Pioneer

IBM

HALDB Online Reorganization Guide and Reference
(Build Date: 10/06/03)

Version 9

IBM

HALDB Online Reorganization Guide and Reference
(Build Date: 10/06/03)

Version 9

Quality Partnership Program (QPP) Edition (January 2004)

This QPP edition applies to Version 9 of IMS (product number 5655-J38) and to all subsequent releases and modifications until otherwise indicated in new editions.

Abstract

This book describes the HALDB Online Reorganization process. Use of this process allows for non-disruptive online reorganization for HALDB PHDAM and PHIDAM partitions.

iii

Contents

v

IBM Confidential

Figures

Tables

ix

About This Book

This new information unit is intended to supply you with as much IMS HALDB Online Reorganization information as possible in a single place.

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

This book is available on the IMS Version 9 Licensed Product Kit (XXXX-XXXX).

To get the most current versions of the PDF and BookManager formats, go to the IMS Web site at www.ibm.com/ims and link to the Library page.

Summary of Contents

This information unit is presented in 3 parts:
- Part 1 contains concepts including an overview of how HALDB Online Reorganization functions and the considerations for using HALDB OLR.
- Part 2 contains HALDB OLR usage information including general IMS commands, HALDB OLR-specific commands, DBRC commands, and RECON records.
- Part 3 contains HALDB OLR diagnostic information including IMS messages, codes, abends, and log records.

Prerequisite Knowledge

To use HALDB OLR support you need a working knowledge of database management, HALDBs, and DBRC.

An understanding of the HALDB Online Reorganization terminology is needed. See "Terminology of HALDB OLR" on page xii for more information.

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 37 or "CHANGE.DB" on page 62..

HALDB OLR can be initiated and managed by issuing the INITIATE, TERMINATE, UPDATE OLREORG commands. These commands are described in Chapter 4, "Setting Up HALDB OLR," on page 31 and in Chapter 6, "Modifying and Tuning HALDB OLRs," on page 49.

HALDB OLR can be monitored by issuing the QUERY OLREORG, /DISPLAY OLR, or QUERY DB commands. These commands are described in Chapter 6, "Modifying and Tuning HALDB OLRs," on page 49.

To activate HALDB OLR using commands issued from the OM API, an IMS Common Service Layer which includes Operations Manager (OM) and Structured Call Interface (SCI) is required.

Related Reading:
- Refer to the *IMS Common Service Layer* for more information on the OM and SCI.

- Refer to the *IMS DBRC Guide and Reference* for more information on DBRC.

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 boundary that delineates the end of the database records that have been copied.

cursor-active status
> Paired data sets are active. Cursor 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.

hardened data
> At an RSR tracking site, data that is guaranteed to have been written to DASD in the shadow database data sets up through a particular milestone.

inactive data sets
> The old input data sets that the data was copied from.

milestone
> A marker recognizable by the log router that delineates a point from which a restart can occur in IMS database tracking.

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, IMS 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 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

This book contains syntax diagrams.

Each syntax diagram 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.

Conventions used in syntax diagrams are described in Table 1:

Table 1. How to Read Syntax Diagrams
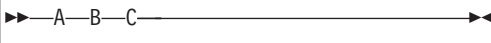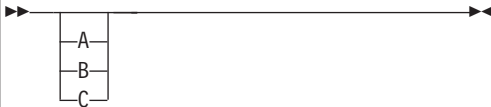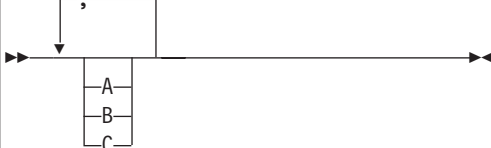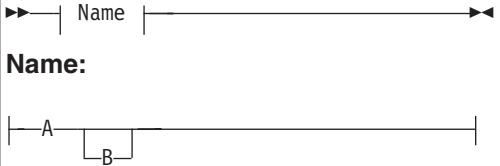
Convention	Meaning
▶▶—A—B—C————————————▶◀	You must specify values A, B, and C. Required values are shown on the main path of a syntax diagram.
▶▶———————————————————▶◀ └A┘	You have the option to specify value A. Optional values are shown below the main path of a syntax diagram.
▶▶—┬A┬————————————————▶◀ ├B┤ └C┘	You must specify value A, B, or C.
▶▶———————————————————▶◀ ┬A┬ ├B┤ └C┘	You have the option to specify A, B, C, or none of these values.
▶▶—┬A┬————————————————▶◀ ├B┤ └C┘	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.
▶▶—┬,┬————————————————▶◀ ┬A┬ ├B┤ └C┘	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.

Table 1. How to Read Syntax Diagrams (continued)

Convention	Meaning
	You have the option to specify value A multiple times. The separator in this example is optional.
  Name: 	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.
Punctuation marks and numbers	Enter punctuation marks (slashes, commas, periods, parentheses, quotation marks, equal signs) and numbers exactly as shown.
Uppercase values	Keywords, their allowable synonyms, and reserved parameters, appear in uppercase letters for z/OS. Enter these values exactly as shown.
Lowercase values without italics	Keywords, their allowable synonyms, and reserved parameters, appear in lowercase letters for UNIX. Enter these values exactly as shown.
Lowercase values in italics (for example, *name*)	Supply your own text or value in place of the *name* variable.
b	A b symbol indicates one blank position.

Other conventions include the following:

- When entering 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 *.

Example Syntax Diagram

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

Name:

Greeting:

```
                        (2)
├──,─your_greeting─────────────────────────────────────────────────┤
```
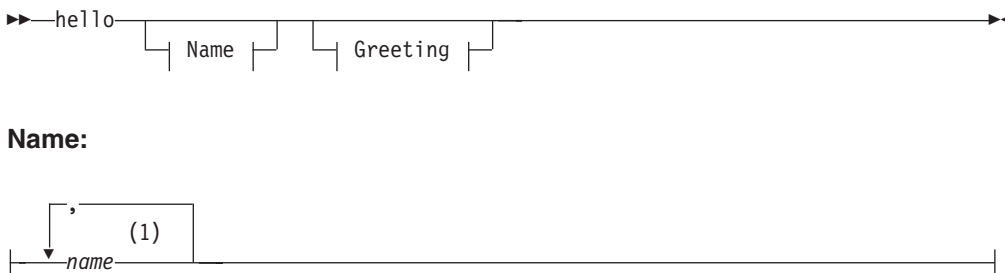
Notes:

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 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*.

Dean, since we coded the comma before the greeting differently from this example, does this last statement still make sense?

How to Send Your Comments

TBD

Change Indicators

Since this is a new information unit. All information herein is new and is presented without change indicators.

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.

Changes to This Book for IMS Version 9

This is a new book that contains new technical information for IMS Version 9.

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

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.

New and Revised Titles

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

-
  The library includes a new book: . This book is available only in PDF and BookManager formats.

-
  The library includes a new book: .

- The book formerly titled *IMS Version 8: IMS Java User's Guide* is now titled .

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 *IMSplex* 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

xvii

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 Documentation

Online documentation 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.

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.

Part 1. Understanding HALDB OLREORG

Chapter 1. Overview of HALDB OLR

HALDB (High-Availability Large Database) Online Reorganization will be known as HALDB OLR from this point on in this document.

The HALDB OLR function provides non-disruptive reorganization of HALDB PHDAM (Partitioned Hierarchical Direct Access Method) and PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) partitions. HALDB OLR provides 100% 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 planned data outage time that to date has constituted the largest amount of time that data was unavailable. HALDB OLR also offers significant benefit toward reducing the cost of owning data (for example: reduction in skills needed to maintain data, less personnel needed, reduction in storage costs, and other costs).

In This Chapter:
- "How HALDB OLR Functions"
  - "When HALDB OLR Begins - Initialization and IMS Subsystem Ownership"
  - "After HALDB OLR Initialization Completes - During the Copying Process" on page 4
  - "After the HALDB OLR Copying Process Completes - **Cursor-Active** Status is Reset" on page 6
- "Naming Conventions for HALDB OLR Data Sets" on page 6
- "DBRC Changes to Support HALDB OLR" on page 8

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 in the DLISAS (Data Language/I separate address space). The online reorganization function occurs non-disruptively while concurrent IMS updates are allowed, small amounts of data are moved from the active input set of data sets to the active output set of data sets and reorganized.

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

When HALDB OLR Begins - Initialization and IMS Subsystem Ownership

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 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 a *cursor-active* status: (CURSOR ACTIVE=YES).

What Happens to Output Data Sets When HALDB OLR Initialization Fails

During initialization, various error conditions, such as an unacceptable preexisting output 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 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 output data sets that were already automatically created are retained along with any preexisting ones.

Preexisting output data sets must meet certain requirements (discussed later in this book in "Required Characteristics for Existing Output Data Sets" on page 27) and any output data sets that do not already exist are created automatically by IMS as part of initialization.

After HALDB OLR Initialization Completes - During the Copying Process

When the *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. Note that during the reorganization, database changes may still be made to the input data sets that have not yet been copied to the output data sets, and to parts of the output data sets to which data has already been copied.

The following diagram, Figure 1 illustrates the conceptual relationship between the input and output data sets at a point during the reorganization.
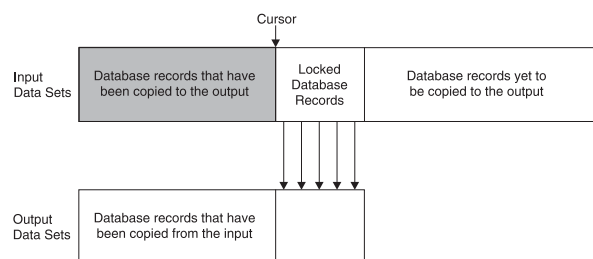
Figure 1. Relationship between Input and Data Sets at a Point During Reorganization

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. This is true even in the input data sets for the database records that are still to be copied; 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.

For a concurrently running application, IMS determines whether to expect a given database record to be in either the input or the output data sets based on whether the record is at a point either beyond or prior to the cursor, but this is completely transparent to an application program accessing the database.

