brief explanation of the TERMINATE OLREORG command. Table 16 on page 48 includes an explanation of the completion codes. Errors unique to the processing of TERMINATE OLREORG command are returned as completion codes. A completion code is returned for each action against a HALDB partition.

Table 15. Return and Reason Code for the TERMINATE OLREORG Command

| Return Code | Reason Code | Meaning   |
|-------------|-------------|---|
| X'00000000' | X'00000000' | The TERMINATE OLREORG command completed successfully. |
| X'00000010' | X'00004014' | Command issued on an RSR tracker.                     |
|             | X'0000400C' | Command issued on an XRF alternate.                   |
| X'00000014' | X'00005000' | A GETMAIN error occurred.                             |

Table 16. Completion Codes for the TERMINATE OLREORG Command

| Completion Code | Meaning   |
|-----------------|---|
| C' 0'           | The TERMINATE OLREORG command completed successfully for the partition. |
| C' 10'          | Resource name invalid.  |
| C' 14'          | Resource not a partition name.  |
| C' 1C'          | Resource is a partitioned secondary index.                              |
| C' 24'          | No HALDB OLR is in progress.  |
| C' 28'          | NO DMB loaded - need to schedule.                                       |

## Examples of TERMINATE OLREORG Input and Output

This section provides type-1 and type-2 input and output examples and explanations for the TERMINATE OLREORG command.

### **EXAMPLE 1:**

#### **Entry ET:**

```
/TERM OLREORG NAME(PDHDOKA)
```

#### **Response ET:**

```
DFS0725I TERMINATE OLREORG COMMAND FOR DB PDHDOKA COMPLETE. CC= 24
```

#### **Explanation:**

The TERM OLREORG command is issued for partition PDHDOKA to stop the OLR in progress. The command is not successful because OLR is not in progress for the partition.

#### **Related Reading:**

- See for more information about message DFS0725I.

### **EXAMPLE 2:**

#### **TSO SPOC Input:**

```
TERM OLREORG NAME(PDHDOKA,PDHDOKC)
```

#### **TSO SPOC Output:**

| Partition | MbrName | CC |
|-----------|---------|----|
| PDHDOKA   | IMSA    | 0  |
| PDHDOKA   | IMS1    | 24 |
| PDHDOKC   | IMSA    | 24 |
| PDHDOKC   | IMS1    | 24 |

#### **OM API Input:**

```
CMD ( TERM OLREORG NAME(PDHDOKA,PDHDOKC))
```

#### **OM API Output:**

```
<imsout>
<ctl>
<omname>OM10M </omname>
<omvsn>1.2.0</omvsn>
<xmlvsn>1 </xmlvsn>
```

```

<statime>2003.168 21:31:13.035976</statime>
<stotime>2003.168 21:31:13.038227</stotime>
<staseq>B9962C747D6C8868</staseq>
<stoseq>B9962C747DF93586</stoseq>
<rqsttkn1>USRT005 10143113</rqsttkn1>
<rc>0200000C</rc>
<rsn>00003008</rsn>
</ctl>
<cmderr>
<mbr name="IMSA  ">
<typ>IMS  </typ>
<styp>DBCTL  </styp>
<rc>0000000C</rc>
<rsn>00003000</rsn>
</mbr>
<mbr name="IMS1  ">
<typ>IMS  </typ>
<styp>DBDC  </styp>
<rc>0000000C</rc>
<rsn>00003004</rsn>
</mbr>
</cmderr>
<cmd>
<master>IMSA  </master>
<userid>USRT005 </userid>
<verb>TERM</verb>
<kwd>OLREORG  </kwd>
<input>TERM OLREORG NAME(PDHDOKA,PDHDOKC) </input>
</cmd>
<cmdrsphdr>
<hdr s1b1="PART" l1b1="Partition" scope="LCL" sort="A" key="1"
  scroll="NO" len="9" dtype="CHAR" align="left" />
<hdr s1b1="MBR" l1b1="MbrName" scope="LCL" sort="N" key="0" scroll="NO"
  len="8" dtype="CHAR" align="left" />
<hdr s1b1="CC" l1b1="CC" scope="LCL" sort="N" key="0" scroll="YES"
  len="4" dtype="INT" align="right" />
</cmdrsphdr>
<cmdrspdata>
<rsp> PART(PDHDOKA ) MBR(IMSA  ) CC( 0) </rsp>
<rsp> PART(PDHDOKC ) MBR(IMSA  ) CC( 24) </rsp>
<rsp> PART(PDHDOKA ) MBR(IMS1  ) CC( 24) </rsp>
<rsp> PART(PDHDOKC ) MBR(IMS1  ) CC( 24) </rsp>
</cmdrspdata>
</imsout>

```

**Explanation:** The TERM OLREORG command is issued to stop the OLR for partitions PDHDOKA and PDHDOKC. The command is routed to IMSA and IMS1. The command is successful for partition PDHDOKA at IMSA where OLR is in progress. The command is not successful for PDHDOKC at IMSA because OLR is not in progress for PDHDOKC on IMSA. The command is not successful for either of the two partitions at IMS1 because OLR is not in progress for PDHDOKA and PDHDOKC on IMS1.

---

## UPDATE OLREORG Command

The UPDATE OLREORG command allows you to change the intensity and impact of HALDB OLRs with the RATE option. UPDATE OLREORG also allows you to specify whether to delete inactive HALDB OLR data sets when HALDB OLR processing completes with the DEL or NODEL options.

Figure 6 on page 50 shows the processing steps of a HALDB OLR where you invoke UPDATE OLREORG to change the RATE parameter and set the NODEL option to ON during copying. This figure contains columns depicting the flow of control

|  
|  
|

between you (as the user) and IMS and the statuses of the sets of data sets as processing events occur.

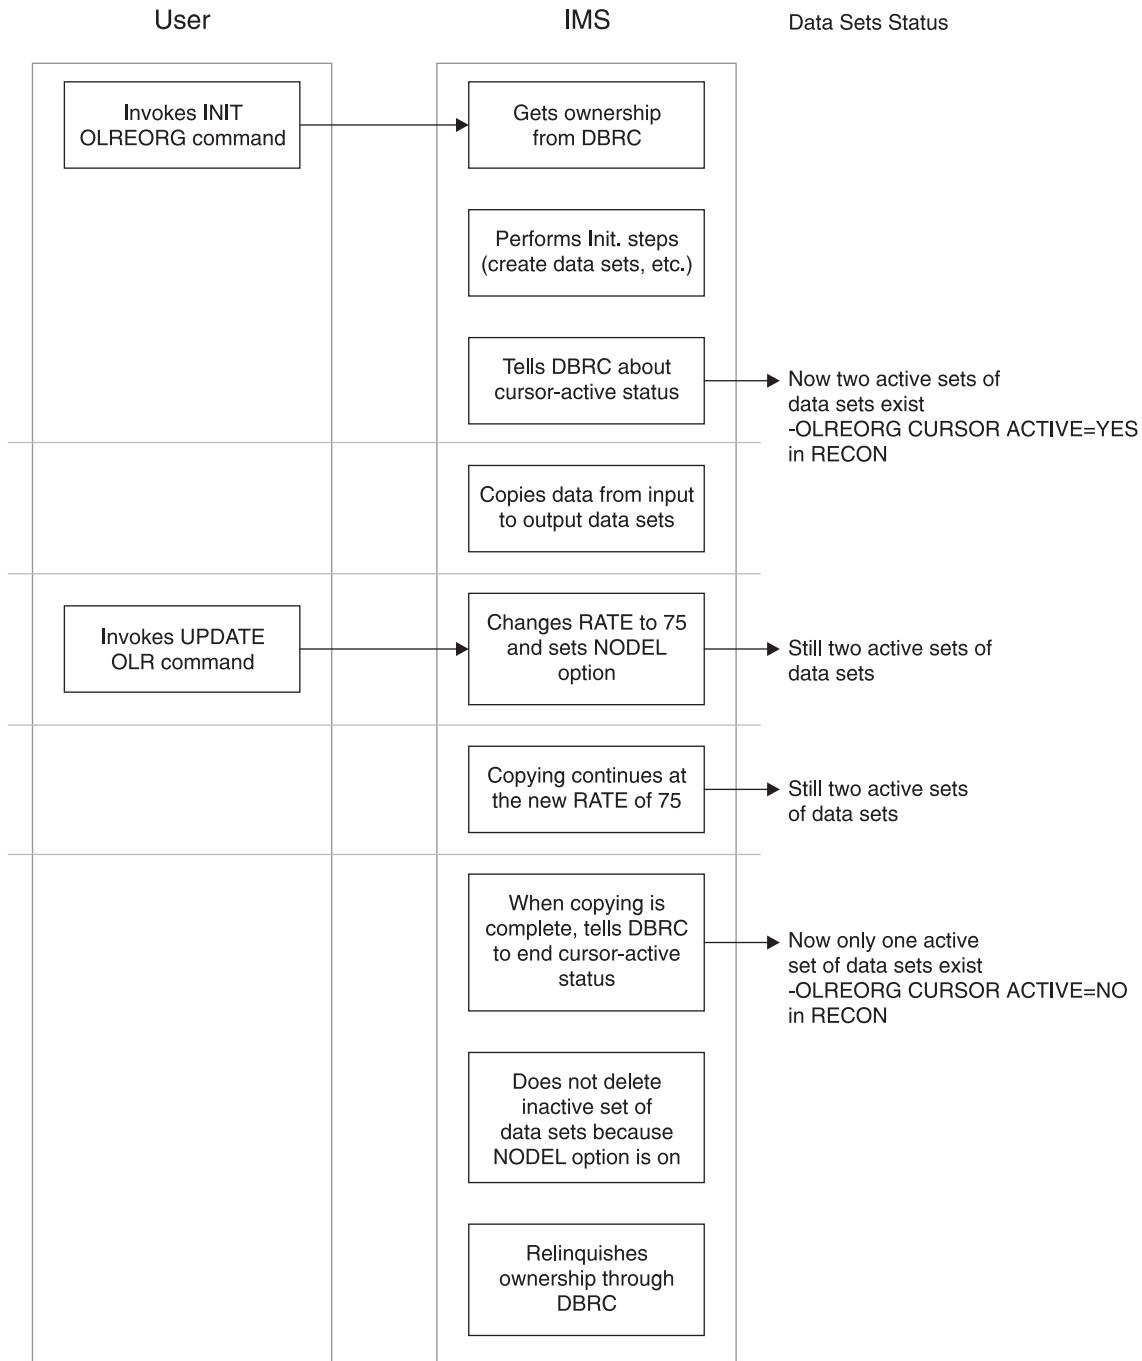


Figure 6. HALDB OLR Process with an UPDATE OLREORG

## /UPDATE OLREORG Command

The type-1 command syntax is the same as the type-2 command syntax issued from OM API but is preceded by a CRC (command recognition character) usually a / (slash).



*Online Reorganization Guide and Reference* for performance considerations related to changing the rate at which the online reorganization is run.

The output for this type-2 format command is defined in XML and is available to automation programs that communicate with OM.

**Related Reading:** See the *IMS Version 9: Common Service Layer Guide and Reference* for sample IMS command XML.

**NAME()** Allows you to specify one or more HALDB partitions for the UPDATE OLRORG type-2 command.

For the type-1 version of the command, you can specify only one NAME keyword for each of these commands that you issue.

**PARTNAME** The name of a HALDB partition. Only PHDAM or PHIDAM HALDB partition names can be specified. Also, the use of the wildcard character (\*), for example DBPART\*, is *not* allowed for PARTNAME.

\* All defined HALDB partitions.

**SET()** Allows you to specify the RATE option as described below.

**RATE()**

Allows you to specify the RATE at which the HALDB OLR is to be run.

**rate** You can specify a rate value of 1 to 100. An online reorganization's impact on the system is affected by the available system resources, by total system utilization (including other online reorganizations), by total logging volume, by log contention, and by the intensity at which this reorganization was requested to run. These same factors also affect the speed at which the reorganization runs.

A rate value of 50 specifies that 50% of the elapsed time be spent copying records and the remaining 50% be spent in a delay. This causes the reorganization to run approximately twice as long as it would have run with a rate value of 100.

You can use the RATE parameter to control the intensity at which the reorganization runs. This can affect both the reorganization's speed and its impact on the rest of the system. The value you specify for rate is the percentage of elapsed time to be devoted to copying records. The remaining time is to be an intentionally introduced delay in the copying process that minimizes the reorganization's impact on other IMS work and on the whole system.

RATE(100) is the default.

**OPTION()** Allows the specification of the DEL or NODEL options as described below. If the partition is tracked at an RSR tracker site, the OPTION value (DEL or NODEL) in effect at the completion of the HALDB



OLR also determines whether the inactive data sets for the shadow partition are deleted at the completion of the tracking of the reorganization.

**Note:** If the HALDB Online Reorganization is stopped prior to completion the DEL or NODEL keyword is *not* remembered and will need to be specified on the INITIATE OLREORG command that is issued to resume the stopped HALDB Online Reorganization.

**DEL** Indicates the deletion of the inactive data sets is to be attempted when the online reorganization completes. The attempted deletion occurs regardless of who created the data sets or when the data sets were created.

**NODEL** Indicates the deletion of the inactive data sets is not to be attempted when the online reorganization completes.

### Command Responses for /UPDATE OLREORG

When the /UPDATE OLREORG command is entered as a type-1 command, the command response is returned in a message format.

When the command completes successfully, message DFS0725I with a completion code of 0 is returned to the system console and master terminal. If the command results in an error, a non-zero completion code or an error message is returned to the master terminal and system console.

If the command results in an error, this error message is returned to the entering terminal:

```
DFS0725I INITIATE|UPDATE|TERMINATE OLREORG COMMAND FOR DB dbnamexx COMPLETE.
CC= nn
  where: dbnamexx is the HALDB partition name entered on the command
         nn       is the completion code
```

For more complete information about this message, see .

### Command Responses for UPDATE OLREORG from OM API

The following sections describe the responses from the OM API for the UPDATE OLREORG command.

#### UPDATE OLREORG Output Fields:

Table 18 shows the UPDATE OLREORG output fields. The columns in the table are as follows:

**Short Label** Contains the short label generated in the XML output.

**Show Keyword**

Identifies keyword on the command that caused the field to be generated.