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 on page 6 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.
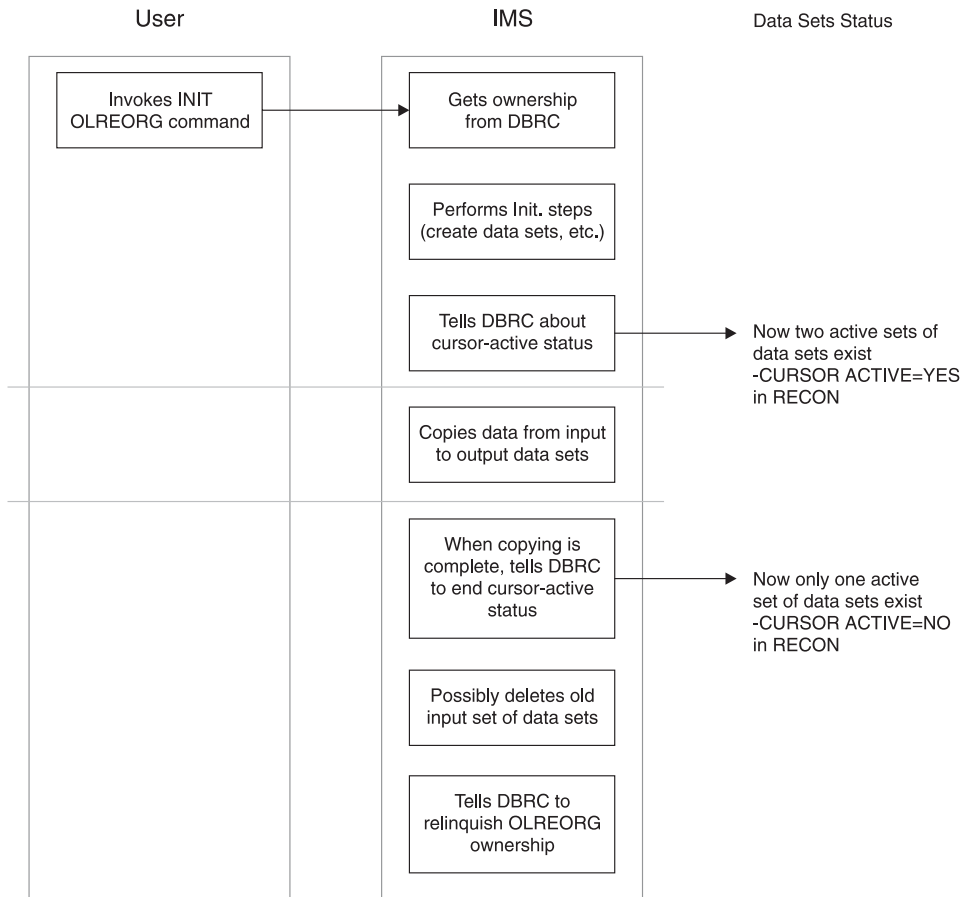
Figure 2. Normal Processing Flow of a HALDB OLR

After the HALDB OLR Copying Process Completes - *Cursor-Active* Status is Reset

After all database records have been copied, 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 DBRC is reset. Unless an initial load or a batch reorganization reload is done, successive online reorganizations for the HALDB partition then alternate between the two sets of data sets.

At successful completion of the reorganization process, an attempt to delete the inactive data sets will occur if the DEL option is in effect. IMS tells DBRC to relinquish ownership of the reorganization.

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 new for this release. Data sets with these names 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 - During the Copying Process" on page 4

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. 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

DBRC Changes to Support HALDB OLR

Changes have been made to some DBRC commands and to some RECON records to support HALDB OLR. Some new RECON records have been added to support HALDB OLR also.

Related Reading:
- For more detailed information on the changed DBRC commands, see Part 2, "Using HALDB OLR," on page 25.
- For more detailed information on the new and changed RECON records, see "RECON Listing HALDB OLR Support" on page 91.

Chapter 2. Considerations

The following sections describe the things that need to be considered for HALDB Online Reorganization support.

In the Chapter:
- "Migration"
- "Coexistence"
- "Fallback" on page 10
- "Restrictions" on page 10
- "Performance" on page 11
- "Recovery" on page 13
- "DBRC Considerations" on page 16
- "Offline Reorganization" on page 20
- "Utilities" on page 20

Migration

The following items list the migration considerations for HALDB Online Reorganization support.
- As part of DBRC RECON upgrade, existing HALDBs are marked to indicate that they are not capable of being reorganized online using HALDB Online Reorganization.
- The CHANGE.DB DBD(HALDB master) OLRCAP command in IMS Version 9 can be used to allow HALDB Online Reorganization.
- IMS Version 8 is allowed full datasharing with HALDBs undergoing online reorganization if the IMS Version 8 OLR Coexistence SPE is applied.
- HALDBs marked as capable of running online reorganization are not known to an IMS Version 8 without the IMS Version 8 OLR Coexistence SPE, or to IMS Version 7.
- IMS Version 8 without the IMS Version 8 OLR Coexistence SPE and IMS Version 7 may access HALDBs marked as not capable of running online reorganization.
- When migrating from a prior release, a `CHANGE.RECON UPGRADE` command is used to upgrade the RECON to an IMS Version 9 RECON. At that time, all existing HALDBs are marked to indicate that they are not capable for HALDB Online Reorganization. Prior to running HALDB Online Reorganization, a `CHANGE.DB DBD(HALDB_master) OLRCAP` command must be issued to change the status of the HALDB and all its partition databases to allow HALDB Online Reorganization.

Coexistence

The following items list the coexistence considerations for HALDB Online Reorganization support.
- IMS Version 7 systems and IMS Version 8 (without the OLR Coexistence SPE applied) may still access the HALDBs defined as not capable of running Online Reorganization. New applications may be created exploiting Online Reorganization capable HALDBs before all IMS systems have been upgraded to IMS Version 9. IMS Version 8 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 Online Reorganization.
- The CHANGE.DB DBD(HALDB master) OLRCAP command in IMS Version 9 can be used to allow HALDB Online Reorganization.
- IMS Version 8 is allowed full datasharing with HALDBs undergoing online reorganization if the IMS Version 8 OLR Coexistence SPE is applied.
- HALDBs marked as capable of running online reorganization are not known to an IMS Version 8 without the IMS Version 8 OLR Coexistence SPE, or to IMS Version 7.
- IMS Version 8 without the IMS Version 8 OLR Coexistence SPE and IMS Version 7 may access HALDBs marked as not capable of running online reorganization.
- To allow full data sharing between IMS Version 8 and IMS Version 9 with HALDBs undergoing online reorganization, the IMS Version 8 OLR Coexistence SPE needs to be applied.

Fallback

The following items describe the fallback considerations for HALDB Online Reorganization support.

For IMS Version 8 with the OLR Coexistence SPE- Full datasharing is allowed and no immediate action is necessary. However, if any M-V and X DBDSs are active or online reorganization is active for any partitions, then the IMS Version 9 Utilities must be run whenever processing those partitions. You will need to run offline reorganization for any partitions with M-V and X DBDSs active or for any partitions with online reorganization active before falling back to use the IMS Version 8 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 the following actions are completed:

- All partitions with the M-V and Y data sets active or online reorganization active must be reorganized using the IMS Version 9 offline reorganization utility (this leaves only the A-J and Y data sets active).
- Change DBRC for the databases to reflect that they are not capable of being reorganized online using IMS Version 9 (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 IMSplex command, an IMS Common Service Layer that includes Operations Manager and Structured Call Interface is required. See the *IMS: Common Service Layer Guide and Reference* for details.
- HALDB Online Reorganization will only run in an `LSO=S` (local storage option-subordinate) environment. An attempt to initiate online reorganization in an `LSO=Y` (local storage option- yes) IMS environment will be rejected.

    For more information about the LSO specification, see the *IMS Installation Volume 2: System Definition and Tailoring*.

- This reorganization process only provides reclustering and space distribution advantages. The online reorganization process does not support data definitional changes.
- The HALDB OLR commands do not support keyword synonyms. Only one version of each keyword can be specified for a given command.
- The following commands are not allowed to process against a HALDB partition on an IMS system while HALDB OLR is running against that partition that IMS.
  - /START DB
  - /DBRECOVERY DB
  - /DBDUMP DB
  - /STOP DB

  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
  - /DBRECOVERY DB
  - /DBDUMP DB

  Related Reading:
  - For more information about message DFS0488I, see Chapter 7, "Messages and Codes of HALDB OLR," on page 67.
  - For more information about these commands see the *IMS Command Reference*.
- Online reorganization cannot be started if it is already running for a particular partition.
- Image copy is not allowed if online reorganization is active.
- Batch backout must be run using a DLI region type to backout in-flight online reorganization work.
- For requirements of existing output data sets, see "Required Characteristics for Existing Output Data Sets" on page 27.
- **NOTE:** You will receive a null response from an Extended Restart Facility (XRF) alternate system and from a Remote Site Recovery (RSR) Tracking system.
- For information on DBRC command restrictions for HALDB OLR, see "GENJCL.IC and GENJCL.OIC Considerations" on page 17 and "NOTIFY.IC and NOTIFY.UIC Considerations" on page 17.

Performance

The following sections describe performance effects of system restart, XRF, and logging.

Restart and XRF Changes

All active HALDB online reorganization tasks are suspended before a shutdown checkpoint. They will be resumed automatically by IMS after the next restart.

All HALDB online reorganization tasks that were active (running) when IMS abnormally terminated will have the new output data set dynamically backed-out to the last sync-point using the same logic that is used to backout online programs.

When backout is complete, the HALDB online reorganization environment will be rebuilt and HALDB online reorganization tasks will be resumed automatically by IMS after the next restart

When an XRF takeover occurs, the HALDB online reorganization tasks will be resumed on the new active IMS. The same resumption is also part of emergency restart and warm start. Partitions undergoing online reorganization are stopped during RSR takeover.

The re-establishment of the HALDB online reorganization tasks will be invoked internally as part of restart or takeover processing.

When an emergency restart or warm start occurs, HALDB OLRs that were active when IMS abnormally terminated or shut down will be automatically resumed on the new active IMS. However, in the case of an RSR, HALDB OLRs that were active at the time of RSR takeover will not be automatically resumed at the takeover /ERestart. 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.

Logging Effects of Reorganization

The logging volume 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 some minor amount of new log records, and because of the prefix size of the x'50' log record, the reorganization of a 2-gigabyte partition would generate considerably more than 2 gigabytes of log data.

The reorganization of multiple partitions in parallel could generate sufficient log data to impact normal transaction processing. The rate of OLDS log switches and their archiving would certainly be affected.

Database Change Log Records
During a HALDB online reorganization, all data movement to the output data sets are tracked in x'50' database change log records for recovery purposes.

New Log Records: x'29' Log Records
Several new x'29' log records reflect the progress of a HALDB online reorganization.

Reorganization Processing Time

An online reorganization's speed and impact on the system can be controlled somewhat by the RATE specified, in the `INITIATE OLREORG` and `UPDATE OLREORG` commands, by system resources, and by system utilization. The amount of updates to the IMS log and system contention can also affect the speed at which an online reorganization runs.

RATE(100) is the default and allows the online reorganization to run as fast as possible (depending on system resources, system contention, logging waits, etc.) with no intentional delays. A RATE of 50 would have intentional delays to cause the reorganization to run twice as long as it would have run at RATE 100. RATE 25 would take 4 times as long, and so on. The RATE can be changed with the `UPDATE OLREORG` command when adjustments are desired depending on system contention, (slow down the online reorganization) or lack of system contention (speed up the online reorganization).

2 Activating Sequential Buffering for OLR

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

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

2 The two forms of the SBONLINE control statement are:

2 `SBONLINE`

2 or

2 `SBONLINE,MAXSB=nnnnn`

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

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

Recovery

The following sections describe the recovery considerations for HALDB Online
Reorganization support.

Database and Data Set

- 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.
  After rebuilding the index and Indirect List data sets, the online reorganization
  may be restarted.

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

Since changes made during online reorganization are logged, database recovery
does not require an image copy as a base for the recovery. Only active DBDSs are
allowed to be recovered unless in they are the middle of an online reorganization.

No Database Image Copy Allowed during Reorganization

During the time that an online reorganization is active for a partition, 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

Once the *must-complete* 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` command for the HALDB partition.

3. Run database change accumulation as necessary. The JCL can be created by the DBRC `GENJCL.RECOV` command, or the Database Change Accumulation Utility can be run from customer-managed 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 customer-managed 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:

7. Enter the `/STA` 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 is similar to the above except that the reorganization doesn't have to be terminated or resumed. Although an output data set can be recovered in this way, it is strongly recommended that an image copy be made of the output data sets as soon as practical after the reorganization. Recovering from this image copy is faster than recovering from the database change records logged during the reorganization.

Recovery of an output data set from other than the beginning of the reorganization, such as from a full dump of a DASD volume, is also possible using existing procedures so long as the reorganization has been either completed or terminated.

For recovery of an input data set before the reorganization completes, the 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 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.

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

Image Copy will always copy 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 need not be aware of which set of data sets is the active set and you need not provide this information for the following 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.

When an emergency restart or warm start of an RSR occurs, HALDB OLRs that were active at the time of RSR takeover will not be automatically resumed at the takeover /EREstart. 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).

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.

Related Reading: See the *IMS: DBRC Guide and Reference* for complete information on these DBRC commands that have changed in relation to online reorganization:

- CHANGE.DB
- GENJCL.IC
- GENJCL.OIC
- INIT.DB
- NOTIFY.IC
- NOTIFY.REORG
- NOTIFY.UIC

Some RECON records that have been changed or added are as follows:

- HALDB Master record

- Partition DB record
- REORG record

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 A-through-J or M-through-V set of data sets.

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 a HALDB OLR is not complete, for example: the online reorganization is active or it is terminated. The NOTIFY.IC and NOTIFY.UIC commands fail if they attempt to record an image copy that falls within the range of a HALDB OLR.

New DBRC Information

The DB record for a HALDB partition contains new information for online reorganization:

- a field that indicates whether the database is capable of being reorganized online.
- a field that indicates whether online reorganization is processing
- the SSID of the IMS currently performing an online reorganization (see LIZ
- a field indicating which set of DBDSs is active for the partition and whether the M-through-V and Y data sets exist
- counters for hardening Online Reorganization 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 '80'x bit on. For RECON list purposes, the '80'x bit is ignored.

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

DSID / DCB Number	A-thru-J, L, and X Data Sets		M-thru-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

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

DSID / DCB Number	A-thru-J, L, and X Data Sets		M-thru-V and Y Data Sets	
	PHIDAM	PHDAM	PHIDAM	PHDAM
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	
14A	J		V	

The M-through-V and Y DBDS records are added to the indicated change accumulation group. For example, 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.

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

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 data set partition- for example A-to-M, B-to-N, C-to-O, and so on. Any of these can be online reorganized at any time and may be active or inactive accordingly. So data sets (or partitions) 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. This 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 must-complete 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 'ONLINE REORG 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 don't 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 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 stoptime
* the DB record for the partition is changed to indicate 'ONLINE REORG 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 is running. DBRC fails the execution of image copy if this is the case.

Since changes made during online reorganization are logged, database recovery does not require an image copy as a base for the recovery. Only active DBDSs are allowed to be recovered unless they are in the middle of an online reorganization.

Offline Reorganization

Offline reorganization can only be run if online reorganization is either not currently running or 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 considerations for IMS utilities: Change Accumulation, Image copy, and Recovery.

Change Accumulation - DFSCUM0

DFSUCUM0 accumulates changes for HALDB partition A-through-J data sets and for M-through-V data sets. A DB0 control 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 header records (type 25 records) may be created for a corresponding A-J and M-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.

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.

Recovery - DFSURDB0

Once the *must-complete* 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 created. 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, it has been automatically stopped. Otherwise, the `TERMINATE OLREORG` command may have to be used to temporarily stop a currently running reorganization.

2. Enter the `/DBR` command for the HALDB partition.

3. Run database change accumulation as necessary. The JCL can be created by the DBRC `GENJCL.RECOV` command, or the Database Change Accumulation Utility can be run from customer-managed 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 customer-managed JCL. In this case, the DD statement for DFSUDUMP must specify DUMMY indicating that there is no image copy from which to restore.

6. Execute the Batch Backout utility, because backout might be required to backout uncommitted data. Note that the operation 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:

7. Enter the `/STA` 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 is similar to the procedure above except that the reorganization doesn't have to be terminated or resumed. Although an output data set can be recovered by this method, it is *strongly recommended* that an image copy be made of the output data sets as soon as is practical after the reorganization. Recovering from this image copy is faster than recovering from the database change records logged during the reorganization.

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

For recovery of an input data set before the reorganization completes, the only additional thing to consider is: to ensure that the reorganization has been terminated before starting the recovery.

Image Copy - DFSUDMP0

Image Copy automatically copies the HALDB partition A-through-J data sets or M-through-V data sets internally.

Image Copy processing has changed in this way for HALDBs: An image copy of a DBDS cannot be taken while HALDB online reorganization is being run on that DBDS. DBRC fails the execution of an image copy in this case.

Image Copy inputs are unchanged. For each request to copy HALDB partition A-through-J data sets or M-through-V data sets, Image Copy issues an error message if HALDB online reorganization is currently active on the HALDB partition. 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.

Generic DDNAME for Database Image Copy

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 the following 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.

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

When online reorganization is running, 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.

Part 2. Using HALDB OLR

25

Chapter 3. The Data Sets of HALDB OLR

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

Restriction: The data sets used by HALDB OLR must of 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 be pre-existing, while the *N* and *O* data sets need to be created by IMS, and the *P* data set might also be pre-existing. So, in this case, IMS creates the *N* and *O* data sets and uses the pre-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 67 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 above under "VSAM:"

Characteristics for Automatically-Created Output Data Sets

The following paragraphs 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)
```

As already stated, a non-SMS-managed VSAM output data set is created on the same single volume as the corresponding input data set.

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 manners, 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 27 and "Characteristics for Automatically-Created Output Data Sets" on page 28 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 condition 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 (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 deletes 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 part of the base of IMS Version 9, so no special installation process is required in order to use it.

There are now 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, ECMS consoles, and APPC/OTMA programs for a single database locally.
- 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 user-defined exit routine, DFSCCMD0. The CIMS class for RACF security 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.

Use the `INITIATE OLREORG` or the `/INITIATE OLREORG` command to initiate or resume HALDB OLR. For more information, see "INITIATE OLREORG Command."

INIT.DB has parameters that you can use to indicate whether HALDB OLR is allowed to run on a database. For more information, see "INIT.DB" on page 37.

NOTIFY.REORG has parameters that you can use to specify whether the reorganization is a HALDB OLR or it is an offline reorganization. If it is a HALDB OLR (ONLINE), whether the HALDB OLR may be used as input for recovery and the time that the HALDB OLR successfully completed. For more information, see "NOTIFY.REORG" on page 38.

INITIATE OLREORG Command

The INTITIATE OLREORG Type-2 command has been developed to support HALDB OLR. This command is issued and processed through the OM (Operations Manager) API (Application Programming Interface).

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

The INITIATE OLREORG command can also be issued as a Type-1 command from terminals, AOI applications, ECMS consoles, and APPC/OTMA programs to initiate or resume a HALDB OLR for a single database locally.

/INITIATE OLREORG Command

The Type-1 command syntax is the same as the IMSplex command syntax issued from OM API 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.

2 Format of the INITIATE OLREORG Command

2   The syntax diagram below illustrates the format of the INITIATE OLREORG command.

2
2

```
                                                   ,
                                              ┌─────◄─────┐         ┌─SET(RATE(100))─┐
►►──┬─INITIATE─┬──OLREORG──NAME──(──▼──partname─┴──)──┴─SET(RATE(rate))─┴──►
    ├─INIT─────┤
    ├─/INITIATE┤
    └─/INIT────┘
```

2
2