N/A appears for output fields that are always returned.

**Meaning** Provides a brief description of the output field.

Table 18. Output Fields for UPDATE OLREORG Command

| Short Label | Show Keyword | Meaning         |
|-------------|--------------|-----------------|
| PART        | N/A          | Partition name. |

Table 18. Output Fields for UPDATE OLREORG Command (continued)

| Short Label | Show Keyword | Meaning                                       |
|-------------|--------------|---|
| MBR         | N/A          | The IMS that built the command response line. |
| CC          | N/A          | Completion code.                              |

### **Return, Reason, and Completion Codes for the UPDATE OLREORG Command:**

The OM return and reason codes that might be returned as a result of the UPDATE OLREORG command are standard for all commands entered through the OM API. See *IMS Version 9: Common Service Layer Guide and Reference* for a list of the OM codes and the code meanings.

Table 7 on page 29 includes the return and reason codes and a brief explanation of the UPDATE OLREORG command. Table 20 includes an explanation of the completion codes. Errors unique to the processing of UPDATE OLREORG command are returned as completion codes. A completion code is returned for each action attempted on a HALDB partition.

Table 19. Return and Reason Code for the UPDATE OLREORG Command

| Return Code | Reason Code | Meaning  |
|-------------|-------------|--|
| X'00000000' | X'00000000' | The UPDATE OLREORG command completed successfully. |
| X'00000008' | X'00002008' | Insufficient number of keywords specified.         |
| X'00000010' | X'0000400C' | Command issued on an XRF alternate.                |
|             | X'00004014' | Command issued on an RSR tracker.                  |
| X'00000014' | X'00005000' | A GETMAIN error occurred.                          |

Table 20. Completion Codes for the UPDATE OLREORG Command

| Completion Code | Meaning  |
|-----------------|--|
| C' 0'           | The UPDATE OLREORG command completed successfully for the partition. |
| C' 10'          | Resource name invalid.   |
| C' 14'          | Resource not a partition name.                                       |
| C' 1C'          | Resource is a partitioned secondary index.                           |
| C' 24'          | No HALDB OLR is in progress.   |

## **Examples of UPDATE OLREORG Input and Output**

This section provides type-1 and type-2 input and output examples and explanations for the UPDATE OLREORG command.

### **EXAMPLE 1:**

#### **Entry ET:**

```
/UPD OLREORG NAME(PDHDOKA) SET(RATE(25))
```

#### **Response ET:**

```
DFS0725I UPDATE OLREORG COMMAND FOR DB PDHDOKA COMPLETE. CC= 0
```

**Explanation:**

The UPDATE OLREORG command is issued for partition PDHDOKA to change the OLR rate to 25. The command is successful as indicated in the message DSF0725 command response.

**EXAMPLE 2:****TSO SPOC Input:**

```
UPD OLREORG NAME(PDHDOKA,PDHDOKB) SET(RATE(25))
```

**TSO SPOC Output:**

| Partition | MbrName | CC |
|-----------|---------|----|
| PDHDOKA   | IMSA    | 0  |
| PDHDOKA   | IMS1    | 24 |
| PDHDOKB   | IMSA    | 0  |
| PDHDOKB   | IMS1    | 24 |

**OM API Input:**

```
CMD (UPD OLREORG NAME(PDHDOKA,PDHDOKB) SET(RATE(25)))
```

**OM API Output:**

```
<imsout>
<ctl>
<omname>OM10M </omname>
<omvsn>1.2.0</omvsn>
<xmlvsn>1 </xmlvsn>
<statime>2003.168 21:17:57.712194</statime>
<stotime>2003.168 21:17:57.713062</stotime>
<staseq>B996297E02942007</staseq>
<stoseq>B996297E02CA6487</stoseq>
<rqsttkn1>USRT005 10141757</rqsttkn1>
<rc>0200000C</rc>
<rsn>00003000</rsn>
</ctl>
<cmderr>
<mbr name="IMS1 ">
<typ>IMS </typ>
<styp>DBDC </styp>
<rc>0000000C</rc>
<rsn>00003004</rsn>
</mbr>
</cmderr>
<cmd>
<master>IMSA </master>
<userid>USRT005 </userid>
<verb>UPD </verb>
<kwd>OLREORG </kwd>
<input>UPD OLREORG NAME(PDHDOKA,PDHDOKB) SET(RATE(25)) </input>
</cmd>
<cmdrsphdr>
<hdr s1b1="PART" l1b1="Partition" scope="LCL" sort="A" key="1"
  scroll="NO" len="9" dtype="CHAR" align="left" />
<hdr s1b1="MBR" l1b1="MbrName" scope="LCL" sort="N" key="0" scroll="NO"
  len="8" dtype="CHAR" align="left" />
<hdr s1b1="CC" l1b1="CC" scope="LCL" sort="N" key="0" scroll="YES"
  len="4" dtype="INT" align="right" />
</cmdrsphdr>
<cmdrspdata>
<rsp> PART(PDHDOKA ) MBR(IMSA ) CC( 0) </rsp>
<rsp> PART(PDHDOKB ) MBR(IMSA ) CC( 0) </rsp>
```

```

<rsp> PART(PDHDOKA ) MBR(IMS1 ) CC( 24) </rsp>
<rsp> PART(PDHDOKB ) MBR(IMS1 ) CC( 24) </rsp>
</cmdrspdata>
</imsout>

```

### Explanation:

The UPDATE OLREORG command is issued for partitions PDHDOKA and PDHDOKB to update the OLR rate to 25. The command is successful at IMSA where OLR is in progress and not at IMS1 where OLR is not in progress. A completion code of 24 is returned in the IMS1 response.

---

## /START DB Command HALDB OLR Support

See the *IMS Version 9: Command Reference* for information on the /START DB command for HALDB Online Reorganization.

---

## /DBRecovery Command HALDB OLR Support

See the *IMS Version 9: Command Reference* for information on the /DBRecovery DB command for HALDB Online Reorganization.

**Related Reading:** Also see the step-by-step recovery process in the section entitled, "Database Recovery - DFSURDB0" on page 20..

---

## /RECover Command HALDB OLR Support

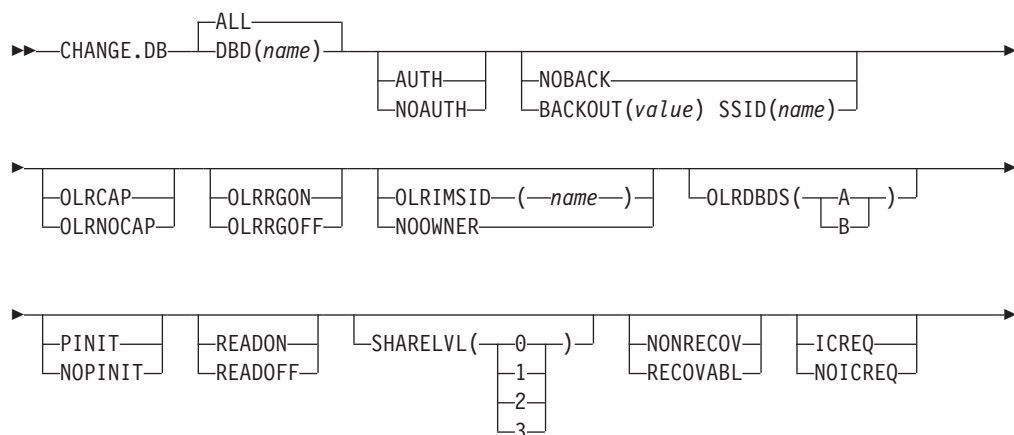
See the *IMS Version 9: Command Reference* for information on the /RECover command for HALDB Online Reorganization.

---

## CHANGE.DB

These are the HALDB OLR support keywords and parameters for the CHANGE.DB command.

### Format of CHANGE.DB





reorganized online using HALDB OLR. There is no default and these parameters are valid only when the DBD is a HALDB master.

**OLRCAP** Specifies in the master database record and the partition database record that HALDB OLR is allowed to run.

**OLRNOCAP** Specifies in the master database record and the partition database record that HALDB OLR is not allowed to run. The command fails if HALDB OLR cursor is active or if any M-through-V and Y DBDSs are active. In this case the M-through-V and Y DBDSs and all of their associated records are deleted from RECON for all partitions.

|  
|  
|  
|  
|  
|  
|

## Chapter 7. Messages and Codes of HALDB OLR

These are the IMS messages that pertain to HALDB OLR.

They also appear in the *IMS Version 9: Messages and Codes, Volume 2* and in the *IMS Version 9: Messages and Codes, Volume 1*.

### In This Section:

- “HALDB OLR DFS Series Messages”
- “HALDB OLR DSP Series Messages” on page 74
- “HALDB OLR Abends and Codes” on page 76

### HALDB OLR DFS Series Messages

These are the IMS DFS series messages that pertain to HALDB OLR.

**DFS047A UNABLE TO OBTAIN AUTHORIZATION FOR DATA BASE xxxxxxxx. REASON CODE = zz. PSB= psbname.**

**Explanation:** IMS was unable to obtain database authorization for the named database which appears in one of the DB PCBs contained in the named PSB during PSB scheduling. The PSB name only appears in the message text if the IMS is an online control region.

| Code (Hex) | Meaning |
|------------|---------|
|------------|---------|

|    |   |
|----|---|
| 01 | The requested state and the current authorization state (see RECON data set) are incompatible. The database is authorized by another active or abnormally terminated IMS subsystem, and its authorization state is incompatible with the current authorization request. |
|----|---|

|    |   |
|----|---|
| 02 | The database is not registered in the RECON data set. |
|----|---|

|    |  |
|----|--|
| 03 | The database is marked as prohibiting further authorization for one of the following reasons: a global /DBRecovery, a global /STOP, an UPDATE DB STOP, or a CHANGE.DB(NOAUTH) command. |
|----|--|

|    |  |
|----|--|
| 04 | The database is authorized only if the requested state is "READ" or "READ-GO" because of a global /DBDUMP or an UPDATE DB STOP(UPDATES) command. |
|----|--|

|    |  |
|----|--|
| 05 | The database is marked as needing an image copy. |
|----|--|

|    |   |
|----|---|
| 06 | The database is marked as needing recovery. |
|----|---|

|    |  |
|----|--|
| 07 | The database is marked as needing backout. In a data sharing |
|----|--|

environment, you can use a LIST.DB command to determine which subsystems require back outs for the named database.

For sharing within a host, the requesting subsystem is not within the authorized host.

One or more subsystems for which the database has been previously authorized are supposed to be active, but the subsystems are not known to IRLM. IRLM=N or another IRLM name might have been specified on the subsystem.

The database has been previously authorized to the subsystem.

An invalid parameter is found during the database usage compatibility evaluation process. Possibly, the database record in the RECON data set is bad.

The current authorization state in DBRC is invalid because of an unauthorization error.

An error occurred in DBRC while trying to process an authorization request.

An error occurred in DBRC while trying to process a change authorization request.

The database is already authorized to an active IMS subsystem.

The area is not marked as needing recovery for the DEDB initialization utility.

The data set name of the database

|    |  |  |    |  |
|----|--|--|----|--|
|    | data set passed to DBRC does not match the data set name in the RECON data set.  |  |    | until a high key is defined. A key is required because the HALDB master does not use a Partition Selection Routine.  |
| 20 | Database authorization failed because the database was not registered with DBRC and the DBRC RECON data sets were initialized with FORCER, which requires that all databases must be registered. |  | 36 | Image copy not allowed during online reorganization processing.  |
|    | This error might be the result of a mismatch of data set names in the RECON data set and in the dynamic allocation blocks (DFSMDA members).  |  | 37 | Loading into an M-through-V DBDS of a partition database is not allowed.   |
| 22 | The database is currently authorized to another subsystem.   |  | 38 | Offline reorganization is not allowed when online reorganization is active and the online reorganization is owned by an IMS subsystem.   |
| 23 | The database is a recovery-level tracking database.  |  |    | <b>System Action:</b> A batch region (DL/I or DBB region) will terminate with abend 0047. The database is stopped. Dependent regions continue processing, but programs cannot access the database. An attempt to access the database will result in either a 'BA' status code, if the program has issued the INIT call, or in abend U3303, if the program has not issued the INIT call.  |
| 24 | There is gap in the log stream received from the RSR active site that can contain update records for the database or area.   |  |    | <b>Programmer Response:</b> Reason codes 01 through 09 and 15 through 19, indicating a temporary authorization failure, can be corrected by one of the following actions:  |
| 25 | All log records have been applied to the tracking database for the authorizing stream and authorization is not for an OFR stream.  |  |    | <ul style="list-style-type: none"> <li>• Wait for a currently authorized subsystem to terminate or enter a /DBRecovery UPDATE DB STOP(ACCESS) to unauthorize the database from the currently authorized subsystems. Batch backout or an emergency restart might be required for an abnormally terminated subsystem. (RC 01, 08 or 15)</li> <li>• Verify that the correct log is input to the batch backout utility if DBRC=C is being used. (RC 01)</li> <li>• Enter a /START DB or UPDATE DB START(ACCESS), a /DBDUMP or UPDATE DB START(UPDATES), or a CHANGE.DB (auth) command. (RC 03)</li> <li>• Enter a /START DB or UPDATE DB START(ACCESS) command. (RC 04)</li> <li>• Run the appropriate IMS utility. (RC 05, 06 or 07)</li> <li>• Run this IMS subsystem on the host where the database is currently authorized. (RC 08)</li> <li>• Wait for the failed subsystems to complete the emergency restart or database backout. (RC 09)</li> <li>• Use the DBRC command CHANGE.SUBSYS to indicate that an IMS subsystem is no longer active. (RC 15)</li> <li>• Use the DBRC CHANGE.DBDS command to mark an area as needing recovery. (RC 16)</li> <li>• Use the DBRC INIT.DBDS command to register a DBDS or DEDB area to DBRC. (RC 18, 19)</li> <li>• Check the RECON data set to determine the correct authorization state.</li> </ul> |
| 26 | Records from an earlier update set ID have been received from the RSR active site, but have not been applied to the database.  |  |    |  |
| 27 | The database does not belong to the invoker's global service group (GSG).  |  |    |  |
| 28 | The update set ID (USID) is earlier than the USID of the most recently received image copy.  |  |    |  |
| 29 | One or more tracking database data sets have not been received.  |  |    |  |
| 30 | The database or area is covered but the requesting subsystem is not signed on to the covering GSG.   |  |    |  |
| 31 | The definition of the HALDB has changed. The caller must get the new definition before authorization can be granted.   |  |    |  |
| 32 | The HALDB needs to be initialized.   |  |    |  |
| 33 | An attempt has been made to authorize the HALDB master. Authorization can only be requested at the partition level.  |  |    |  |
| 34 | Partition Initialization is requesting authorization for a partition that does not need to be initialized.   |  |    |  |
| 35 | The DB partition cannot be authorized  |  |    |  |
|    |  |  |    | Reason codes 10 through 14 indicate that a system error has occurred. If these problems persist, contact IBM for problem determination assistance.   |



For reason code 33, check the DBDLIB level of the databases to make sure they match the RECON level.