```
   ┌─OPTION(DEL)───┐
►──┼───────────────┼─────────────────────────────────────────►◄
   └─OPTION(NODEL)─┘
```

2
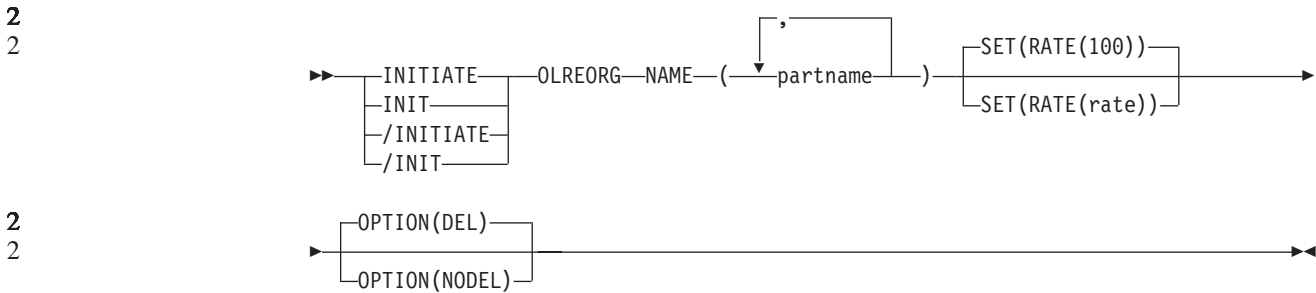
2 Environments, Keywords, and Parameters of the INITIATE OLREORG
2 Command

2   Table 4 below lists the environments, keywords, and parameters of the INITIATE
2   OLREORG command.

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

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

2   *Table 4. Valid Environments for the INITIATE OLREORG command Keywords, and*
2   *Parameters*

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

2 Usage of the INITIATE OLREORG Command

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

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

2   Part of the HALDB OLR processing involves:
2   • the creation of the output data sets at the beginning of the online reorganization
2   • the optional deletion of the inactive data sets when the online reorganization
2   completes

2
2
2
2
2

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 from a SPOC or IMSPLEX environment (not issued as a Type-1 command).

Related Reading: Refer to the Appendix of the *IMS Common Service Layer* for example XML for IMS commands.

NAME() Allows the specification of one or more HALDB partitions that the HALDB OLR is to run for.

Note: For /INITIATE OLREORG or /INITIATE OLREORG, only one NAME keyword can be specified 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 the specification of the RATE option as described below.

RATE()
Allows the specification of the RATE at which the HALDB OLR is to be run.

RATE A rate value of 1 to 100 can be specified. RATE(100) is the default that allows the HALDB OLR to run as fast as possible (depending on system resources, system contention, logging waits, etc.) without any coded waits. The value is the approximate percentage of elapsed time to be devoted to copying records. The remaining time (100 - value) is wait or idle time, that allows other IMS applications to run faster than they can run with RATE(100) specified.

For example: RATE(50) causes the HALDB OLR to run about twice as long as RATE(100). The RATE can be changed with the UPDATE OLREORG IMSplex command when adjustments are desired, depending on system contention (to slow down the HALDB OLR) or lack of system contention (to speed up the HALDB OLR).

OPTION() Allows the specification of the DEL/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 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 input 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 entered 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 or 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 yyyyyyyy (tracker, XRF alt)
DFS128  - TOO FEW KEYWORDS
DFS136  - COMMAND xxxxxxxxxxx INVALID FOR yyyyyyyy (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 "HALDB OLR DFS Series Messages" on page 67 and the *IMS 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 5 on page 35 shows the INITIATE 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 /INITIATE or to the /INIT command.

2 | **Keyword** | Identifies keyword on the command that caused the field to be generated.
N/A appears for output fields that are always returned.

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

2 *Table 5. Output Fields for INITIATE OLREORG Command*

Short Label	Keyword	Meaning
PART	N/A	Partition name.
MBR	N/A	The IMS from which the comand was issued.
CC	N/A	Completion code.

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

2 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 *IMS Common Service Layer* for a list of the OM codes and the code meanings.

2 Table 6 includes the return and reason codes and a brief explanation of the `INITIATE OLREORG` command. Table 7 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.

2 *Table 6. 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

2 *Table 7. 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 partition secondary index
C' 20'	HALDB OLR is already in progress

Examples of INITIATE OLREORG Input and Output

2 This section provides OM API and classic-format input and output examples and explanations for the `INITIATE OLREORG` command.

2 ***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 DB 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 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 parition PDHDOKA is issued from a terminal, resulting in the message, DFSO58I 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 Ouput:

```
Partition MbrName    CC
PDHDOKA   IMSA         0
PDHDOKA   IMS1        C3
PDHDOKB   IMSA         0
PDHDOKB   IMS1        C3
```

Explanation:

TBD

OM API Input:

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

2   **OM API Output:**
2

```
<imsout>
<ctl>
<omname>OM1OM    </omname>
<omvsn>1.2.0</omvsn>
<xmlvsn>1    </xmlvsn>
<statime>2003.168 21:19:06.827362</statime>
<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 slbl="PART" llbl="Partition" scope="LCL" sort="A" key="1"
 scroll="NO" len="9" dtype="CHAR" align="left" />
<hdr slbl="MBR" llbl="MbrName" scope="LCL" sort="N" key="0" scroll="NO"
 len="8" dtype="CHAR" align="left" />
<hdr slbl="CC" llbl="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>
```

Figure 3. OM API Output Example of the INITIATE OLREORG Command

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

2

2 INIT.DB

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

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

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

2 Format of INIT.DB

```
2    +------------------------------------------------------------------------------+
2    |                                                                              |
2    |                         .-RECOVABL---------------.            .-0-.          |
2    | >>--INIT.DB--DBD(name)--|------------------------|--SHARELVL(|-1-|)--------->|
2    |                         |            .-ICREQ---. |           |-2-|           |
```

```
2                                        '-NONRECOV--|---------|--'              '-3-'                    |
2                                                    '-NOICREQ-'
2                                                                                                         |
2          .-TYPEIMS---------------------------------------.
2        >--.--------------------------------------------------.------------------>
2          | -TYPEFP--------------------------------------------|
2          |            .-OLRCAP-----------.                    |
2          '-TYPHALDB--|-----------------|--.-HIKEY------------.----------->       |
2          |
2                              .-DBTRACK--.
2        >--.-------------------.--.----------.--->
2          '-GSGNAME(gsgname)- '  '-RCVTRACK-'
2          +---------------------------------------------------------------------+
```

INIT.DB Command Usage

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

2 **OLRCAP | OLRNOCAP**
2   Mutually exclusive, optional parameters used to indicate that HALDB OLR is
2   allowed to run.

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

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

NOTIFY.REORG

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

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

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

Format of NOTIFY.REORG

```
2     +---------------------------------------------------------------------+
2     |                                                                     |
2     |                                 .-CURRENT-----------.               |
2     |>>--NOTIFY.REORG--DBD(name)--DDN(name)--+--------------------+----------> |
2     |                                 '-RUNTIME(timestamp)-'       |
2     |
2     |>--.-------------------.--.---------------------.------------------> |
2     |   |         .-1-----. |  |         .-1-----. |                    |
2     |   '-FILESEQ(-'-value-'-)-'  '-FILESEQ2(-'-value-'-)-'              |
2     |
2     |>--.--------------.--.-------------.------------------------------> |
2     |   '-ICDSN(name)-'  '-ICDSN2(name)-'                               |
2     |
2   1 |   .-OFFLINE-------------------------------.                        |
2   1 |>--+-------------------------------------------+-----------------> |
2   1 |   '-ONLINE-.-----------------------------------.--'              |
2   1 |           |                    .-NORECOV-. |                      |
2   1 |           '-STOPTIME(timestamp)-+---------+--'                    |
2   1 |                                 '-RECOV---'                       |
2     |
2     |>--.--------------.--.----------------------.----------------------> |
```

```
2          '-RECDCT(value)-'  |        .-3400-----.  |
2                             '-UNIT(-'-unittype-'-)-'
2
2     >--.----------------------.--.------------.-------------------------->
2       |         .-3400-----.  |  '-USID(value)-'
2       '-UNIT2(-'-unittype-'-)-'
2
2     >--.----------------------.--.------------------------.---------------><
2       |          <-,------<   |  |           <-,------<   |
2       '-VOLLIST(--(volser)--)-'  '-VOLLIST2(--(volser)--)-'
2
2     +------------------------------------------------------------------+
```

2 # NOTIFY.REORG Command Usage

2 **ONLINE | <u>OFFLINE</u>**
2 Mutually exclusive, optional parameters used to specify whether the
2 reorganization is a HALDB OLR or an offline reorganization.

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

2 **<u>OFFLINE</u>** specifies that the DBDS was reorganized offline (default)

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

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

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

2 **STOPTIME(TIMESTAMP)**
2 Optional parameter used to specify the time that a HALDB OLR successfully
2 completed. This parameter may only be specified with the ONLINE keyword.

Chapter 5. Monitoring HALDB OLR

Use the QUERY OLREORG command to see status information about active HALDB OLRs in an IMSplex environment.

Use the /DISPLAY DB command to see the status of the specified databases or partitions locally (including those with HALDB OLRs in progress).

Use the /DISPLAY DB OLR command to see all HALDB partitions that have HALDB OLRs in progress locally.

Use the QUERY DB STATUS(OLR) command to see a list of all databases with HALDB OLR in progress.

QUERY OLREORG Command

The QUERY OLREORG command returns status information about active HALDB OLRs in an IMSplex 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 46..

Format of the QUERY OLREORG Command

The syntax diagram below illustrates the format of the QUERY OLREORG command.

```
                        ┌─NAME─(──*──)──────────┐
►►──┬─QUERY─┬──OLREORG──┤                       ├──►
    └─QRY───┘           │        ┌─,────────┐   │
                        └─NAME─(─▼─partname─┴─)─┘

►──┬────────────────────────────────┬──┬──────────────────────────┬──►◄
   │          ┌─,─────────┐          │  │        ┌─,─────────┐     │
   └─STATUS(──▼─┬─RUNNING──┬┴──)─────┘  └─SHOW(──▼─┬─STATUS─┬┴──)──┘
                ├─OWNED────┤                       ├─ALL────┤
                ├─NOTOWNED─┤                       ├─RATE───┤
                └─UNKNOWN──┘                       └─BYTES──┘
```
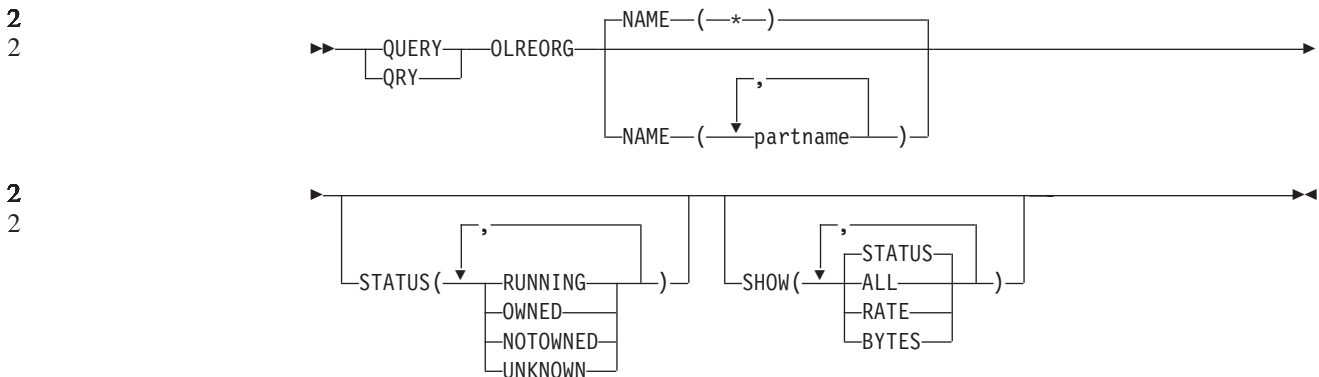
Environments, Keywords, and Parameters of the QUERY OLREORG Command

Table 8 on page 42 below 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 8. 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 specific
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.

Attention: A non-zero return code and non-zero reason code will be returned
when the command is routed to an XRF (Extended Restart Facility) alternate
system and when the command is routed from 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 this command is defined in XML and is available to automation
programs which communicate with OM.

Related Reading: Refer to the Appendix of the *IMS: Common Service Layer Guide
and Reference* for example IMS command XML.

NAME() Displays the Online Reorganization status for the specified HALDB
partitions.

> **PARTNAME** 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.

SHOW() Specifies the output fields to return. If SHOW is not specified only
the partname is returned (depending on what is specified in the
STATUS keyword).

> **ALL** Returns STATUS, RATE, and BYTES on each
> response line.
>
> **BYTES** Returns the number of bytes that have been moved
> to the output data set.
>
> **RATE** Returns the rate from 1 to 100, at which the HALDB
> OLR is running.
>
> **STATUS** Online reorganization status.

STATUS() Displays online reorganizations that posess 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.

	NOTOWNED	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.
	OWNED	Specifies that the output is for HALDB OLRs that are owned by any IMS. Those owned by this IMS display a `STATUS` of `RUNNING`.
	RUNNING	Specifies the output will be for the HALDB OLRs that are owned by each IMS for the specified partname or partnames. This option can be used to determine which IMS has an online reorganization running for which partname.
	UNKNOWN	Specifies that the output is for those partnames 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 by the IMS not being authorized to the named partname because of an IRLM failure.

QUERY OLREORG Output Fields

Table 9 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.

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 9. Output Fields of QUERY OLREORG

Short Label	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.
RATE	RATE	The speed at which HALDB OLR runs. A value of 1 to 100 percent.
BYTES	BYTES MOVED	Number of bytes moved

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 for a list of the OM codes and the code meanings.

Table 6 on page 35 includes the return and reason codes and a brief explanation of the `QUERY OLREORG` command. Table 11 on page 44 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 10. 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 11. 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 partition secondary index
C' 24'	No HALDB OLR is in progress
C' CB'	Partition not in specified status

Examples of QUERY OLREORG Input and Output

This section provides OM API and Classic-format input and output examples and explanations for the QUERY OLREORG command and the /DISPLAY DB OLR command.

Entry ET:

```
/DIS DB OLR
```

Comments can be entered enclosed within slashes and asterisks (/* */) after the command.

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:

- partition name
- master database name
- OLR rate
- number of bytes moved
- status of the OLR

When this command is entered as a Type 1 command, the response is returned in the message format. One or more of the following messages might be returned:

```
2    DFS0725I INITIATE|UPDATE|TERMINATE OLREORG COMMAND FOR DB dbnamexx COMPLETE.
2    CC=  nn
2      where: dbnamexx is the HALDB partition name entered on the command
2             nn      is the completion code
2
2    DFS058  - INITIATE COMMAND IN PROGRESS
2    DFS107  - REQUIRED KEYWORD NOT PRESENT
2    DFS110  - COMMAND KEYWORD OLREORG INVALID FOR yyyyyyyy (tracker, XRF alt)
2    DFS128  - TOO FEW KEYWORDS
2    DFS136  - COMMAND xxxxxxxxxxx INVALID FOR yyyyyyyy (tracker, XRF alt)
2    DFS165  - COMMAND VERB IS INVALID
2    DFS2026 - NO WKAP STORAGE
2    DFS2038 - INSUFFICIENT STORAGE AVAILABLE TO PROCESS COMMAND
2    DFS2262 - TOO MANY DATABASE PARAMETERS
2    DFS3630 - GENERIC PARAMETER NOT ALLOWED
```

2 **TSO SPOC Input:**

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

2 **TSO SPOC Output:**

```
2    Partition MbrName    CC LCLSTAT  RATE  BYTES-MOVED
2    PDHDOKA   IMSA         0 RUNNING     5       115260
2    PDHDOKB   IMSA         0 RUNNING     5          634
```

2 **OM API Input:**

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

2 **OM API Output:**
2

```
<imsout>
<ctl>
<omname>OM1OM    </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 slbl="PART" llbl="Partition" scope="LCL" sort="A" key="1"
 scroll="NO" len="9" dtype="CHAR" align="left" />
<hdr slbl="MBR" llbl="MbrName" scope="LCL" sort="N" key="0" scroll="NO"
 len="8" dtype="CHAR" align="left" />
<hdr slbl="CC" llbl="CC" scope="LCL" sort="N" key="0" scroll="YES"
 len="4" dtype="INT" align="right" />
<hdr slbl="LSTT" llbl="LCLSTAT" scope="LCL" sort="NO" key="0"
 scroll="YES" len="*" dtype="CHAR" align="left" />
<hdr slbl="RATE" llbl="RATE" scope="LCL" sort="N" key="0" scroll="YES"
 len="4" dtype="INT" align="right" />
<hdr slbl="BYTES" llbl="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>
```

Figure 4. OM API Output Example of the QUERY OLREORG Command

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 command.

```
DATABASE  TYPE  TOTAL UNUSED  TOTAL UNUSED ACC  CONDITIONS
DBVHDJ05  PHDAM                             UP
PVHDJ5C   PART                             UP   OLR
DBOHIDK5  PHIDAM                           UP
POIHDKD   PART                             UP   OLR
```

Figure 5. /DISPLAY DB 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 following information is also returned:

- master db name
- partition name
- OLR rate
- number of bytes moved
- status of OLR

The /DISPLAY DB OLR information is returned with the following 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.

Example of /DISLAY DB OLR Output:

Here is an example of output returned to the originating terminal from a /DISPLAY DB PVHDJ5C POIHDKD command.

```
DATABASE     PART  RATE        BYTES  STATUS
DBOHIDK5  POHIDKD     1  123456789999 RUNNING
DBVHDJ05  PVHDJ5C    50             0 RUNNING
```

Figure 6. /DISPLAY DB OLR output

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 Command Reference*.

2

2 # Chapter 6. Modifying and Tuning HALDB OLRs

2 Use the UPDATE OLREORG command to change the RATE and DEL or NODEL
2 option, or the RATE or DEL or NODEL option of a HALDB OLR that is owned. For
2 more information, see "UPDATE OLREORG Command" on page 55.

2 The TERMINATE OLREORG command is used to stop one or more HALDB OLRs
2 in an IMSplex environment. For more information, see "TERMINATE OLREORG
2 COMMAND."

2 Use the /START DB command to perform the following tasks:
2 • allocate a HALDB partition the first time it is referenced
2 • allocate a HALDB partition if a /START DB command is issued for a partition that
2 has EEQEs or for a partition that is authorized

2 For more information, see "/START DB Command HALDB OLR Support" on page
2 61.

2 There are some effects to a HALDB master database that has been taken offline
2 with a /DBRecovery DB command, regarding QUERY OLREORG. For more
2 information, see "/DBRecovery Command HALDB OLR Support" on page 62.

2 Issue the /RECover command to recover the data set to prepare that data set that
2 was created by HALDB OLR and continue the online reorganization. For more
2 information, see "/RECover Command HALDB OLR Support" on page 62.

2 Issue the CHANGE.DB command to modify the status of a HALDB OLR in the
2 RECON data set. For more information, see "CHANGE.DB" on page 62.

2

2 ## TERMINATE OLREORG COMMAND

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

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

2 Figure 7 on page 50 shows the processing steps of a HALDB OLR where you
2 invoke TERMINATE OLREORG to temporarily stop a HALDB OLR during copying
2 and then invoke INITIATE OLREORG to resume the online reorganization. This
2 figure contains columns depicting your (User) and IMS control flow and the statuses
2 of the sets of data sets as processing occurs.
2
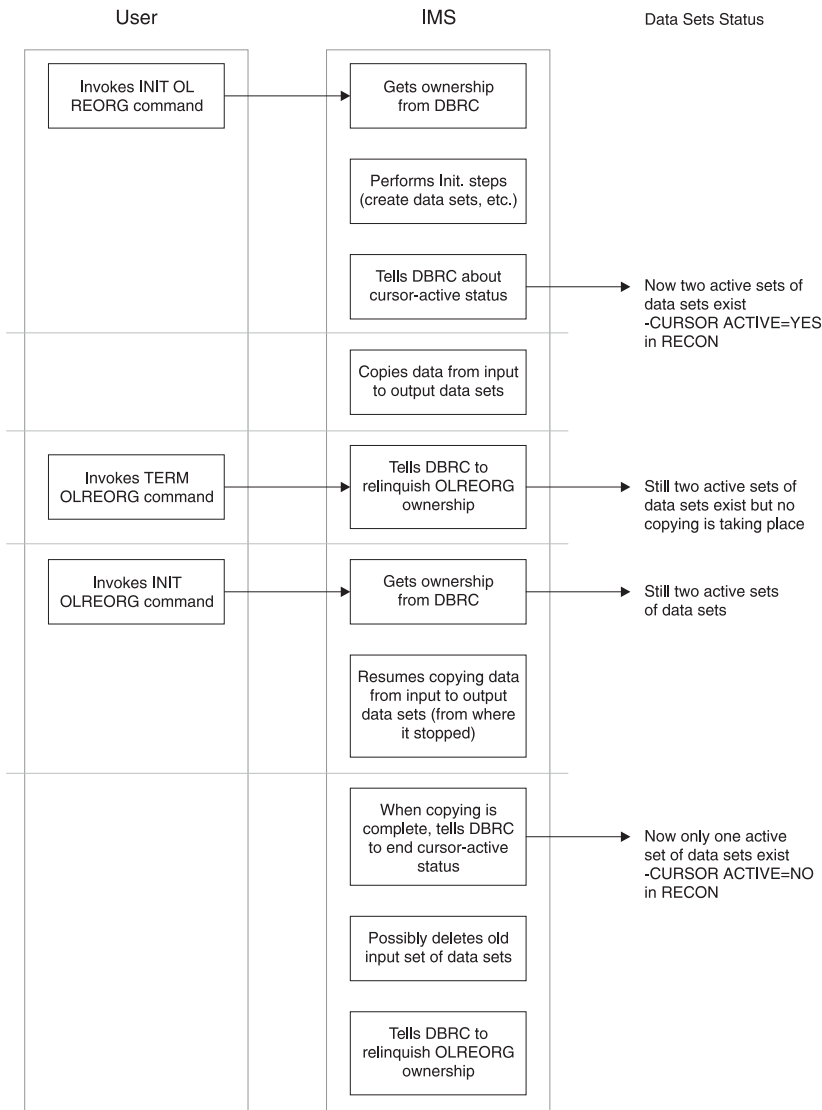
Figure 7. HALDB OLR Process with a TERMINATE OLREORG

2

2 /TERMINATE OLREORG Command

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

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

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

2 **Related Reading:**

2
2

- See Chapter 7, "Messages and Codes of HALDB OLR," on page 67 for more information about message DFS0725I.

Format of the TERMINATE OLREORG Command

2
2
The syntax diagram below illustrates the format of the `TERMINATE OLREORG` command.
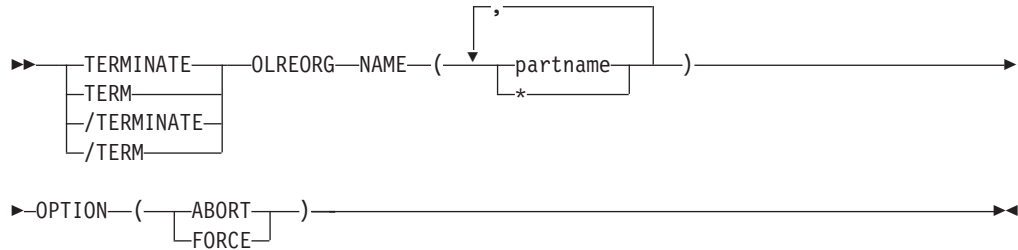
2
2