**Problem Determination:** 35

**Module:** DFSDBAU0

---

**DFS132 ALL DATABASE AREA PARAMETERS ARE INVALID**

**Explanation:** This message is issued for one of the following reasons:

- The database or area names entered are invalid.
- A parameter is valid, but has failed processing (for example, an allocation error has occurred on the database).
- A valid parameter was specified, but the particular type of IMS you are running does not support that parameter. For example, the /DISPLAY AREA or DB command with the OFR keyword is not valid on an RSR active subsystem.

**System Action:** The command is ignored.

**System Programmer Response:** Reenter the correct database or area names that have been defined to IMS. If the parameter is valid, correct or investigate other error messages and reenter the command. Make sure any command keywords or parameters are valid for the flavor of IMS you are running.

**Module:** DBFCRAR0, DFSDBDR0, DFSICL20, DFSIC430, DFSIDPJ0, DFSIDP50

---

**DFS135 INVALID COMMAND ATTRIBUTE COMBINATION**

**Explanation:** One or more command attributes specified are invalid with another specified.

**System Action:** The command is ignored.

**Operator Response:** Reenter the command with the proper attribute combination.

**Module:** DFSIDP50

---

**DFS0488I DBR | DBD | STA | STO | LOCK | UNLOCK COMMAND COMPLETED. DBN=*database name* RC=*nn*.**

---

**DFS0488I DBR | DBD | STA | STO | VUN COMMAND COMPLETED. DBN | AREA | ADS=*database name* | *areaname* | *adsname* RC=*nn***

---

**DFS0488I UNLOCK COMMAND COMPLETED. RC=*nn***

---

**DFS0488I DBR | DBD | STA | STO COMMAND COMPLETED. KEYWORD ALL**

---

**DFS0488I DBR | DBD | STA | STO COMMAND COMPLETED. DATAGROUP (*groupname*)**

**Explanation:** The /DBRecovery, /DBDUMP, /START, or /STOP command processing is completed. A zero return code indicates that the named database has been successfully processed. The ADS specification is effective only for the /STOP command. The area is only for the /DBRecovery, /START, and /STOP commands. A nonzero return code indicates that an unexpected error was encountered while processing.

If the datagroup contains both full function and fast path databases, you might see a message indicating that the command is complete before seeing any messages indicating fast path areas did not start successfully.

The /VUNLOAD command completed processing for area xxxxxxxx. DFS0488I is only issued for areas for which the /VUNLOAD command failed, along with the reason code. The reason code indicates the result of the command. A final DFS0488I message is issued to indicate the command is completed.

For the second form of the message, the return code *nn* can have the value 00, 04, or 20.

The third and fourth forms of the message are issued to avoid excessive individual messages for successful operations.

When the keyword ALL is specified, one DFS0488I is issued for every database name, area name, or ads name in error, followed by a final DFS0488I DBR|DBD|STA|STO|VUN COMMAND COMPLETED message. DFS0488I is not issued for each database name, area name, or ADS name with return code of 00, (RC=00 indicates that the command was successful).

**Restriction:** This message will only be received for the LOCK command issued from an OM API.

**Return Code (Dec)**

|           | <b>Meaning</b>  |
|-----------|---|
| <b>01</b> | The requested state and the current authorization state (see the RECON data set) are incompatible. The database is authorized by another active or abnormally terminated IMS subsystem, and its authorization state is incompatible with the current authorization request. |
| <b>02</b> | The global command was not able to synchronize with all data sharing partners. The command was completed on the subsystem that received this message, but the   |

|           |   |  |           |   |
|-----------|---|--|-----------|---|
|           | partners need to be interrogated to determine if the command completed successfully.  |  |           | VUNLOAD lock for an area in a coupling facility structure failed. The /VUNLOAD, /STOP AREA GLOBAL, /DBRecovery AREA GLOBAL, or /STOP ADS command is terminated. The VUNLOAD, STOP or DBR processing is not done.  |
| <b>03</b> | The database is marked as prohibiting further authorization for one of the following reasons; a global /DBRrecovery, a global /ST0, or a DBRC CHANGE.DB(NOAUTH) command.  |  |           |   |
| <b>04</b> | An error occurred at the close of the database. For detailed information about the reason codes, refer to message DFS730I.<br><br>Module DFSMDA00 can also issue return code 4 if dynamic allocation or deallocation failed. This indicates an interface problem between another IMS module and the IMS dynamic allocation module, DFSMDA00. This can also indicate a missing or bad DFSMDA00 member.<br><br>For a /VUNLOAD command, the area specified was not found in any data space or coupling facility structure. Verify that the area name is correct. Issue a /DISPLAY FPVIRTUAL command to find which areas are in data spaces or coupling facility structures.<br><br>For the second form of the message, this return code indicates that the /UNLOCK SYSTEM command was only partially successful. At least one tolerated block or CI could not be purged. The /UNLOCK SYSTEM command should be entered again at a later time.<br><br>Return code 04 can also mean the database is authorized only if the requested state is "READ" or "READ-GO" because of a global /DBD command. |  | <b>11</b> | An invalid parameter was found during the evaluation process for database-usage compatibility. The database record in the RECON data set might be invalid.  |
|           |   |  | <b>12</b> | Dynamic allocation or deallocation failed because no storage was available from DFSMDA00.<br><br>Return code 12 can also mean that the current authorization state in DBRC is invalid because of an unauthorization error.  |
|           |   |  | <b>14</b> | The named database cannot be started because of one of the following possible conditions: <ul style="list-style-type: none"> <li>• DBRC has already authorized other READ access users in the other IMS subsystem to access the database. The granted authorization level for the other subsystem is incompatible with the original authorization level for the current subsystem because of the last /DBDUMP command issued. As a consequence, the original authorization level cannot be restored for the current subsystem.</li> <li>• The changed database authorization level has been requested through the /STA DB ACCESS command. The request for the change in the ACCESS parameter failed because of the incompatibility of the database with the authorization level currently held in the other IMS subsystem.</li> </ul> |
| <b>05</b> | The database is marked as needing an image copy.  |  |           |   |
| <b>06</b> | The database is marked as needing recovery.   |  |           |   |
| <b>07</b> | The database is marked as needing backout.  |  |           |   |
| <b>08</b> | Dynamic allocation failed (from DFSMDA00). This indicates that the IMS dynamic allocation module (DFSMDA00) received a nonzero return code from MVS dynamic allocation.   |  | <b>15</b> | The duplicate AREA entry on the /DBRecovery AREA command is ignored. The area was already unauthorized by DBRC.   |
|           | For a /VUNLOAD, /STOP AREA GLOBAL, /DBRecovery AREA GLOBAL, or /STOP ADS command, an attempt to get the   |  | <b>16</b> | Dynamic allocation or deallocation failed because of a system error. For detailed information about the reason codes, refer to message DFS2498I.  |

|    |  |  |  |  |
|----|--|--|--|--|
| 20 | <p>For the /UNLOCK SYSTEM command, one of the following conditions occurred:</p> <ul style="list-style-type: none"> <li>• There were no EEQEs to purge.</li> <li>• The /UNLOCK command was entered on an IMS system that is not XRF-capable.</li> <li>• The /UNLOCK command was entered on an active system that was not a previous alternate system. This command must be entered only on the alternate system.</li> <li>• The /UNLOCK command was entered on the alternate system before the alternate system completed the XRF takeover. This command must be entered on the alternate system only after one or more of the following events have occurred: <ul style="list-style-type: none"> <li>– the failing active system has terminated</li> <li>– the processor containing the failing active system has been reset</li> <li>– I/O prevention on the active system is complete</li> </ul> </li> </ul>  |  |  | <p>command was not able to delete the failed persistent connection. In this case, the following messages are generated:</p> <ul style="list-style-type: none"> <li>– DFS2783A UNABLE TO DELETE FAILED-PERSISTENT CONNECTION TO STR:<br/>SSSSSSSSSSSSSSSSSS<br/>REASON=RRRR</li> <li>– DFS0488I STA COMMAND COMPLETED. AREA=XXXXXXXX RC=21</li> </ul>   |
| 21 | <p>One of the following has occurred:</p> <ul style="list-style-type: none"> <li>• A VLOCK request cannot be granted when processing a /START AREA, or a /VUNLOAD AREA command. The area is marked as STOPPED. One sharing IMS system is holding the retained VLOCK exclusive.</li> <li>• While processing either a /VUNLOAD AREA, or a /START AREA command, an IMS sharing partner noted an inconsistency in the VSO options between the DEDB area control block (DMAC) on DASD and the DMAC in storage. The area is marked as STOPPED.<br/><br/>Issue a /START AREA command to correct the inconsistency.</li> <li>• During /VUNLOAD, or a /START AREA command processing, an IMS sharing partner noted an inconsistency in VSO options between the DMAC on DASD and the DMAC in storage. The AREA is stopped. Issue a /START AREA command to correct the inconsistency.</li> <li>• During /START AREA command processing, an IXLFORCE command is issued to delete any failed persistent connections from a prior failed system. The IXLFORCE</li> </ul> |  |  | <p>A command with the ALL keyword is not processed for a Fast Path MSDB. If /START DB ALL access is entered, module DFSDBDR0 does not include MSDBs.</p> <p>The named database cannot be started because of the following conditions:</p> <ul style="list-style-type: none"> <li>• The changed database authorization level to a data entry database (DEDB) has been requested through the /START DB ACCESS command. The change of an authorization level to a DEDB cannot be done because the referenced DEDB is in use by a region.</li> <li>• Changing the DEDB access intent requires stopping the scheduling of all PSBs that have access to the target DEDB. This might require stopping regions in which PSBs associated with WFI transactions are currently scheduled, and stopping the transactions or PSBs for non-WFI transactions.</li> </ul> <p>A /START AREA command failed because the area is marked as <i>recovery-needed</i> in the DBRC RECON data set, or the DBRC RECON indicates that there are no available area data sets for the area.</p> <p>Because a /START DB ACCESS command was issued, a change database authorization request was made from IMS to DBRC. The encoded database state returned from DBRC was set to a value of zero.</p> <p>A backout error occurred for the database. A /START command for the database is invalid until a /DBRECOVERY command is issued and batch backout is performed.</p> <p>One of the following events has occurred:</p> |
|    | 24   |  |  |  |
|    | 28   |  |  |  |
|    | 31   |  |  |  |
|    | 34   |  |  |  |
|    | 35   |  |  |  |
|    | 37   |  |  |  |

- A write error occurred for the database. The command issued is invalid until a /DBRECOVERY command is issued.
  - A request for a new dataspace was issued during the open processing for a shared VSO area. Refer to message DFS2444I, that precedes message DFS0488I.
- 38** A /START DB command was not executed because the database was being extended during an IMS XRF takeover, and I/O Prevention was not yet complete. You cannot start the database until IMS knows that I/O Prevention is complete.
- 39** A /START DB command specifying ACCESS=RO or RD is invalid while incomplete, restartable back outs exist for the database. IMS must have update or exclusive access to the database to redo the incomplete back outs. The command is not processed for the specified database.
- 41** The area is not loaded into the Coupling Facility. This does not prevent access to the area unless this code is preceded by message DFS2819A.
- 42** The named HALDB partition cannot be processed in response to a /START DB command due to an error in the Partition Structure Manager. Check the system console for related DFS0415I messages.
- 44** When the /STOP ADS command is entered, the ADS of the named ddname is already stopped. This /STOP ADS command is ignored.
- 50** A /DBRECOVERY, /START, or /STOP command was issued with AREAGROUP or DBGROUP, where the AREAGROUP contained a specified database or the DBGROUP contained a specified area. The command is not processed for databases in an AREAGROUP or areas in a DBGROUP.
- On a system defined to support Fast Path, the keyword 'MSDB' was entered on a /START, /STOP, or /DBRecovery command. ('MSDB' is allowed on /DBDUMP.) The keyword is ignored and the remainder of the command is processed.
- 51** A /DBRecovery command was entered on the RSR tracker for a Fast Path
- 52** DEDB. The DEDB name is ignored and the remainder of the command is processed.
- A /DBRecovery or /START command was entered for a database that was in the process of being recalled from HSM. The command process is not performed at this time. The /DBRecovery command must be entered again after the database is recalled. The /START command will be automatically performed after the recall is complete.
- 53** A command specifying a full-function DL/I database was entered on an RSR tracker executing with LSO=Y. Full-function databases are not available to LSO=Y tracker systems, so the command could not be processed for the named database.
- 54** There is a discrepancy between the number of ADS allocated by IMS and the number of ADS known to DBRC. Correct the discrepancy, then reissue the /DBRecovery or /STOP command to deallocate the data sets.
- 55** The command failed because the FF DB or FP area is currently being recovered.
- 56** An error occurred while attempting to open the named database or HALDB partition. Check the System Console for message DFS0730I to determine the reason for the failure.
- 57** The named database is a HALDB master and is invalid with the OPEN keyword.
- 58** The named HALDB partition cannot be processed in response to a /DBDUMP, /DBRECOVERY, /START, or /STOP command. It is still in use by an active Online Reorganization. The command is not processed for the HALDB partition named in the message, but processing continues for any other databases or HALDB partitions specified in the command.
- 60** The command failed because the area has been deleted by Online Change.
- 62** The named database failed directory initialization or the DBB or PSB pool space was not large enough to accommodate the blocks. Check the System Console for messages referencing this database.
- 70** The named database or HALDB

partition failed authorization. Check the System Console for messages related to the authorization failure.

- | 71 Open keyword processing could not continue for the named HALDB partition because a /DBRECOVERY command has been issued for the HALDB master. Start the HALDB master and then reissue the /START DB command for the HALDB partition.
- | 72 Partition DBM build failed for the named HALDB partition. Check the System Console for DFS0415I messages to determine the reason for the failure.
- | 73 DMB build processing failed for the named database or HALDB partition. Check the System Console for DFS0561I or DFS0592I messages to determine the reason for the failure.
- | 74 DMB build processing failed for the named database or HALDB partition because the data management block pool (DMB pool) is too small to hold all of the required data management blocks. Change the appropriate parameter to make the pool larger and restart the system. Message DFS0992I might also be issued.
- | 75 DMB build processing failed for the named database or HALDB partition. IMS detected a DMB pool shortage while attempting to load the DMB block for the named database or HALDB partition. The size of the DMB pool must be large enough to accommodate all DMBs that are actively being used. Change the appropriate parameter to make the pool larger and restart the system.
- | 76 DMB build processing failed for the named database or HALDB partition. Check the System Console for DFS0564I messages to determine the reason for the failure.
- | 77 DMB build processing failed for the named database or HALDB partition. An unexpected return code was received. Check the System Console for any related messages and contact the IBM Support Center.
- | 79 The command was issued for a database that is still being processed by Restart Parallel Database Open. Wait for that process to complete and try the command again.

**80** The named DEDB database cannot be started in a non-FP system.

**System Action:** IMS continues processing.

**Problem Determination:** 6, 11

| **Module:** DBFARD10, DBFARD40, DBFARDA0, DFSDBDR0, DFSRDSH0, DBFCVUN0

---

| **DFS0725I INITIATE | UPDATE | TERMINATE  
OLREORG COMMAND FOR DB  
dbnamexx COMPLETE.  
CC= nn**

| **Explanation:**

| **dbnamexx** is the HALDB partition name entered in the command

| **nn** is the completion code

| The /INITIATE, /UPDATE, or /TERMINATE OLREORG command processing is complete. A zero completion code indicates that the command has been successfully processed. Message DFS2970I is issued after OLR initialization completes to indicate that the online reorganization has started. A nonzero completion code indicates that an unexpected error was encountered during command processing.

| **Completion**

| <u>Code</u> | <u>Meaning</u>  |
|-------------|---|
| <b>0</b>    | REQUEST SUCCESSFUL  |
| <b>10</b>   | RESOURCE NAME INVALID   |
| <b>14</b>   | RESOURCE NOT A PARTITION NAME                                 |
| <b>1C</b>   | RESOURCE IS A PARTITION SECONDARY INDEX                       |
| <b>20</b>   | INIT OLREORG ISSUED, OLR IS ALREADY IN PROGRESS               |
| <b>24</b>   | NO OLREORG IS IN PROGRESS                                     |
| <b>28</b>   | NO DMB LOADED - need to schedule a PSB that references the DB |
| <b>C1</b>   | OLREORG DDIR MISSING OR DFSPNT ZERO                           |
| <b>C2</b>   | OLREORG ALREADY OWNED   |
| <b>C3</b>   | OLREORG PARTITION STOPPED, or LOCKED, or DBDUMPED             |
| <b>C4</b>   | OLREORG PARTITION BEING RECOVERED                             |
| <b>C5</b>   | OLREORG PARTITION ACCESS NOT UPDATE/EXCLUSIVE                 |
| <b>C6</b>   | OLREORG SCHEDULING FAILED - Check HALDB Master status         |



|           |  |
|-----------|--|
| <b>C7</b> | OLREORG PARTITION NOT<br>AUTHORIZED - Check HALDB<br>Master status |
| <b>C8</b> | OLREORG MASTER STOPPED, or<br>LOCKED, or DBDUMPED                  |
| <b>C9</b> | OLREORG MASTER ACCESS NOT<br>UPDATE or EXCLUSIVE                   |
| <b>CA</b> | OLREORG SET OWNERSHIP<br>FAILED                                    |
| <b>CB</b> | PARTITION NOT IN SPECIFIED<br>STATUS                               |

| **System Action:** IMS continues processing.

| **Problem Determination:** 6, 11

| **Module:** DFSORC10, DFSORP50

---

**DFS0832I ABEND Uwww REASON CODE xxxx  
yyyyyyyyy zzzzzzzz**

**Explanation:** This is an informational message issued to assist you to debug problems that cause abends. This message is issued when an abend condition is detected by module DFSDLA00. This condition can be caused by either:

- A non-zero return code on a partition selection or validation.
- A pseudo-abend set by module DFSDLR00, DFSDDL00, DFSDXMT0, DFSDLD00, or DFSRCHB0.

The value of xxxx is the detected user abend. For these abends the reason code is yyyy and it is related to a problem that was detected while processing a HALDB:

- 0880
- 3303
- 3498

Where:

|             |  |
|-------------|--|
| <b>www</b>  | is the pseudo-abend code   |
| <b>xxxx</b> | is the reason code for HALDB abends, or N/A for non-HALDB abends   |
| <b>yyyy</b> | If xxxx (the abend) is not one of the above, the value of yyyy is N/A otherwise it is one of the following reason codes: |
| <b>4001</b> | No more partitions are defined   |
| <b>4020</b> | DB authorization error (data unavailable)  |
| <b>4025</b> | DB unload action error   |
| <b>4026</b> | DB initialization action error   |
| <b>4059</b> | The access level is not compatible with the intent level   |
| <b>8001</b> | Target partition requires initialization   |
| <b>8010</b> | Target partition not found   |
| <b>8031</b> | Open DMB failure   |
| <b>8051</b> | User partition selection failed  |
| <b>C001</b> | Non-partitioned DB access  |
| <b>C002</b> | Partitions not active  |

|             |  |
|-------------|--|
| <b>C003</b> | Partition stopped                                |
| <b>C004</b> | Partition structure terminating                  |
| <b>C005</b> | Partition structure not initialized              |
| <b>C006</b> | Partition structure rebuild failed               |
| <b>C009</b> | ILDS DSG not provided                            |
| <b>C010</b> | Unusual structure status                         |
| <b>C030</b> | Invalid authorization reason code                |
| <b>C040</b> | Invalid partition DDIR                           |
| <b>C041</b> | Invalid partition set operation                  |
| <b>C042</b> | Unable to allocate PSET storage                  |
| <b>C043</b> | Invalid partition set selection action           |
| <b>C050</b> | User selection exit not invoked                  |
| <b>C060</b> | Invalid function code                            |
| <b>C070</b> | Buffer handler byte locate failed                |
| <b>C071</b> | Buffer handler locate by key failed              |
| <b>C073</b> | Buffer handler mark altered failed               |
| <b>C074</b> | Buffer handler byte locate & mark altered failed |
| <b>C075</b> | IRLM root lock request (GRID) failed             |
| <b>C076</b> | Request to obtain BCB storage (BCBGET) failed    |

#### **PARTITION | DATABASE**

is PARTITION if HALDB, or DATABASE if non-HALDB

#### **zzzzzzzz**

is the DBD name for one of these:

- a HALDB partition
- a HALDB master
- a non-HALDB database

or NOTAPPLI if no DBD is available

**System Action:** The application program is terminated.

**Programmer Response:** Correct the situation that caused the application to terminate.

**Module:** DFSDLA00

---

#### **DFS2262 TOO MANY DATABASE PARAMETERS**

| **Explanation:** Too many database parameters were entered for this command. Only one can be entered.

| **System Action:** The command is ignored.

| **Operator Response:** Reenter the command with only one database parameter.

| **Module:** DFSORC10

---

#### **DFS2970I OLR STARTED | RESUMED FOR NAME=partname**

| **Explanation:** HALDB Online Reorganization has started or resumed for the named partition.

| **System Action:** IMS continues processing.

| **Module:** DFSORP00

| **Note:** Message DFS2970I is issued after the type-1 or type-2 INIT OLREORG command is processed to the extent that OLREORG ownership is

established in RECON. If the online reorganization cursor is active then *RESUMED* is indicated; otherwise, *STARTED* is indicated.

**DFS2971W OLR INITIALIZATION FAILED FOR NAME=partname RC=xx RS=yy**

**DFS2971W OLR TERMINATED FOR NAME=partname RC=xx RS=yy**

**Explanation:** First format of the message: For the reason indicated by the return and reason codes, the HALDB online reorganization did not complete the OLR initialization process to either start or resume the OLR.

Second format of the message: For the reason indicated by the return and reason codes, the OLR terminated before all records were copied from the input to the output set of data sets.

See the following return and reason codes for the exact cause of the termination:

**Problem Determination:**

| Return Code (Hex) | Reason Code (Dec) | Description   |
|-------------------|-------------------|---|
| 04                | 04                | Unexpected return code from ISWITCH                     |
|                   | 08                | DFSBCB GET for QSAV failed                              |
|                   | 12                | DFSCWU create work unit failed                          |
|                   | 16                | IMS was restarted with COLDSYS or COLDBASE              |
|                   | 20                | Restart authorization failed                            |
|                   | 24                | Deferred stop or close due to error during XRF takeover |

**Module:** DFSORP00

| Return Code (Hex) | Reason Code (Hex) | Description   |
|-------------------|-------------------|---|
| 08                | 01                | GETMAIN for OLR work area failed                        |
|                   | 02                | Scheduling of the OLR PSB failed                        |
|                   | 03                | Output data set creation or validation failed           |
|                   | 04                | DFSBCB GET for QSAV failed                              |
|                   | 05                | Timer ITASK creation failed                             |
|                   | 06                | Database in error because backout failed                |
| 0C                |                   | Partition selection failed                              |
|                   | 8031              | Open error for partition                                |
|                   | C030              | Authorization failure                                   |
|                   | Cxxx              | Other partition selection failures                      |
| 10                | abend code        | Pseudo-abend during partition selection (decimal value) |

**Module:** DFSORP10

| Return Code (Hex) | Reason Code (Hex) | Description   |
|-------------------|-------------------|---|
| 14                | 01                | OLR terminated by a TERMINATE OLREORG command           |
|                   | 02                | Buffer handler call for PHDAM failed                    |
|                   | 03                | Buffer handler call for PHIDAM failed                   |
|                   | 04                | GBIDP lock failure during cursor update                 |
|                   | 05                | GZIDP lock failure during cursor update                 |
|                   | 06                | Buffer handler failure during cursor update             |
|                   | 07                | Notify failure during cursor update                     |
|                   | abend code        | Pseudo-abend during DFSORP20 processing (decimal value) |

**Module:** Issued by DFSORP10 for errors detected in DFSORP20

| Return Code (Hex) | Reason Code (Hex) | Description   |
|-------------------|-------------------|---|
| 1C                | 01                | Backout failure   |
|                   | 02                | OLR terminated by a TERMINATE OLREORG command with the FORCE option |
|                   | 03                | Key with all X'FF' found, or error doing chain maintenance          |
| 20                | abend code        | Pseudo-abend during DFSORP30 processing (decimal value)             |
| 24                | status code       | Unexpected status code from internal DL/I call                      |

**Module:** Issued by DFSORP10 for errors detected in DFSORP30

**System Action:** First format of the message: For an INITIATE OLREORG command that tried to start a new online reorganization, the cursor-active status is not recorded in the RECON. And the database records in the partition are still contained in the original set of data sets that were in use before the command was entered.

For an automatic or command-initiated resumption of an online organization for which the cursor-active status was already recorded in the RECON:

- The cursor-active status is still in effect.
- The database records in the partition are still contained in two sets of data sets, both of which are required for continued use of the partition.
- No additional database records were copied from the input to the output set of data sets.

Second format of the message: The cursor-active status for the partition is already recorded in the RECON, and

the database records in the partition are now contained in two sets of data sets, both of which are required for continued use of the partition.

**System Programmer Response:** After the problem is corrected, start or resume the HALDB online reorganization with the INITIATE OLREORG command or else run an offline reorganization.

**Problem Determination:** 1, 6, 27

---

#### DFS2972E OLR PSB BUILD FAILED. *reason-text*

**Explanation:** IMS was building an internal PSB for an online reorganization. This process encountered the error indicated by *reason-text*. The possible values of *reason-text* are listed here:

- DATABASE IS NOT A PARTITION
- MASTER IS NOT PARTITIONED
- GETMAIN FAILED FOR PSB
- DMB LOAD FAILED

**System Action:** The online reorganization is abandoned, and the *must-complete* status is not recorded in the RECON.

**Programmer Response:** In the case of a DMB LOAD FAILED situation, ensure that the ACBLIB contains the DMB with the name of the HALDB master.

In the case of a GETMAIN FAILED FOR PSB situation, storage may be constrained so run the online reorganization later.

The other errors are probably IMS errors, so contact the IBM support center.

**Module:** DFSOPSB0

---

#### DFS2973W NO RESPONSE FROM SHARING IMS AFTER ONLINE REORG. PART=partname

**Explanation:** HALDB online reorganization has completed for the named HALDB partition and the inactive data sets were to be deleted. Data set deletion was attempted because OPTION(DEL) was specified or defaulted on the INITIATE OLREORG command or was subsequently specified on the UPDATE OLREORG command. However, before data set deletion was attempted, all sharing IMS systems were notified, and they were to indicate that they were no longer using these data sets. One or more of these systems did not respond even after three requests.

**System Action:** To complete the processing, IMS attempts to delete the inactive data sets anyway. An error may result during data set deletion if some job, such as an IMS batch job, has the data sets allocated but is unable to respond within the time allowed for a response. If this happens, one or more of the inactive data sets won't be deleted.

**System Programmer Response:** Check the MVS system console log to determine whether error

messages indicate that any of the inactive data sets from the HALDB online reorganization were not deleted. After these data sets are no longer in use, delete them.

**Module:** DFSORP40

---

#### DFS2974I OLR COMPLETED FOR NAME=partname

**Explanation:** HALDB Online Reorganization has completed for the named HALDB partition.

**System Action:** IMS continues processing.

**Module:** DFSORP10

---

#### DFS2990E *macro* MACRO, RETURN=X'*returncode*', [ REASON=X'*reasoncode*', ], PART=*partname* DSN=*dsname*

**Explanation:** IMS internally invoked a macro instruction to gather information about the data set *dsname* for the HALDB partition *partname* as part of either:

- Starting a HALDB online reorganization for the HALDB partition *partname*.
- Tracking these changes at the RSR tracking site.