```
>>─┬─TERMINATE──┬──OLREORG──NAME──(─┬─┬─partname─┬─┬──)──────────────────>
   ├─TERM───────┤                   │ └─*────────┘ │
   ├─/TERMINATE─┤                   └──────,───────┘
   └─/TERM──────┘
```

2
2

```
>──OPTION──(──┬─ABORT─┬──)──────────────────────────────────────────────><
              └─FORCE─┘
```

2

Environments, Keywords, and Parameters of the TERMINATE OLREORG Command

2
2
Table 12 below lists the environments, keywords, and parameters of the `TERMINATE OLREORG` command.

2
2
2
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.

2
2
Restriction: This command does not support keyword synonyms. Only one version of each keyword can be specified for each command issued.

2
2
Table 12. 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

2
2
2
2
2
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 cursor-active status until the online reorganization is resumed with an INITIATE OLREORG command and completes, or an offline reorganization is run.

2
2
2
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 a SPOC, IMSPLEX environment and not as a classic-format command.

2
2
Related Reading: Refer to the Appendix of the *IMS Version 9: Common Service Layer Guide and Reference* for example XML for IMS commands.

NAME() Allows you to specify one or more HALDB partitions for which HALDB OLR is to be stopped.

Note: Only one NAME keyword is allowed for /TERMINATE OLREORG and /TERM OLREORG.

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 below.

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

??Question for Developers:
Does anything else belong under this heading?

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 or 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 67.

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 13 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.

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 13. Output Fields for TERMINATE OLREORG Command

Short Label	Keyword	Meaning
PART	N/A	Partition name
MBR	N/A	The IMS from which the command was issued
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 IMS Common Service Layer* for a list of the OM codes and the code meanings.

Table 6 on page 35 includes the return and reason codes and a brief explanation of the `TERMINATE OLREORG` command. Table 15 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 14. 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 15. 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

2

2 *Table 15. Completion Codes for the TERMINATE OLREORG Command (continued)*

Completion Code	Meaning
C' 14'	Resource not a partition name
C' 1C'	Resource is a partition secondary index
C' 24'	No HALDB OLR is in progress
C' 28'	NO DMB loaded - need to schedule

2
2
2
2
2
2

2 # Examples of TERMINATE OLREORG Input and Output

2
2 This section provides OM API and classic-format input and output examples and
explanations for the `TERMINATE OLREORG` command.

2 ***EXAMPLE 1:***

2 **Entry ET:**
2 ```
/TERM OLREORG NAME(PDHDOKA)
```

2 **Response ET:**
2 ```
DFS0725I TERMINATE OLREORG COMMAND FOR DB PDHDOKA  COMPLETE. CC=  24
```

2 **Explanation:**

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

2 **Related Reading:**
2 • See Chapter 7, "Messages and Codes of HALDB OLR," on page 67 for more
2 information about message DFS0725I.

2 ***EXAMPLE 2:***

2 **TSO SPOC Input:**
2 ```
TERM OLREORG NAME(PDHDOKA,PDHDOKC)
```

2 **TSO SPOC Output:**
2 ```
Partition MbrName   CC
PDHDOKA    IMSA      0
PDHDOKA    IMS1     24
PDHDOKC    IMSA     24
PDHDOKC    IMS1     24
```

2 **Explanation:**

2 TBD

2 **OM API Input:**
2 ```
CMD (TERM OLREORG NAME(PDHDOKA,PDHDOKC))
```

2 **OM API Output:**
2

```
<imsout>
<ctl>
<omname>OM1OM   </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 slbl="PART" llbl="Partition" scope="LCL" sort="A" key="1"
 scroll="NO" len="9" dtype="CHAR" align="left" />
<hdr slbl="MBR" llbl="MbrName" scope="LCL" sort="N" key="0" scroll="NO"
 len="8" dtype="CHAR" align="left" />
<hdr slbl="CC" llbl="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>
```

Figure 8. OM API Output Example of the TERMINATE OLREORG Command

Explanation: The TERM OLREORG command is issued to stop the OLR for databases 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 or for either of the two partitions at IMS1 because OLR is not in progress on IMS1.

UPDATE OLREORG Command

The `UPDATE OLREORG` command allows you to change the RATE and DEL or NODEL option, or the RATE or DEL or NODEL option of a HALDB OLR that is owned.

2
2
2
2
2
2

Figure 9 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 (User) and IMS and the statuses of the sets of data sets as processing events occur.

User   IMS   Data Sets Status

```
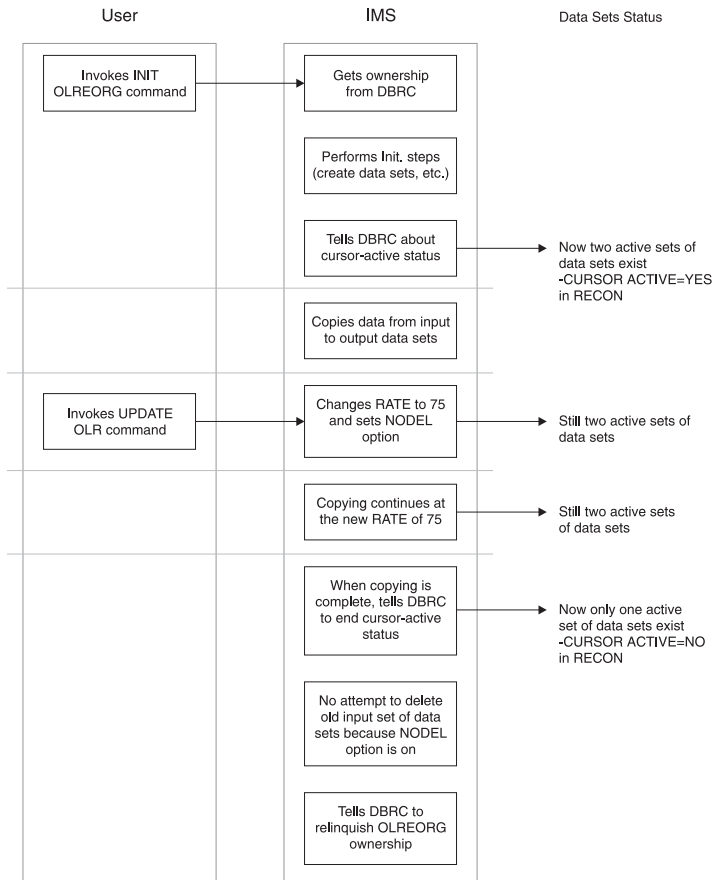┌─────────────────┐      ┌──────────────────┐
│ Invokes INIT    │ ───▶ │ Gets ownership   │
│ OLREORG command │      │ from DBRC        │
└─────────────────┘      └──────────────────┘

                         ┌──────────────────┐
                         │ Performs Init.   │
                         │ steps (create    │
                         │ data sets, etc.) │
                         └──────────────────┘

                         ┌──────────────────┐      Now two active sets of
                         │ Tells DBRC about │ ───▶ data sets exist
                         │ cursor-active    │      -CURSOR ACTIVE=YES
                         │ status           │      in RECON
                         └──────────────────┘

                         ┌──────────────────┐
                         │ Copies data from │
                         │ input to output  │
                         │ data sets        │
                         └──────────────────┘

┌─────────────────┐      ┌──────────────────┐      Still two active sets of
│ Invokes UPDATE  │ ───▶ │ Changes RATE to  │ ───▶ data sets
│ OLR command     │      │ 75 and sets      │
└─────────────────┘      │ NODEL option     │
                         └──────────────────┘

                         ┌──────────────────┐      Still two active sets
                         │ Copying continues│ ───▶ of data sets
                         │ at the new RATE  │
                         │ of 75            │
                         └──────────────────┘

                         ┌──────────────────┐      Now only one active
                         │ When copying is  │ ───▶ set of data sets exist
                         │ complete, tells  │      -CURSOR ACTIVE=NO
                         │ DBRC to end      │      in RECON
                         │ cursor-active    │
                         │ status           │
                         └──────────────────┘

                         ┌──────────────────┐
                         │ No attempt to    │
                         │ delete old input │
                         │ set of data sets │
                         │ because NODEL    │
                         │ option is on     │
                         └──────────────────┘

                         ┌──────────────────┐
                         │ Tells DBRC to    │
                         │ relinquish       │
                         │ OLREORG ownership│
                         └──────────────────┘
```

Figure 9. HALDB OLR Process with an UPDATE OLREORG

2 /UPDATE OLREORG Command

2
2
2

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

2
2
2

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

2
2

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

2

Related Reading:

2
2

- See Chapter 7, "Messages and Codes of HALDB OLR," on page 67 for more information about message DFS0725I.

Format of the UPDATE OLREORG Command

The syntax diagram below illustrates the format of the UPDATE OLREORG (update HALDB OLR) command.

```
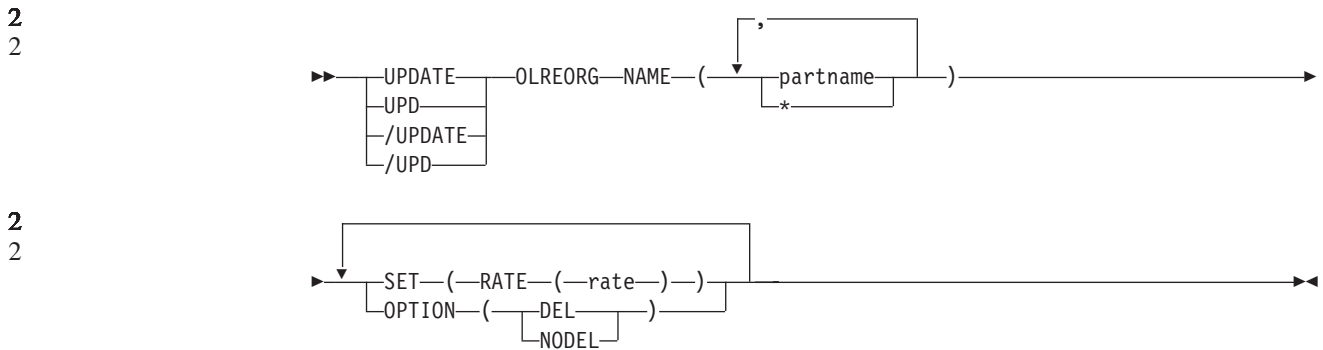>>──UPDATE─────OLREORG──NAME──(──┬─▼─partname─┬──)───────────────────>
     ├─UPD───┤                   └─*─────────┘
     ├─/UPDATE┤
     └─/UPD──┘
```

```
>──┬─▼─SET────(──RATE──(──rate──)──)─┬──────────────────────────><
   └─OPTION──(──┬─DEL──┬──)─┘
                └─NODEL─┘
```

Environments, Keywords, and Parameters of the UPDATE OLREORG Command

Table 16 below lists the environments, keywords, and parameters of the UPDATE 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 16. Valid Environments for the UPDATE 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 UPDATE OLREORG Command

The UPDATE OLREORG command allows you to change the rate and/or the DEL or NODEL data set disposition flags of an owned HALDB OLR. See the section, "Reorganization Processing Time" on page 12 for performance considerations related to changing the RATE at which the online reorganization is run.