This macro instruction completed with a register 15 return code value of *returncode* and a register 0 value of *reasoncode*. These values were not expected by IMS.

**System Action:** If this error occurred in starting the reorganization, then the reorganization is not started, but the partition remains accessible with the original active set of data sets--either the A-through-J and X or the M-through-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline.

In the first two cases, one or more of the output data sets could have already been created automatically.

**Programmer Response:** Based on the macro indicated in the message, see the appropriate publication in the following table to determine the cause of the error, and follow any recommended actions:

| Macro   | Publication  |
|---------|--|
| DEVTYPE | DFSMS/MVS DFSMSdftp Advanced Services                                |
| GETDSAB | z/OS MVS Programming: Authorized Assembler Services Reference, Vol 2 |
| LOCATE  | DFSMS/MVS DFSMSdftp Advanced Services                                |
| OBTAIN  | DFSMS/MVS DFSMSdftp Advanced Services                                |
| SWAREQ  | z/OS MVS Programming: Authorized Assembler Services Reference, Vol 4 |



| Macro   | Publication                          |
|---------|--------------------------------------|
| TRKCALC | DFSMS/MVS DFSMSdfp Advanced Services |

If the error was detected in processing the INITIATE OLREORG command, reissue the command for partition *partname*.

If the error was detected at the RSR tracking site, issue a /START DB command for the partition in order for OFR (Online Forward Recovery) to continue updating the partition.

**Module:** DFSORA00, DFSORA10

**DFS2991E UNACCEPTABLE DATA SET, REASON CODE** *reasoncode*.  
**PART=***partname* **DSN=***dsname*

**Explanation:** In determining whether to create the data set *dsname* for the HALDB partition *partname*, IMS found either:

- A data set of that same name. However, because of certain characteristics of this preexisting data set, the data set could not be used for the intended purpose.
- No data set of that name when there should have been a preexisting data set.

**Case 1:** In starting a HALDB online reorganization for the HALDB partition *partname*, IMS found that the data set that was to be used as the output of the reorganization either:

- Was not of a type that could be used for the applicable data set.
- Had certain characteristics that did not match those of the corresponding input data set.
- Did not exist.

**Case 2:** On the RSR tracking site, IMS found that in tracking database changes for a HALDB online reorganization for the HALDB partition *partname*, the data set that was to be used as the shadow of an output data set either:

- Was not of a type that could be used for the applicable data set.
- Had certain characteristics that did not match those of the corresponding input data set at the active site.
- Did not exist.

In any case, *reasoncode* is a reason code the specific problem. These reason codes are shown in the following table. The "Environment" column indicates whether the error is detected at the beginning of either the reorganization itself or the tracking of the reorganization at the RSR tracking site.

Table 21. Reason Codes, Environments, and Errors Detected for Message DFS2991I

| Reason Code | Environment    | Errors Detected for Pre-Existing Data Set  |
|-------------|----------------|--|
| 1           | Both           | The data set is not a DASD data set.   |
| 2           | Both           | The data set is a VSAM data set, but the database is defined to use OSAM.  |
| 3           | Both           | The data set is a PDS, PDSE, or a member of a PDS or PDSE.   |
| 4           | Both           | The data set is not a VSAM ESDS, but the database is defined to use VSAM.  |
| 5           | Both           | The VSAM data set does not have the REUSE attribute.   |
| 6           | Reorganization | The VSAM data set has a different record length than that of the corresponding input data set.   |
| 7           | Reorganization | The VSAM data set has a different control interval size than that of the corresponding input data set.   |
| 8           | Reorganization | The VSAM ESDS has a different control interval size than that of the corresponding input data set.   |
| 10          | Reorganization | The primary index data set of a PHIDAM database (the KSDS), is defined with a different key offset or key length than that of the corresponding input data set.  |
| 11          | Tracking       | The VSAM data set has a different record length than that of the tracked output data set at the active site.   |
| 12          | Tracking       | The VSAM data set has a different control interval size than that of the tracked output data set at the active site.   |
| 13          | Tracking       | The primary index data set of a PHIDAM database (the KSDS), is defined with a different key offset or key length than that of the tracked output data set at the active site.  |
| 18          | Both           | Either: <ul style="list-style-type: none"> <li>• The data set is cataloged but it is not on the volume to which the catalog entry refers.</li> </ul> or <ul style="list-style-type: none"> <li>• The data set was allocated with a DD statement that had a VOLUME parameter, but the data set was not on the referenced volume.</li> </ul> |

Table 21. Reason Codes, Environments, and Errors Detected for Message DFS2991I (continued)

| Reason Code | Environment              | Errors Detected for Pre-Existing Data Set   |
|-------------|--------------------------|---|
| 19          | Both                     | The data set was allocated with a DD statement, but the data set name in the DSNAMES parameter does not follow the data set naming convention defined for this partition's data sets. |
| 20          | Tracking, Reorganization | The output data set does not exist, but the corresponding input data set resides on multiple volumes and is not an SMS-managed data set.  |

**System Action:** If this error occurred in starting the reorganization, then the reorganization is not started, but the partition remains accessible with the original active set of data sets--either the A-through-J and X or the M-through-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline.

In both cases, one or more of the output data sets could have already been created automatically.

**Programmer Response:**

**Case 1:** Before restarting the HALDB online reorganization, either:

- Delete the preexisting output data set, thus allowing IMS to create it automatically. Note that when IMS creates the data set, the size and placement of the data set is determined automatically by IMS, MVS, and DFSMS/MVS, and also by parameters defined by your site's storage administrator. Therefore, do not let IMS create the data set if you need precise control over the data set's size and placement.
- Create a data set that is acceptable as the output of the reorganization:

Ensure that for an OSAM-defined database the output data set for the database segments (one of the A-through-J or the M-through-V data sets):

- Is a DASD data set.
- Is not a VSAM data set.
- Is neither a PDS, PDSE, nor a member of a PDS or PDSE or that for a VSAM-defined database the output data set for the database segments (one of the A-through-J or the M-through-V data sets):
  - Is a VSAM entry-sequenced data set (ESDS).
  - Has the REUSE attribute.
  - Has a fixed-length record length identical to that of the input data set.
  - Has a control interval size identical to that of the input data set or that for a PHIDAM database the output data set for the primary index (the X or Y data set):

- Is a VSAM key-sequenced data set (KSDS).
- Has the same key offset and key length as the corresponding input data set.
- Has the REUSE attribute.

Then reissue the INITIATE OLREORG command to start the HALDB online reorganization for partition *partname*.

**Case 2:** Before restarting the partition at the RSR tracking site, either:

- Delete the preexisting output data set at the tracking site, thus allowing IMS to create it automatically. Note that when IMS creates the data set, the size and placement of the data set is determined automatically by IMS, MVS, and DFSMS/MVS, and also by parameters defined by your site's storage administrator. Therefore, don't let IMS create the data set if you need precise control over the data set's size and placement.

- Create a data set that is acceptable for tracking the output of the reorganization:

Ensure that for an OSAM-defined database the output data set for the database segments (one of the A-through-J or the M-through-V data sets) at the RSR tracking site:

- Is a DASD data set.
- Is not a VSAM data set.
- Is neither a PDS, PDSE, nor a member of a PDS or PDSE or that for a VSAM-defined database the output data set for the database segments (one of the A-through-J or the M-through-V data sets) at the RSR tracking site:

- Is a VSAM entry-sequenced data set (ESDS).
- Has the REUSE attribute.
- Has the same fixed-length record length as the tracked output data set at the active site (as shown in the X'2930' log record).
- Has the same control interval size as the tracked output data set at the active site (as shown in the X'2930' log record) or that for a PHIDAM database the output data set for the primary index (the X or Y data set) at the RSR tracking site:

- Is a VSAM key-sequenced data set (KSDS).
- Has the REUSE attribute.
- Has the same fixed-length record length as the tracked output data set at the active site (as shown in the X'2930' log record).
- Has the same key offset and key length as the tracked output data set at the active site (as shown in the X'2930' log record).

Then issue a /START DB command for partition *partname* in order for Online Forward Recovery to continue updating the partition.

| **Module:** DFSORA00

---

| **DFS2992E** CSI RETURN CODE=4, [ { CATALOG  
| ERROR, | DATA SET ERROR, } ]  
| CATALOG MANAGEMENT RETURN  
| CODE=*cmreturncode* REASON  
| CODE=*cmreasoncode*  
| MODULE=*cmmoduleid* PART=*partname*  
| DSN=*dsname*

---

| **DFS2992E** CSI RETURN CODE=*csireturncode* [  
| REASON CODE=*csireasoncode* ]  
| PART=*partname* DSN=*dsname*

| **Explanation:** IMS internally invoked the Catalog  
| Search Interface (CSI) to gather information about the  
| data set *dsname* for the HALDB partition *partname* as  
| part of either:

- | • Starting a HALDB online reorganization for the  
| partition.
- | • Tracking these changes at the RSR tracking site.

| The CSI completed with a CSI return code value as  
| shown in the message text. Information on these CSI  
| return codes is in *DFSMS/MVS Managing Catalogs*.

| For the first form of the message (when the CSI return  
| code is 4), a Catalog Management error was detected.  
| In this case, *cmreasoncode*, *cmreturncode*, and  
| *cmmoduleid* are the Catalog Management reason code,  
| return code, and two-character module identification,  
| respectively. Information on this Catalog Management  
| error information is in *z/OS MVS System Messages, V6*  
| (*GOS-IEA*) under message IDC300 9I.

| For the second form of the message (when the CSI  
| return code, *csireturncode* is other than 4), a CSI error  
| was detected. For CSI return code 8, *csireasoncode* is  
| the CSI reason code. Information on this CSI error  
| information is in *DFSMS/MVS Managing Catalogs*.

| These values were not expected by IMS.

| **System Action:** If this error occurred in starting the  
| reorganization, then the reorganization is not started,  
| but the partition remains accessible with the original  
| active set of data sets--either the A-through-J and X or  
| the M-through-V and Y data sets. If this error occurred  
| at the RSR tracking site, then the shadow partition is  
| taken offline.

| In both cases, one or more of the output data sets could  
| have already been created automatically.

| **Programmer Response:** Based on value of the CSI  
| return code, refer to the publication indicated above to  
| determine the cause of the error, and follow any  
| recommended actions.

| Then if the error was detected in processing the

| INITIATE OLREORG command, reissue the command for  
| partition *partname*.

| If the error was detected at the RSR tracking site, issue  
| a /START DB command for the partition in order for  
| Online Forward Recovery to continue updating the  
| partition.

| **Module:** DFSORA80

---

| **DFS2993E** NON-DASD DATA SET. DEVICE  
| CLASS=X'*devclass*'  
| PART=*partname* DSN=*dsname*

| **Explanation:** The data set *dsname*, which is an input  
| data set for a HALDB online reorganization of the  
| HALDB partition *partname*, is not a DASD data set.  
| (The device class reported by the internally executed  
| DEVTYPE macro instruction is *devclass* (hexadecimal)  
| rather than X'20'.) This is inconsistent with its use as  
| part of a database.

| **System Action:** If this error occurred in starting the  
| reorganization, then the reorganization is not started.  
| The partition remains accessible, but successful use of  
| it is not likely.

| If this error occurred at the RSR tracking site, then the  
| shadow partition is taken offline.

| **Module:** DFSORA00

---

| **DFS2994E** IDCAMS RETURN CODE *returncode*  
| CREATING AN ONLINE REORG  
| OUTPUT DATA SET. PART=*partname*  
| DSN=*dsname*

| **Explanation:** IMS tried to create the VSAM output  
| data set *dsname* for the HALDB partition *partname* as  
| part of either:

- | • Starting a HALDB online reorganization for the  
| partition.
- | • Tracking these changes at the RSR tracking site.

| Access Method Services (IDCAMS) gave a return code  
| of *returncode*.

| **System Action:** The message, if any, produced by  
| IDCAMS was displayed prior to this message.

| For return code 4, the data set is assumed to have  
| been created. In this case, none of the following actions  
| apply, and processing continues.

| If this error occurred in starting the reorganization, then  
| the reorganization is not started, but the partition  
| remains accessible with the original active set of data  
| sets--either the A-through-J and X or the M-through-V  
| and Y data sets. If this error occurred at the RSR  
| tracking site, then the shadow partition is taken offline.

| In both cases, one or more of the output data sets could  
| have already been created automatically.

| **Programmer Response:** See the *z/OS MVS System*

| *Messages, V6 (GOS-IEA)* to determine the cause of the error, and follow any recommended actions.

| Then if the error was detected in processing the INITIATE OLREORG command, reissue the command for partition *partname*.

| If the error was detected at the RSR tracking site, issue a /START DB command for the partition in order for Online Forward Recovery to continue updating the partition.

| **Module:** DFSORA20

---

| **DFS2995E IDCAMS RETURN CODE returncode**  
 | **DELETING AN ONLINE REORG INPUT**  
 | **DATA SET.**  
 | **PART=partname DSN=dsname**

| **Explanation:** IMS tried to delete the VSAM input data set *dsname* for the HALDB partition *partname* as part of either:

- | • Completing a HALDB online reorganization for the partition.
- | • Completing the tracking of these changes at the RSR tracking site.

| Access Method Services (IDCAMS) gave a return code of *returncode*.

| **System Action:** The message, if any, produced by IDCAMS was displayed prior to this message.

| The reorganization or the tracking of the reorganization at the RSR tracking site is still considered to have completed successfully, and there is now a single new active set of data sets for partition *partname*--either the A-through-J and X or the M-through-V and Y data sets.

| **Programmer Response:** See *z/OS MVS System Messages, V6 (GOS-IEA)* to determine the cause of the error, and follow any recommended actions, then delete the data set if it has not already been deleted.

| **Module:** DFSORA20

---

| **DFS2996E NOT ENOUGH DASD SPACE. ERROR**  
 | **REASON=X's99errorcode'**  
 | **[ SMS REASON=s99ersn ]**  
 | **BLOCKS=blocks BLKSIZE=blksize**  
 | **VOLUMES=volcount**  
 | **{ UNIT=SYSALLDA | STORAGE**  
 | **CLASS=storageclassname }**  
 | **PART=partname DSN=dsname**

| **Explanation:** IMS tried to create the OSAM output data set *dsname* for the HALDB partition *partname* as part of either:

- | • Starting a HALDB online reorganization for the partition.
- | • Tracking these changes at the RSR tracking site.

| The internally executed DYNALLOC macro instruction (SVC 99) that was used for this purpose reported that the data set could not be created because of DASD space or DASD volume constraints. The requested space was *blocks* blocks, each with a block size of *blksize*. There were *volcount* volumes requested.

| If the text showing the group name (SYSALLDA) appears in the message text, the attempt was to create a non-SMS-managed data set because the corresponding input data set was not SMS-managed. If the text showing the storage class appears, then the attempt was to create an SMS-managed data set using storage class *storageclassname* because the corresponding input data set was SMS-managed and was created using that storage class.

| If the reorganization was being started, these space allocation values were derived from those of the corresponding input data set. On the RSR tracking site, these space allocation values were derived from those of the tracked output data set at the active site where the reorganization occurred.

| For further information of this error, see the *z/OS MVS Programming: Authorized Assembler Services Guide*. The value *s99errorcode* is the hexadecimal value of the error reason code (field S99ERROR in the SVC 99 request block) that is provided when the DYNALLOC macro instruction (SVC 99) fails with return code 4. In the SMS case, *s99ersn* is the SMS reason code (field S99ERSN in the SVC 99 request block extension). The SMS reason code corresponds to the message number of an IGD message describing the error. For further information on the IGD message, see the *z/OS MVS System Messages, Vol 8 (IEF-IGD)*.

| **System Action:** If this error occurred in starting the reorganization, then the reorganization is not started, but the partition remains accessible with the original active set of data sets--either the A-through-J and X or the M-through-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline.

| In both two cases, one or more of the output data sets could have already been created automatically.

| **Programmer Response:** Determine whether the requested space (*blocks* blocks of size *blksize*) is really required for the output data set. In some cases, this space allocation request could include space acquired in a secondary allocation for the input data set and might not be required for the reorganized database. Then either:

- | • Ensure that there is sufficient DASD space on the volumes referenced by the group name SYSALLDA (for the non-SMS-managed case) or referenced by the storage class *storageclassname* (for the SMS-managed case), possibly by scratching unneeded data sets. Note that the primary allocation amount (*blocks* blocks of size *blksize*) must be



available in a total of no more than five extents on a single volume. In the SMS-managed case, if the storage class definition includes the GUARANTEED SPACE attribute, then the primary allocation amount must be available and will be allocated on each of the volumes (rather than on just the first volume).

- Create and catalog this output data set yourself. In this case, you can reserve whatever amount of space you choose, and you can control the number of volumes. See the *IMS Version 9: Administration Guide: Database Manager*, for information on allocating OSAM data sets along with a technique for pre-allocating multi-volume data sets.

Then if the error was detected in processing the INITIATE OLREORG command, reissue the command for partition *partname*.

If the error was detected at the RSR tracking site, issue a /START DB command for the partition in order for OFR to continue updating the partition.

**Module:** DFSORA20

---

**DFS2997E DATA SET ALLOCATED TO ANOTHER JOB AND CANNOT BE CREATED.**  
**PART=partname DSN=dsname**

**Explanation:** IMS tried to create the OSAM output data set *dsname* for the HALDB partition *partname* as part of either:

- Starting a HALDB online reorganization for the partition.
- Tracking these changes at the RSR tracking site.

Creating the data set required that the data set be allocated as though DISP=NEW had been specified on a DD statement through JCL. However, a data set of the same name was already in use by another job, causing the allocation to fail.

**System Action:** If this error occurred in starting the reorganization, then the reorganization is not started, but the partition remains accessible with the original active set of data sets--either the A-through-J and X or the M-through-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline.

In both cases, one or more of the output data sets could have already been created automatically.

**Programmer Response:** Ensure that no other job is using a data set of the name *dsname* even if that data set isn't cataloged and even if that data set is not intended to be the same data set as the output of the reorganization.

Then if the error was detected in processing the INITIATE OLREORG command, reissue the command for partition *partname*.

If the error was detected at the RSR tracking site, issue

a /START DB command for the partition in order for OFR to continue updating the partition.

**Module:** DFSORA20

---

**DFS2998E [ DATA SET CREATION ERROR. I UNALLOCATION FAILURE. ]**  
**RETURN=returncode**  
**[ { INFO. REASON=X's99infocode' | ERROR REASON=X's99errorcode' } ]**  
**[ SMS REASON=s99ersn ]**  
**BLOCKS=blocks BLKSIZE=blksize**  
**VOLUMES=volcount**  
**{ UNIT=SYSALLDA | STORAGE**  
**CLASS=storageclassname }**  
**PART=partname DSN=dsname**

**Explanation:** IMS tried to create the OSAM output data set *dsname* for the HALDB partition *partname* as part of either:

- Starting a HALDB online reorganization for the partition.
- Tracking these changes at the RSR tracking site.

As indicated by the message text, the internally executed DYNALLOC macro instruction (SVC 99) reported one of these two problems:

1. The data set could not be created.
2. The data set that was just created successfully could not be unallocated from the IMS job.

The requested space was *blocks* blocks, each with a block size of *blksize*. There were *volcount* volumes requested.

When the text showing the group name (SYSALLDA) appears in the message text, the attempt was to create a non-SMS-managed data set because the corresponding input data set was not SMS-managed. If the text showing the storage class appears, then the attempt was to create an SMS-managed data set using storage class *storageclassname* because the corresponding input data set was SMS-managed and was created using that storage class.

If the reorganization was being started, these space allocation values were derived from those of the corresponding input data set. On the RSR tracking site, these space allocation values were derived from those of the tracked output data set at the active site where the reorganization occurred.

For the meaning of the various error codes shown in this message, see the *z/OS MVS Programming: Authorized Assembler Services Guide*.

**returncode**

Is the decimal value of the register 15 return code provided by the DYNALLOC macro instruction (SVC 99).

**s99infocode**  
Is the hexadecimal value of the information reason code (field S99INFO in the SVC 99 request block).

**s99errorcode**  
Is the hexadecimal value of the error reason code (field S99ERROR in the SVC 99 request block).

**s99ersn**  
Is, for the SMS case, the SMS reason code (field S99ERSN in the SVC 99 request block extension). The SMS reason code corresponds to the message number of an IGD message describing the error. For further information on the IGD message, see the *z/OS MVS System Messages, Vol 8 (IEF-IGD)*.

**System Action:** If this error occurred in starting the reorganization, then the reorganization is not started, but the partition remains accessible with the original active set of data sets--either the A-through-J and X or

the M-through-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline.

In the first two cases, one or more of the output data sets could have already been created automatically.

**Programmer Response:** Using the information found in the *z/OS MVS Programming: Authorized Assembler Services Guide*, determine the meaning of returncode, *s99info*, *s99errorcode*, and *s99ersn*, as applicable. Follow any recommended actions to correct the error.

Then if the error was detected in processing the INITIATE OLREORG command, reissue the command for partition *partname*.

If the error was detected at the RSR tracking site, issue a /START DB command for the partition in order for Online Forward Recovery to continue updating the partition.

**Module:** DFSORA20

---

## HALDB OLR DSP Series Messages

These are the IMS DSP series messages that pertain to HALDB OLR. They also appear in the *IMS Version 9: Messages and Codes, Volume 1*.

---

**DSP0197I ONLINE REORG NOTIFICATION NOT VALID FOR DATABASE ORGANIZATION**

**Explanation:** A command to update RECON regarding an online reorganization was issued for a database that is not the correct type. Online reorganization is valid for HALDB PHDAM or PHIDAM databases only.

**System Action:** Processing is terminated. The command fails. Review RECON contents and reissue the command appropriately.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP0198I RECOVERY NOT ALLOWED DURING ONLINE REORGANIZATION**

**Explanation:** Recovery when online reorganization is active is not allowed.

**System Action:** Processing is terminated. The command fails. Review RECON contents and reissue the command appropriately.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP0199I IMAGE COPY RUNTIME COINCIDES WITH ONLINE REORG EXECUTION**

**Explanation:** A NOTIFY command to add an image copy was issued. Online reorganization is active.

**System Action:** Processing is terminated. The command fails. Review RECON contents.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP0292I ONLINE REORG ACTIVE FOR DBD=XXXXXXXX**

**Explanation:** A command was issued to add or change information in RECON. The command cannot be processed when online reorganization is active. If online reorganization is active and has an owning IMS which is currently active, parameters OLRRGON | OLRRGOFF, OLRIMSID() | NOOWNER, and OLRDBDS(A|M) are not allowed. Parameter OLRDBDS(A|M) is not allowed if online reorganization is active.

**System Action:** Processing is terminated. The command fails. Review RECON contents.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP0293I ACTIVE DBDS=M-V FOR  
DBD=XXXXXXX**

**Explanation:** A command was issued to add or change information in RECON. The command cannot be processed when the M-V DBDS data sets are active.

**System Action:** Processing is terminated. The command fails. Review RECON contents.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP0295I SUBSYSTEM RECORD NOT FOUND IN  
RECON**

**Explanation:** A command was issued to change the owning IMS for a HALDB online reorganization but the subsystem record for OLRIMSID name was not found.

**System Action:** Processing is terminated. The command fails. Review RECON contents. The subsystem record must exist.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP0830I RECOVERY REQUEST INCONSISTENT  
WITH ONLINE REORG RUNTIME  
TIMESTAMP**

**Explanation:** The requested recovery is invalid due to inconsistent information found in RECON. One of the following occurred: An online reorganization has completed prior to the full recovery and the online reorganization cannot be used as input to recovery or the recovery-to-time falls within an online reorganization and time stamp recoveries are not allowed.

**System Action:** Processing is terminated. The RECON data sets are not updated. If the online reorganization cannot be used as input to recovery, select a time stamp that is valid or recover the active DBDS. Choose a valid time stamp if the recovery-to-time is within an online reorganization.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP0831I ONLINE REORG INPUT NOT VALID  
FOR RECOVERY**

**Explanation:** A GENJCL.RECOV command specifying USEOLRRG was issued and DBRC determines that an image copy is needed instead.

**System Action:** Processing is terminated. The command fails. Review RECON contents and reissue the command appropriately.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP0832I FULL RECOVERY NOT ALLOWED FOR  
AN INACTIVE DBDS**

**Explanation:** A full recovery of a DBDS was requested. However, the specified DBDS is an inactive DBDS of a HALDB partition.

**System Action:** The GENJCL.RECOV command or Database Recovery utility terminates processing for the DBDS.

**Programmer Response:** Determine if the GENJCL.RECOV command or the SYSIN control statement is correct. If not, make the correction and rerun the job.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

**Module:** DSPURVIN

---

**DSP0833I ONLINE REORG INPUT IS BEING USED  
FOR RECOVERY DBD=dbname  
DDN=ddname**

**Explanation:** This message indicates that the recovery for the identified DBDS will be performed using an online reorganization as input.

**System Action:** Processing continues.

**Programmer Response:** None.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

**Module:** DSPURVIN

---

**DSP1077I THE ACTIVE DBDS IS BEING IMAGE  
COPIED**

**Explanation:** The inactive DBDS data set was requested to be image copied. Only active DBDSs may be image copied. The active DBDS will be image copied.

**System Action:** Processing continues. This is an informational message.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP1078I    INACTIVE M-V DBDS RECORDS FOR  
DBD=XXXXXXXX DELETED FROM  
RECON**

**Explanation:** A command was issued to change information in RECON. As a result of the command processing, the inactive M-V DBDSs were deleted. If Y DBDSs existed, they were also deleted.

**System Action:** Processing continues. This is an informational message.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP1079I    RECON INCONSISTENCY, DBDS  
RECORD CANNOT BE FOUND FOR  
DBD=XXXXXXXX DBN=XXXXXXXX**

**Explanation:** An inconsistency was detected in the RECON. A DBDS record for the indicated DBD and DBN is not found.

**System Action:** Processing is terminated. The RECON data sets are not updated.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

**DSP1080I    Online Reorg with runtime timestamp  
has been reset to active status**

**Explanation:** An unknown command was used to change the status of an online reorganization record.

**System Action:** Informational message.

**Programmer Response:** Contact the IBM Support Center if you require additional assistance.

**Problem Determination:** 2, 3, 8, and listings of the RECON data sets.

---

## HALDB OLR Abends and Codes

The section lists the abends that include HALDB OLR support.

**Abend U0369** This abend is set when the buffer handler router detects an RBA that is inconsistent with the data set characteristics of either the A-through-J or M-through-V set of data sets. If it is on the A-through-J side, the RBA must be even; if it is on the M-through-V side, the RBA must be odd.

**Abend U0474** The ABORT option of the TERMINATE OLREORG command causes the HALDB OLR to be stopped immediately, possibly with a completion code of Abend U0474, without waiting for the current unit of reorganization to complete.

**Abend U381** There are some additional Error Codes for Abend U381.

**Abend U0880** There are some additional user abend reason codes for Abend U0880.

**Related Reading:** For detailed information on these or other abends, see the *IMS Version 9: Messages and Codes, Volume 1* and the *IMS Version 9: Failure Analysis Structure Tables (FAST) for Dump Analysis*.



---

## Chapter 8. Log Records of HALDB OLR

These are the log records that support HALDB OLR.

### **Related Reading:**

- See the topic entitled, "Logging Effects of Reorganization" on page 11 for more important information on performance and logging of HALDB OLR.
- See the *IMS Version 9: Diagnosis Guide and Reference* for more details on these log records.

---

### DFSLOG - Database Change Log Records

During a HALDB OLR, all segments that are copied to the output data sets are logged in X'50' log records for recovery purposes.

DDSID-- Existing data set ID (DCB number). This definition is added within the existing data set ID field:

#### **DDSIDAM**

When this high-order bit (the X'80' bit) is on, then this DCB number represents one of the M-through-V or Y data sets. When this high-order bit (the X'80' bit) is off, then this DCB number represents one of the A-through-J or X data sets.