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

Related Reading: Refer to the Appendix of the *IMS: Common Service Layer Guide and Reference* for example IMS command XML.

NAME()   Allows you to specify one or more HALDB partitions for the UPDATE OLREORG IMSplex command.

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.

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 A rate value of 1 to 100 can be specified. RATE(100) allows the online reorganization to run as fast as possible (depending on system resources, system contention, logging waits, etc.) without any coded waits. A RATE other than 100 causes coded waits and causes the online reorganization to run at a reduced speed proportional to 100%. For example: RATE(50) causes the online reorganization to run twice as long as RATE(100).

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.

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

> **??Question for Developers:**
> Does anything else belong under this heading?

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 or 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
```

2
2

For more complete information about this message, see "HALDB OLR DFS Series Messages" on page 67.

2
2
2

Command Responses for UPDATE OLREORG from OM API

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

2

UPDATE OLREORG Output Fields:

2
2

Table 17 shows the UPDATE OLREORG output fields. The columns in the table are as follows:

2

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

2
2

> **Note:** This field does not apply to the /UPDATE or to the /UPD command.

2
2

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

2

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

2

Meaning   Provides a brief description of the output field.

2

Table 17. Output Fields for UPDATE OLREORG Command

Short Label	Keyword	Meaning
PART	N/A	Partition name.
MBR	N/A	The IMS from which the command was issued.
CC	N/A	Completion code.

2
2

2
2

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

2
2
2

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 for a list of the OM codes and the code meanings.

2
2
2
2
2

Table 6 on page 35 includes the return and reason codes and a brief explanation of the UPDATE OLREORG command. Table 19 on page 60 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 against a HALDB partition.

2

Table 18. 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

2   *Table 19. 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 partition secondary index
C' 24'	No HALDB OLR is in progress

Examples of UPDATE OLREORG Input and Output

This section provides OM API and Classic-format 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 DFS0725 command response.

Related Reading:
*   See Chapter 7, "Messages and Codes of HALDB OLR," on page 67 for more information about message DFS0725I.

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
```

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 there is OLR is not in progress. A completion code of 24 is returned in the IMS1 response.

OM API Input:

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

2   **OM API Output:**

2

```
<imsout>
<ctl>
<omname>OM1OM    </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 slbl="PART" llbl="Partition" scope="LCL" sort="A" key="1"
 scroll="NO" len="9" dtype="CHAR" align="left" />
<hdr slbl="MBR" llbl="MbrName" scope="LCL" sort="N" key="0" scroll="NO"
 len="8" dtype="CHAR" align="left" />
<hdr slbl="CC" llbl="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>
```

Figure 10. OM API Output Example for the UPDATE OLREORG Command

2   **Explanation:**

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

2

2 /START DB Command HALDB OLR Support

2   The following list describes the results of issuing a /START DB command for a
2   HALDB partition:

2
2

- The partition will be allocated by issuing the /START DB command if the OPEN keyword is used and the ID of the partition is authorized.

2
2

- The partition is allocated the first time it is referenced.

2
2
2

- The partition is also allocated if a /START DB command is issued for a partition that has EEQEs or for a partition that is authorized.

2
2
2
2

Except for migrated data sets, if allocation fails during partition selection or during a /START DB command with the OPEN keyword, the partition that remains unallocated is also authorized.

2 /DBRecovery Command HALDB OLR Support

2
2
2

If a HALDB master database has been taken offline with a /DBRecovery DB command, QUERY OLREORG is not able to determine its status and returns an UNKNOWN status for that offline HALDB master database.

2
2

If an output data set needs to be recovered before a HALDB OLR is completed then the /DBRecovery command is part of that recovery process.

2
2

Related Reading: See the step-by-step process recovery in the section entitled, "Recovery - DFSURDB0" on page 21..

2 /RECover Command HALDB OLR Support

2
2
2

If an error is detected while recovering a data set that was created by HALDB OLR, issue the /RECover command to recover the data set to prepare that data set and continue the online reorganization.

2 CHANGE.DB

2
2

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

2 Format of CHANGE.DB

```
+---------------------------------------------------------------------------+
|                                                                           |
|                      .-ALL-------.                                         |
|      >>--CHANGE.DB--+-----------+--.-------.-----------------------------> |
|                     '-DBD(name)-'  |-AUTH---|                              |
|                                    '-NOAUTH-'                              |
|                                                                           |
|      >--.-------------------------.--.--------.--.---------.-------------> |
|         |-NOBACK------------------|  |-OLRCAP-- |  |-OLRRGON--|            |
|         '-BACKOUT(value)SSID(name)-'  '-OLRNOCAP-'  '-OLRRGOFF-'           |
|                                                                           |
|      >--.---------------------------.--.----------------.--.---------.---> |
|         '-OLRIMSID(name)NOOWNER-'  '-OLRDBDS(-.-A-.-)-'  |-PINIT---|       |
|                                              '-M-'       '-NOPINIT-'       |
|                                                                           |
|      >--.----------.--.----------.--.------------------.--.---------.----> |
|         |-READON--|   |-NONRECOV-|  '-SHARELVL(-.-0-.-)-'  |-TYPEFP--|     |
|         '-READOFF-'   '-RECOVABL-'             |-1-|       '-TYPEIMS-'     |
|                                                |-2-|                       |
|                                                '-3-'                       |
|                                                                           |
|      >--.------------------.--.----------.-------------------------------><|
|         |-GSGNAME(gsgname)-|  |-RCVTRACK-|                                 |
```

```
2                         |   '-NOTCOVER---------'  '-DBTRACK--'                              |
2                         |                                                                   |
2                         +-------------------------------------------------------------------+
```

CHANGE.DB Command Usage

The following keywords are specified when changing the status of a HALDB OLR in the RECON data set. They can only be specified for a HALDB partition (TYPE=PART). HALDB Online Reorganization must have already been attempted for this partition, so the M-through-V data sets must exist.

Use of the CHANGE.DB command for these parameters is not necessary for normal operations. The ALL parameter is not allowed when changing the status of an online reorganization, except for CHANGE.DB ALL OLRCAP. All HALDBs that are available to be marked as OLRCAP are marked as OLRCAP. Those unavailable for change (such as HALDBs that are currently authorized) are skipped and are not marked as OLRCAP.

CHANGE.DB Command HALDB OLR Parameters

OLRRGOFF | OLRRGON

Mutually exclusive, optional parameters used to specify whether HALDB OLR is in progress.

OLRRGON specifies that HALDB OLR is in progress. Both the A-through-J and the M-through-V data sets contain active data.

OLRRGOFF specifies that HALDB OLR has completed and that either the A-through-J set of data sets or the M-through-V set of data sets contain active data. The flag will be changed to indicate that the inactive data sets are now the active data sets. `ONLINE REORG ACTIVE = NO` will be set in the DB record. The OLRIMSID will be cleared.

Note: Manual intervention may be needed to either delete the online REORG records or to add a stoptime. See the *IMS DBRC Guide and Reference* for more information on this and the DELETE.REORG and NOTIFY.REORG commands.

OLRIMSID(NAME) | NOOWNER

Mutually exclusive, optional parameters you use to specify or change the name of the IMS responsible for the HALDB OLR.

OLRDBDS(A | M)

Optional parameter you use to indicate the active set of DBDSs for this partition either A-through-J and X or M-through-V and Y.

A specifies that the A-through-J and X data sets are active

M specifies that the M-through-V and Y data sets are active

Valid for HALDB master(type=HALDB) only:

OLRCAP | OLRNOCAP

Mutually exclusive, optional parameters that specify if the HALDB can be 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

2   command fails if HALDB OLR is active or if any M-V or Y
2   DBDSs are active. In this case the M-V or Y DBDSs and all of
2   their associated records are deleted from RECON for all
2   partitions.

Part 3. Diagnosing HALDB OLR

Chapter 7. Messages and Codes of HALDB OLR

These are the IMS messages that pertain to HALDB OLR.

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

In This Chapter:
- "HALDB OLR DFS Series Messages"
- "HALDB OLR DSP Series Messages" on page 83
- "HALDB OLR Abends and Codes" on page 85

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/DBR`, a `global/STO`, 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 `/DBD` 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 backouts for the named database.
08	For sharing within a host, the requesting subsystem is not within the authorized host.
09	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.
10	The database has been previously authorized to the subsystem.
11	An invalid parameter is found during the database usage compatibility evaluation process. Possibly, the database record in the RECON data set is bad.
12	The current authorization state in DBRC is invalid because of an unauthorization error.
13	An error occurred in DBRC while trying to process an authorization request.
14	An error occurred in DBRC while trying to process a change authorization request.
15	The database is already authorized to an active IMS subsystem.
16	The area is not marked as needing recovery for the DEDB initialization utility.
19	The data set name of the database

data set passed to DBRC does not match the data set name in the RECON data set.

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.

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).

22 The database is currently authorized to another subsystem.

23 Recovery-level tracking database.

24 There is gap in the log stream received from the RSR active site that can contain update records for the database or area.

25 All log records have been applied to the tracking database for the authorizing stream and authorization is not for an OFR stream.

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.

2 **31** The definition of the HALDB has
2 changed. The caller must get the new
2 definition before authorization can be
2 granted.

32 The HALDB needs to be initialized.

2 **33** Attempting to authorize the HALDB
2 master. Authorization can only be
2 requested at the partition level.

34 Partition Initialization is requesting authorization for a partition that does not need to be initialized.

2 **35** The DB partition cannot be authorized
2 until a high key is defined. A key is

2 required because the HALDB master
2 does not use a Partition Selection
2 Routine.

36 Image copy not allowed during online reorganization processing.

37 Loading into an M-V DBDS of a partition database is not allowed.

38 Offline reorg is not allowed when online reorg is active and the online reorg is owned by an IMS subsystem.

System Action: 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.