---

### DFSLOG20 - Database Open and Close (X'20' and X'21') Log Records

DOCLOGDC-- Existing data set ID (DCB number). This definition is added within the existing data set ID field:

#### **DOCLOGAM**

When this high-order bit (the X'80' bit) is on, then this DCB number represents one of the M-through-V or Y data sets. When this high-order bit (the X'80' bit) is off, then this DCB number represents one of the A-through-J or X data sets.

---

### DFSLOG27 - Database Extension (X'27) Log Records

DFSLOG27 - Database Extension Log Records. When the high-order bit (the X'80' bit) is on in the DBXDCBN field, then the DCB number represents one of the M-through-V or Y data sets. When this high-order bit is off, then the DCB number represents one of the A-through-J or X data sets.

---

### DFSLOG53 - CI/CA and Space Manager (X'53) Log Records

DFSLOG53 - CI/CA and Space Manager Log Records. When the high-order bit (the X'80' bit) is on in the SPLDCB field, then the DCB number represents one of the M-through-V or Y data sets. When this high-order bit is off, then the DCB number represents one of the A-through-J or X data sets.

---

### Log Records: X'29' OLR Event Log Records

Several X'29' log records reflect the progress of a HALDB OLR.

These log records are created at the indicated points in the online reorganization process.

Table 22. X'29' OLR Event Log Record Subtypes

| TYPE or<br>SUB-TYPE | WHY WRITTEN   | ISSUING<br>MODULE    |
|---------------------|---|----------------------|
| X'2910'             | Ownership of the reorganization for a partition was established through DBRC.   | DFSORP60             |
| X'2920'             | The UPDATE OLREORG command updated either the RATE option or the [NO]DEL option for a HALDB partition.  | DFSORC00             |
| X'2930'             | The output data sets were successfully validated or created. One record includes all output data sets.  | DFSORA00<br>DFSRDBL0 |
| X'2940'             | OLREORG Cursor active. Initialization of the reorganization of the partition was completed successfully, two sets of data sets exist, and copying is about to begin. The reorganization is in a cursor-active status. | DFSORP60             |
| X'2950'             | The cursor was updated, but the unit of reorganization was not committed.   | DFSORP20             |
| X'2970'             | OLREORG Cursor inactive. Copying from the input to the output data sets has completed. The output data sets become active, and the input data sets become inactive.   | DFSORP60             |
| X'2990'             | Ownership of the reorganization for a partition was relinquished. This is followed by the X'07' log record for OLREORG Itask termination.   | DFSORP60             |

### **X'2910': Ownership Established**

Ownership of the online reorganization for a partition was established through DBRC.

### **X'2920': UPDATE OLREORG Command**

The UPDATE OLREORG command was processed. The X'2920' log record is written once for each HALDB partition affected by the UPDATE OLREORG command.

### **X'2930': Output Data Set Information**

The output data sets have been successfully validated or created. This X'2930' log record contains various characteristics of all of the output data sets, both those that were preexisting and those that were automatically created. There is enough information to recreate any of these output data sets.

### **X'2940': Cursor-Active Status**

Initialization of the reorganization of the partition was completed successfully, two sets of data sets exist, and copying is about to begin. The reorganization was recorded through DBRC as being in a cursor-active status.

### **X'2950': Cursor Movement**

The cursor was updated. The X'2950' log record appears before the X'3730' log record that indicates that a unit of reorganization was committed.

### **X'2970': Cursor-Active Status Reset**

Copying from the input to the output data sets has completed. The reorganization was recorded through DBRC as no longer being in a cursor-active status.

### **X'2990': Ownership Relinquished**

Ownership of the reorganization for a partition was relinquished through DBRC.

## Chapter 9. RECON Listing HALDB OLR Support

These are the LIST command output (RECON) records that support HALDB OLR.

### HALDB MASTER RECORD

Here is an example of a HALDB Master record listing. The flag for OLREORG CAPABLE = YES | NO is included in the HALDB master DB record and in the partition DB record. This output shows the counters for online reorganization activity not yet reflected in the databases at the RSR tracking site. These counters at the RSR active site are always zero.

```

-----
DB
DBD=DBHIDJ03                DMB#=2    CHANGE#=2    TYPE=HALDB
SHARE LEVEL=3                GSGNAME=**NULL**
PSNAME=**NULL**            CURRENT PARTITION ID=00004
FLAGS:                        COUNTERS:
  RECOVERABLE                =YES          PARTITIONS                =4
  ONLINE REORG CAPABLE       =YES          DATA SET GROUP MEMBERS  =1
  
```

### PARTITION DB RECORD

Here is an example of a HALDB Partition record listing, showing the DB record fields for HALDB OLR. Also included in the listing of TYPE=PART databases is information from the partition record (DSPPTNRC). The flag for OLREORG CAPABLE = YES | NO is included in the master HALDB and in the DBD partition record.

This output shows the counters for online reorganization activity not yet reflected in the databases at the tracking site.

These counters at the active site are always zero.

The OLRIMSID field reflects the name of the IMS that owns the online reorganization. OLRIMSID is always null at the tracking site.

The OLREORG CURSOR ACTIVE flag is on when the paired sets of data sets are active. At the tracking site, the flag is on when the output data sets have been validated or created and that validation or creation process occurs within a hardened milestone.

If the OLR cursor is active, ACTIVE DBDS=A-J and M-V or ACTIVE DBDS=M-V and A-J is displayed. The input data set is listed first and the output data set is listed second.

```

-----
DB
DBD=DBHIDJ03                DMB#=2    CHANGE#=2    TYPE=HALDB
SHARE LEVEL=3                GSGNAME=**NULL**
PSNAME=**NULL**            CURRENT PARTITION ID=00004
FLAGS:                        COUNTERS:
  RECOVERABLE                =YES          PARTITIONS                =4
  ONLINE REORG CAPABLE       =YES          DATA SET GROUP MEMBERS  =1
  
```

Partition DB record:

```

DB
DBD=POHIDJA  MASTER DB=DBHIDJ03  IRLMID=**NULL  CHANGE#=2    TYPE=PART
USID=0000000003  AUTHORIZED USID=0000000003  HARD USID=0000000003
  
```



DSN PREFIX=IMSTESTS.DBOHIDK5 PARTITION ID=00001  
 PREVIOUS PARTITION=\*\*NULL\*\* NEXT PARTITION=POHIDKB  
 OLRIMSID=\*\*NULL\*\* ACTIVE DBDS=M-V

FREE SPACE:  
 FREE BLOCK FREQ FACTOR=0 FREE SPACE PERCENTAGE=0

PARTITION HIGH KEY/STRING (CHAR): (LENGTH=5 )  
 K2000  
 PARTITION HIGH KEY/STRING (HEX):  
 D2F2F0F0F040

OSAM BLOCK SIZE:  
 A = 4096  
 B = 4096

|                        |      |                          |    |
|------------------------|------|--------------------------|----|
| FLAGS:                 |      | COUNTERS:                |    |
| BACKOUT NEEDED         | =OFF | RECOVERY NEEDED COUNT    | =0 |
| READ ONLY              | =OFF | IMAGE COPY NEEDED COUNT  | =0 |
| PROHIBIT AUTHORIZATION | =OFF | AUTHORIZED SUBSYSTEMS    | =1 |
|                        |      | HELD AUTHORIZATION STATE | =6 |
| DATABASE LEVEL TRACK   | =YES | EEQE COUNT               | =0 |
| TRACKING SUSPENDED     | =NO  | RECEIVE REQUIRED COUNT   | =0 |
| OFR REQUIRED           | =NO  | OLR ACTIVE HARD COUNT    | =0 |
| PARTITION INIT NEEDED  | =NO  | OLR INACTIVE HARD COUNT  | =0 |
| OLREORG CURSOR ACTIVE  | =NO  |                          |    |
| PARTITION DISABLED     | =NO  |                          |    |
| ONLINE REORG CAPABLE   | =YES |                          |    |

ASSOCIATED SUBSYSTEM INFORMATION:

|        |                 |         |            |           |
|--------|-----------------|---------|------------|-----------|
|        |                 | ENCODED | B/O NEEDED |           |
| -SSID- | -ACCESS INTENT- | -STATE- | -COUNT-    | -SS ROLE- |
| IMS1   | UPDATE          | 6       | 0          | ACTIVE    |

The DBDS record now uses the OTHER DDN= xxxxxxx field. If a partner DBDS exists, this OTHER DDN= xxxxxxx field points to it.

DBDS  
 DSN=IMSTESTS.DBOHIDK5.A00001 TYPE=PART  
 DBD=POHIDKA DDN=POHIDKAA DSID=001 DBORG=HIDAM DSORG=OSAM  
 CAGRP=\*\*NULL\*\* GENMAX=2 IC AVAIL=0 IC USED=1 DSSN=00000004  
 NOREUSE RECOVPD=0 **OTHER DDN=POHIDKAM**  
 DEFLTJCL=\*\*NULL\*\* ICJCL=PI CJCL OICJCL=POICJCL RECOVJCL=PRECOJCL  
 RECVJCL=PRECVJCL  
 FLAGS: COUNTERS:  
 IC NEEDED =OFF  
 IC RECOMMENDED =ON  
 RECOV NEEDED =OFF  
 RECEIVE NEEDED =OFF EEQE COUNT =0

DBDS  
 DSN=IMSTESTS.DBOHIDK5.M00001 TYPE=PART  
 DBD=POHIDKA DDN=POHIDKAM DSID=001 DBORG=HIDAM DSORG=OSAM  
 CAGRP=\*\*NULL\*\* GENMAX=2 IC AVAIL=0 IC USED=0 DSSN=00000004  
 NOREUSE RECOVPD=0 **OTHER DDN=POHIDKAA**  
 DEFLTJCL=\*\*NULL\*\* ICJCL=PI CJCL OICJCL=POICJCL RECOVJCL=PRECOJCL  
 RECVJCL=PRECVJCL  
 FLAGS: COUNTERS:  
 IC NEEDED =OFF  
 IC RECOMMENDED =ON  
 RECOV NEEDED =OFF  
 RECEIVE NEEDED =OFF EEQE COUNT =0

REORG  
 RUN = 03.287 13:16:25.3 \* USID = 0000000004  
 STOP = 03.287 13:16:26.4 ONLINE RECOV = YES

---

## REORG RECORD

Here is an example of a REORG record showing a HALDB OLR. STOP= is only present for online reorganizations. STOP= is zero for online reorganizations that have not yet completed.

```
-----  
REORG  
RUN      = 2002.365 11:59:12.0 -08:00    * RUN USID      = 0000000037  
STOP     = 2003.001 12:34:56.7 -08:00    ONLINE RECOV=YES  
-----
```

---

## 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:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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 is for planning purposes only. The information herein is subject to change before the products described become available.

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. 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 IBM's application programming interfaces.



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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

---

## Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

|             |       |
|-------------|-------|
| BookManager | IMS   |
| CICS®       | MVS   |
| DFSMS       | RACF  |
| IBM         | z/OS® |

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

Other company, product, and service names, which may be denoted by a double asterisk (\*\*), may be trademarks or service marks of others.



## Bibliography

This bibliography includes all the publications cited in this book, including the publications in the IMS library.

- *DFSMS/MVS® DFSMSdfp™ Advanced Services*
- *DFSMS/MVS Managing Catalogs*
- *z/OS MVS Programming: Authorized Assembler Services Guide*
- *z/OS MVS Programming: Authorized Assembler Services Reference, Vol 2*
- *z/OS MVS Programming: Authorized Assembler Services Reference, Vol 4*
- *z/OS MVS System Messages, Vol 6 (GOS-IEA)*
- *z/OS MVS System Messages, Vol 8 (IEF-IGD)*

### IMS Version 9 Library

|           |        |   |
|-----------|--------|---|
| ZES1-2330 | ADB    | <i>IMS Version 9: Administration Guide: Database Manager</i>                      |
| ZES1-2331 | AS     | <i>IMS Version 9: Administration Guide: System</i>                                |
| ZES1-2332 | ATM    | <i>IMS Version 9: Administration Guide: Transaction Manager</i>                   |
| ZES1-2333 | APDB   | <i>IMS Version 9: Application Programming: Database Manager</i>                   |
| ZES1-2334 | APDG   | <i>IMS Version 9: Application Programming: Design Guide</i>                       |
| ZES1-2335 | APCICS | <i>IMS Version 9: Application Programming: EXEC DLI Commands for CICS and IMS</i> |
| ZES1-2336 | APTM   | <i>IMS Version 9: Application Programming: Transaction Manager</i>                |
| ZES1-2337 | BPE    | <i>IMS Version 9: Base Primitive Environment Guide and Reference</i>              |
| ZES1-2338 | CR     | <i>IMS Version 9: Command Reference</i>   |
| ZES1-2339 | CQS    | <i>IMS Version 9: Common Queue Server Guide and Reference</i>                     |
| ZES1-2340 | CSL    | <i>IMS Version 9: Common Service Layer Guide and Reference</i>                    |
| ZES1-2341 | CG     | <i>IMS Version 9: Customization Guide</i>   |
| ZES1-2342 | DBRC   | <i>IMS Version 9: DBRC Guide and Reference</i>                                    |
| ZES1-2343 | DGR    | <i>IMS Version 9: Diagnosis Guide and Reference</i>                               |

|           |        |  |
|-----------|--------|--|
| ZES1-2344 | FAST   | <i>IMS Version 9: Failure Analysis Structure Tables (FAST) for Dump Analysis</i> |
| ZES1-2346 | OLR    | <i>IMS Version 9: HALDB Online Reorganization Guide and Reference</i>            |
| ZES1-2347 | JGR    | <i>IMS Version 9: IMS Java Guide and Reference</i>                               |
| ZES1-2348 | IIV    | <i>IMS Version 9: Installation Volume 1: Installation Verification</i>           |
| ZES1-2349 | ISDT   | <i>IMS Version 9: Installation Volume 2: System Definition and Tailoring</i>     |
| ZES1-2350 | INTRO  | <i>IMS Version 9: An Introduction to IMS</i>                                     |
| ZES1-2351 | MIG    | <i>IMS Version 9: Master Index and Glossary</i>                                  |
| ZES1-2352 | MC1    | <i>IMS Version 9: Messages and Codes, Volume 1</i>                               |
| ZES1-2353 | MC2    | <i>IMS Version 9: Messages and Codes, Volume 2</i>                               |
| ZES1-2354 | OTMA   | <i>IMS Version 9: Open Transaction Manager Access Guide and Reference</i>        |
| ZES1-2355 | OG     | <i>IMS Version 9: Operations Guide</i>   |
| GC17-7831 | RPG    | <i>IMS Version 9: Release Planning Guide</i>                                     |
| ZES1-2358 | URDBTM | <i>IMS Version 9: Utilities Reference: Database and Transaction Manager</i>      |
| ZES1-2359 | URS    | <i>IMS Version 9: Utilities Reference: System</i>                                |

#### Supplementary Publications

|           |     |   |
|-----------|-----|---|
| GC17-7825 | LPS | <i>IMS Version 9: Licensed Program Specifications</i> |
| ZES1-2357 | SOC | <i>IMS Version 9: Summary of Operator Commands</i>    |

#### Publication Collections

|           |                 |  |
|-----------|-----------------|--|
| LK3T-7213 | CD              | IMS Version 9 Softcopy Library   |
| LK3T-7144 | CD              | IMS Favorites  |
| LBOF-7789 | Hardcopy and CD | Licensed Bill of Forms (LBOF): IMS Version 9 Hardcopy and Softcopy Library |
| SBOF-7790 | Hardcopy        | Unlicensed Bill of Forms (SBOF): IMS Version 9 Unlicensed Hardcopy Library |
| SK2T-6700 | CD              | OS/390 Collection  |
| SK3T-4270 | CD              | z/OS Software Products Collection  |

**Publication Collections**

| SK3T-4271 DVD z/OS and Software Products  
| DVD Collection  
|

**Accessibility Titles Cited in this Book**

| SA22-7787 z/OS V1R1.0 TSO Primer  
| SA22-7794 z/OS V1R1.0 TSO/E User's  
| Guide  
| SC34-4822 z/OS V1R1.0 ISPF User's  
| Guide, Volume 1  
|

# Index

## Special characters

/DBRecovery DB command 56  
 /DISPLAY DB command 40  
 /DISPLAY DB OLR command 40  
 /EREstart command 11  
 /INITIATE OLREORG 25  
 /RECover command 56  
 /TERMINATE OLREORG command 45  
 /UPDATE OLREORG command 51

## A

A-through-J (and X) data sets 2, 4  
 abends 76  
 ABORT option 76  
 acknowledgement, sharing system 24  
 ACS 22  
 active set of data sets xi, 1, 4  
 ALL parameter 36  
 AOI application 25  
 API, OM 25  
 APPC/OTMA programs 25  
 automatic class selection 22  
 automatically-created data sets 22

## B

backup, database 12  
 block size, OSAM 80  
 buffering, sequential 12  
 BYTES parameter 36

## C

capable, HALDB OLR 7  
 CHANGE.DB command 56  
 CIMS class 25  
 codes  
   completion, INITIATE OLREORG 29  
   completion, QUERY OLREORG 38  
   completion, TERMINATE OLREORG 48  
   completion, UPDATE OLREORG 54  
   reason, INITIATE OLREORG 29  
   reason, QUERY OLREORG 38  
   return, INITIATE OLREORG 29  
   return, QUERY OLREORG 38  
 coexistence SPE 7  
 command recognition character 25  
 command security exit routine, DFSCCMD0 25  
 commands  
   /DBRecovery DB 56  
   /DISPLAY DB 40  
   /DISPLAY DB OLR 40  
   /EREstart 11  
   /INITIATE OLREORG 25  
   /RECover 56  
   /TERMINATE OLREORG 45

commands (*continued*)  
 /UPDATE OLREORG 51  
 CHANGE.DB 56  
 CHANGE.DB ALL OLRCAP 7  
 database 9  
 format  
   INITIATE OLREORG 25  
   QUERY OLREORG 35  
   TERMINATE OLREORG 45  
   UPDATE OLREORG 51  
 formats 25  
 GENJCL.IC 15  
 GENJCL.OIC 15  
 INIT.DB 31  
 INITIATE OLREORG  
   NAME keyword 27  
   OPTION keyword 27  
   output fields 28  
   SET keyword 27  
 NOTIFY.IC 15  
 NOTIFY.REORG 32  
 NOTIFY.UIC 15  
 QUERY DB STATUS(OLR) 42  
 QUERY OLREORG 35  
   NAME keyword 36  
   output fields 37  
   SHOW keyword 36  
   STATUS keyword 37  
 security 25  
 slash 25  
 START 9  
 STOP 9  
 TERMINATE OLREORG 43  
 the types of 25  
 type-2 25  
 UPDATE OLREORG 49  
   DEL option 4  
   NODEL option 49  
   RATE option 49  
 usage  
   INITIATE OLREORG 26  
   QUERY OLREORG 36  
 completion codes, INITIATE OLREORG 29  
 completion codes, QUERY OLREORG 38  
 completion codes, TERMINATE OLREORG 48  
 completion codes, UPDATE OLREORG 54  
 completion of HALDB OLR 24  
 convention, data set naming 4  
 copying, database records 2  
 CRC 25  
 CRC commands 25  
 cursor xi, 2  
 cursor-active flag 79  
 cursor-active status xi  
 cursor-active status, OLR 1

**D**

data sets  
 A-through- J (and X) 2  
 A-through-J 4  
 active set xi, 4  
 automatically-created 22  
 creating output 21  
 deletion of 24  
 ILDS 5  
 inactive set 4  
 inactive sets of, xi  
 M-through-V 4  
 M-through-V (and Y) 2  
 multiple-volume 23  
 naming convention 4  
 non-SMS-managed 22  
 OSAM 21  
 partition ID 5  
 PHDAM 1  
 PHIDAM 1  
 pre-existing 21  
 primary index 5  
 PRIMARY INDEX 21  
 single volume 22  
 SMS-managed 22  
 VSAM 21

data sharing 7  
 database change log records 77  
 Database Recovery utility 12  
 DB commands 9  
 DDSID 77  
 DDSIDAM 77  
 DEL option  
 UPDATE OLREORG 4  
 DEL parameter 28  
 deletion of data sets 24  
 DFS messages 59  
 DFSCCMD0 25  
 DFSCCMD0 command security exit routine 25  
 DFSDLOG 77  
 DFSORA00 78  
 DFSORC00 78  
 DFSORP20 78  
 DFSORP60 78  
 DFSPREC0 15  
 DFSRDBL0 78  
 DFSVSMxx SBONLINE statement 12  
 DFSVSNxx 10  
 DLISAS 1  
 DOCLOGAM 77  
 DOCLOGDC 77  
 DSP messages 74

**E**

ECMS console 25  
 EMCS console 25  
 environments  
 INITIATE OLREORG command 26  
 QUERY OLREORG command 35

environments (*continued*)  
 TERMINATE OLREORG command 45  
 UPDATE OLREORG command 51  
 error  
 data set allocation 24  
 initialization 2  
 ESDS 21  
 example  
 OM API  
 TERMINATE OLREORG 48  
 type-1  
 /TERMINATE OLREORG 48

**F**

failure  
 initialization 2, 24  
 fallback 8  
 format  
 INITIATE OLREORG command 25  
 QUERY OLREORG command 35  
 TERMINATE OLREORG command 45  
 UPDATE OLREORG command 51  
 formats of commands 25

**G**

GENJCL.IC command 15  
 GENJCL.OIC command 15  
 GUARANTEED SPACE attribute 23

**H**

HALDB master 57  
 HALDB Master record 79  
 HALDB OLR xi, 1  
 completion of 24  
 HALDB OLREORG xi  
 hardened data 79  
 how to use  
 INITIATE OLREORG 26  
 QUERY OLREORG 36

**I**

ILDS data set 5  
 Image Copy, Database 12  
 inactive set of data sets 1, 4  
 inactive sets of data sets xi  
 INIT.DB command 31  
 initialization 1  
 INITIATE OLREORG  
 completion codes 29  
 INITIATE OLREORG command 25  
 output fields 28  
 INITIATE OLREORG command usage 26  
 input data sets 2

**K**

keyword  
 synonyms 26, 36, 43, 45  
 keywords  
 MAXSB 12  
 NAME 27, 36  
 OPTION 27  
 reserved 26, 35, 45  
 SET 27  
 SHOW 36  
 STATUS 37  
 KSDS 21

**L**

LIST command output 79  
 listing  
 HALDB Master record 79  
 Partition record 79  
 REORG record 82  
 locking, global 9  
 log records  
 database change 77  
 x'20' 77  
 x'21' 77  
 x'29' 77  
 x'2910' 78  
 x'2920' 78  
 x'2930' 78  
 x'2940' 78  
 x'2950' 78  
 x'2970' 78  
 x'2990' 78  
 x'3730' 78  
 x'50' 77  
 log-only recovery 12  
 logging volume 11

**M**

M-through-V (and Y) data sets 2, 4  
 MAXSB keyword 12  
 messages  
 DFS series 59  
 DSP series 74  
 HALDB OLR 59  
 migrating 7  
 milestone 79  
 multiple-volume data sets 23

**N**

NAME keyword 27, 36  
 naming convention, data set 4  
 NODEL option  
 UPDATE OLREORG 49  
 NODEL parameter 28  
 non-SMS-managed data sets 22  
 NOOWNER parameter 57  
 NORECOV parameter 33

normal process, HALDB OLR 3  
 NOTIFY.IC command 15  
 NOTIFY.REORG command 32  
 NOTIFY.UIC command 15  
 NOTOWNED parameter 37

**O**

OFFLINE parameter 33  
 offline reorganization 8  
 OLR xii  
 OLRCAP parameter 32, 57  
 OLRDBDS parameter 57  
 OLRIMSID parameter 57, 79  
 OLRNOCAP parameter 32, 57  
 OLRRGOFF parameter 57  
 OLRRGON parameter 57  
 OM API 25  
 ONLINE parameter 33  
 online reorganization xii  
 See HALDB OLR  
 OPTION keyword 27  
 OSAM block size 80  
 OSAM data sets 21  
 output  
 HALDB Master record 79  
 LIST command 79  
 Partition record 79  
 REORG record 82  
 output data sets 2  
 output fields  
 INITIATE OLREORG command 28  
 QUERY OLREORG command 37  
 owned xii  
 OWNED parameter 37  
 ownership, HALDB OLR 1

**P**

paired sets of data sets 2  
 parameters  
 ALL 36  
 BYTES 36  
 DEL 28  
 NODEL 28  
 NOOWNER 57  
 NORECOV 33  
 NOTOWNED 37  
 OFFLINE 33  
 OLRCAP 32, 57  
 OLRDBDS 57  
 OLRIMSID 57, 79  
 OLRNOCAP 32, 57  
 OLRRGOFF 57  
 OLRRGON 57  
 ONLINE 33  
 OWNED 37  
 PARTNAME 27, 36  
 RATE 10, 11, 27, 36  
 RECOV 33  
 RUNNING 37

parameters (*continued*)  
 STATUS 37  
 STOPTIME 33  
 UNKNOWN 37  
 PARTITION HIGH KEY 80  
 partition ID 5  
 Partition record listing 79  
 PARTITION SELECTION STRING 80  
 PARTNAME parameter 27, 36  
 performance 9  
 PHDAM data sets 1  
 PHIDAM data sets 1  
 primary index 5  
 PRIMARY INDEX data set 21  
 primary space allocation 23  
 process  
 copying 2  
 HALDB OLR, normal 3  
 processing  
 UPDATE OLREORG 49  
 processing time 11

## Q

QUERY DB STATUS(OLR) command 42  
 QUERY OLREORG  
 completion codes 38  
 QUERY OLREORG command 35  
 output fields 37

## R

RACF 25  
 RANDOMIZER values 80  
 RATE option  
 UPDATE OLREORG 49  
 RATE parameter 10, 11, 27, 36  
 reason codes  
 INITIATE OLREORG 29  
 QUERY OLREORG 38  
 reclustering of data 1  
 RECON records 79  
 RECOV parameter 33  
 recovery, database 12  
 recovery, log-only 12  
 REORG record listing 82  
 reserved keywords 26, 35, 45  
 restart, system 9, 10  
 restrictions of HALDB OLR 8  
 resume online reorganization 25  
 return codes  
 INITIATE OLREORG 29  
 QUERY OLREORG 38  
 RSR 10  
 RUNNING parameter 37

## S

SBONLINE statement 12  
 secondary space allocation 23  
 security 25

sequential buffering 12  
 SET keyword 27  
 SHAREPLEX, IMS xii  
 sharing system acknowledgement 24  
 SHOW keyword 36  
 single volume data sets 22  
 slash commands 25  
 SMS-managed data sets 22  
 space allocation  
 primary 23  
 secondary 23  
 STATUS keyword 37  
 status of OLR, cursor-active 1  
 STATUS parameter 37  
 status, cursor-active xi  
 STOP= 82  
 STOPTIME 33  
 storage class 23  
 storage volumes 22  
 syntax  
 INITIATE OLREORG command 25  
 TERMINATE OLREORG command 45  
 UPDATE OLREORG command 51  
 system restart 9, 10

## T

TERMINATE OLREORG  
 completion codes 48  
 TERMINATE OLREORG command 43  
 TERMINATE OLREORG command ABORT option 76  
 time, processing 11  
 type-1 commands 25  
 type-2 commands 25  
 type=HALDB 57

## U

unit of reorganization xii, 2  
 UNKNOWN parameter 37  
 UPDATE OLREORG  
 completion codes 54  
 UPDATE OLREORG command 49  
 UPDATE OLREORG command processing 49  
 upgrade, DBRC RECON 7  
 usage  
 INITIATE OLREORG command 26  
 QUERY OLREORG command 36  
 utilities  
 Database Change Accumulation 13, 19  
 Database Image Copy 14  
 Database Recovery 12, 20  
 HALDB Index/ILDS Rebuild 19

## V

values, RANDOMIZER 80  
 volume, logging 11  
 volumes, storage 22  
 VSAM data sets 21



## X

x'20' log record 77  
x'21' log record 77  
x'29' log record 77  
x'2910' log record 78  
x'2920' log record 78  
x'2930' log record 78  
x'2940' log record 78  
x'2950' log record 78  
x'2970' log record 78  
x'2990' log record 78  
XRF 9, 10







Program Number: 5655-J38

IBM Confidential  
Printed in USA

ZES1-2346-01



Spine information:



IMS

HALDB Online Reorg. Guide and Reference

Version 9