Programmer Response: Reason codes 01 through 09 and 15 through 19, indicating a temporary authorization failure, can be corrected by one of the following actions:

• Wait for a currently authorized subsystem to terminate or enter a /DBR 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)

• Verify that the correct log is input to the batch backout utility if DBRC=C is being used. (RC 01)

• Enter a /START DB, a /DBD, or a CHANGE.DB (auth) command. (RC 03)

• Enter a /START DB command. (RC 04)

• Run the appropriate IMS utility. (RC 05, 06 or 07)

• Run this IMS subsystem on the host where the database is currently authorized. (RC 08)

• Wait for the failed subsystems to complete the emergency restart or database backout. (RC 09)

• Use the DBRC command CHANGE.SUBSYS to indicate that an IMS subsystem is no longer active. (RC 15)

• Use the DBRC CHANGE.DBDS command to mark an area as needing recovery. (RC 16)

• Use the DBRC INIT.DBDS command to register a DBDS OR DEDB area to DBRC. (RC 18, 19)

• Check the RECON data set to determine the correct authorization state.

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

DFS135 **ALL PARAMETERS ARE INVALID**

Explanation: A /DIS DB OLR command was issued on the RSR tracker or the XRF alternate.

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 /DBR, /DBD, /STA, or /STO 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 /STO command. The area is only for the DBR, STA, and STO 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 indicated 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, which indicates that the command was successful.

Code (Dec)	Meaning
01	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.
02	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.
03	The database is marked as prohibiting further authorization for one of the following reasons; global a /DBR, a global /STO, or a CHANGE.DB(NOAUTH) command.
04	An error occurred at the close of the database. For detailed information about the reason codes, refer to message DFS730I.

Module DFSMDA00 can also issue return code 4 if dynamic allocation/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.

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.

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.

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.

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.

08 Dynamic allocation failed (from DFSMDA00). This indicates that the IMS dynamic allocation module (DFSMDA00) received a nonzero return code from MVS dynamic allocation.

For a /VUNLOAD command, an attempt to get the /VUNLOAD lock for an area in a coupling facility structure failed. The /VUNLOAD command is terminated and the structure is not unloaded.

11 An invalid parameter was found during the evaluation process for database-usage compatibility. The database record in the RECON data set might be invalid.

12 Dynamic allocation/deallocation failed because no storage was available. (from DFSMDA00)

Return code 12 can also mean the current authorization state in DBRC is invalid because of an unauthorization error.

14 The named database cannot be started because of one of the following possible conditions:

- 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 /DBD command issued. As a consequence, the original authorization level cannot be restored for the current subsystem.
- The changed database authorization level has been

requested through the /STA DB ACCESS command. The request for the change in ACCESS parameter failed because of the incompatibility of the database with the authorization level currently held in the other IMS subsystem.

Return code 14 can also mean an error occurred in DBRC while trying to process a change authorization request.

15 The duplicate AREA entry on the /DBR AREA command is ignored. The area was already unauthorized by DBRC.

16 Dynamic allocation or deallocation failed because of a system error. For detailed information about the reason codes, refer to message DFS2498I.

20 For the /UNLOCK SYSTEM command, one of the following conditions occurred:

- There were no EEQEs to purge.
- The /UNLOCK command was entered on an IMS system that is not XRF-capable.
- 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.
- 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 the failing active system has terminated, the processor containing the failing active system has been reset, or I/O prevention on the active system is complete.

21 One of the following has occurred:

- A VLOCK request cannot be granted when processing a /STA AREA or a /VUN AREA command. The area is marked as STOPPED. One sharing IMS system is holding the retained VLOCK exclusive.
- While processing either a /VUN AREA or /STA 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. Issue a `/STA AREA` command to correct the inconsistency.

- During /VUNLOAD or /STA 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 /STA AREA command to correct the inconsistency.

- During /STA AREA command processing, an IXLFORCE command is issued to delete any failed persistent connections from a prior failed system. The IXLFORCE was not able to delete the failed persistent connection. In this case, the following messages are generated:

  – DFS2783A UNABLE TO DELETE FAILED-PERSISTENT CONNECTION TO STR: SSSSSSSSSSSSSSSS REASON=RRRR

  – DFS0488I STA COMMAND COMPLETED. AREA= XXXXXXXX RC=21

24 A command with the "ALL" keyword will not be processed for a Fast Path MSDB. If /STA DB ALL access is entered, module DFSDBDR0 will not include MSDBs.

28 The named database cannot be started because of the following conditions:

- The changed database authorization level to a data entry database (DEDB) has been requested through the `/STA 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.

- Changing the DEDB access intent requires stopping the scheduling of all PSBs that have access to the DEDB in question. 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.

31 A `/START AREA` command failed because the area is marked as

"recovery-needed" in the DBRC RECON data set, or the DBRC RECON indicates that there are no available area data sets for the area.

34 Because a `/STA 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 zero value.

35 A backout error occurred for the database. A `/START` command for the database is invalid until a `/DBRECOVERY` command has been issued and batch backout has been performed.

37 One of the following has occured:

- 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 `DFS2444I`, which precedes `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 backouts exist for the database. IMS must have update or exclusive access to the database to redo the incomplete backouts. The command is not processed for the specified database.

41 The area is not loaded into the Coupling Facility. This will not prevent access to the area unless 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 `/STO ADS` command is entered, the ADS of the named ddname is already stopped. This `/STO 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 /STA, /STO, or /DBR command. ('MSDB' is allowed on /DBD.) The keyword is ignored and the remainder of the command is processed.

51 A /DBR command was entered on the RSR tracker for a Fast Path DEDB. The DEDB name is ignored and the remainder of the command is processed.

52 A /DBR or /STA 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 /DBR command must be entered again after the database is recalled. The /STA 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 /DBR or /STO command to deallocate the data sets.

55 The command failed because the FF DB or FP area is currently being recovered.

56 An error occured 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, /STOP, /LOCK, or /UNLOCK 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.

59 The named HALDB partition cannot be processed in response to a /DBDUMP, /DBRECOVERY, /START, /STOP, /LOCK, or /UNLOCK command. Restart processing has not completed opening the named HALDB partition. 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. Reissue the command after Restart processing is complete.

62 The named database failed directory initialization or the DBB or PSB pool space was not large enough to accomodate 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 accomodate 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: DFSDBDR0, DBFARDA0, DFSRDSH0, DFSDBDR0, DBFARDA0, 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 /INIT, /UPD, or /TERM 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 the Online Reorganization has started. A nonzero completion code indicates that an unexpected error was encountered during command processing.

Completion Code Meaning   **Meaning**

0   REQUEST SUCCESSFUL

10   RESOURCE NAME INVALID

14   RESOURCE NOT A PARTITION NAME

1C   RESOURCE IS A PARTITION SECONDARY INDEX

20   INIT OLREORG ISSUED, OLR IS ALREADY IN PROGRESS

24   NO OLREORG IS IN PROGRESS

28   NO DMB LOADED - need to schedule a PSB that references the DB

C1   OLREORG DDIR MISSING OR DFSPNT ZERO

C2   OLREORG ALREADY OWNED

C3   OLREORG PARTITION STOPPED, or LOCKED, or DBDUMPED

C4   OLREORG PARTITION BEING RECOVERED

C5   OLREORG PARTITION ACCESS NOT UPDATE/EXCLUSIVE

C6   OLREORG SCHEDULING FAILED - Check HALDB Master status

C7   OLREORG PARTITION NOT AUTHORIZED - Check HALDB Master status

C8   OLREORG MASTER STOPPED, or LOCKED, or DBDUMPED

C9   OLREORG MASTER ACCESS NOT UPDATE or EXCLUSIVE

CA   OLREORG SET OWNERSHIP FAILED

CB   PARTITION NOT IN SPECIFIED STATUS

System Action: IMS continues processing.

Problem Determination: 6, 11

Module: DFSORC10, DFSORP50

DFS0832I   ABEND Uwwww REASON CODE xxxx yyyyyyyyy zzzzzzzz

Explanation: This is an informational message to assist in debugging 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, DFSDDLE0, 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:

wwww   is the pseudo abend code

xxxx   is the reason code for HALDB abends, or N/A for non-HALDB abends

yyyy   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:

4001	No more partitions are defined
4020	DB authorization error (data unavailable)
4025	DB unload action error
4026	DB initialization action error
4059	The access level is not compatible with the intent level
8001	Target partition requires initialization
8010	Target partition not found
8031	Open DMB failure
8051	User partition selection failed
C001	Non-partitioned DB access
C002	Partitions not active
C003	Partition stopped
C004	Partition structure terminating
C005	Partition structure not initialized
C006	Partition structure rebuild failed
C009	ILDS DSG not provided
C010	Unusual structure status
C030	Invalid authorization reason code
C040	Invalid partition DDIR
C041	Invalid partition set operation
C042	Unable to allocate PSETE storage
C043	Invalid partition set selection action
C050	User selection exit not invoked
C060	Invalid function code
C070	Buffer handler byte locate failed
C071	Buffer handler locate by key failed
C073	Buffer handler mark altered failed
C074	Buffer handler byte locate & mark altered failed
C075	IRLM root lock request (GRID) failed
C076	Request to obtain BCB storage (BCBGET) failed

PARTITION | DATABASE
is PARTITION if HALDB, or DATABASE if non-HALDB

zzzzzzzz
is the DBD name for either:
- 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 HALDB partition.

System Action: IMS continues processing.

Module: DFSORP00

Note: Message DFS2970I is issued after OM INIT OLREORG command is processed to the extent that Online Reorganization ownership is established in RECON. If the Online Reorganization cursor is active then RESUMED is indicated; otherwise, STARTED is indicated.

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

Explanation: HALDB online reorganization has terminated for the named HALDB partition. Refer to the following return and reason codes to determine why the termination occurred.

Problem Determination:

Return Code (Hex)	Description
04	HALDB online reorganization internal error.

	Reason Code (Hex)	Description
	04	ISWITCH target failed.

08		DFSBCB GET QSAV failed.
12		Create work unit failed.
16		IMS Cold start.
20		Restart AUTH failed, or XRF deferred STOP/CLOSE.

System Action: IMS continues processing.

System Programmer Response: Since HALDB online reorganization must complete it can be resumed with the INIT OLREORG command or offline reorganization will need to be run.

Problem Determination: 1, 6, 27

Module: DFSORP00

Problem Determination:

Return Code (Hex)	Description
08	HALDB online reorganization internal error.

Reason Code (Hex)	Description
01	GETMAIN for the DFSOLRW area failed.
02	Schedule call to DFSOSCH0 failed.
03	CREATE for output datasets failed.

Return Code (Hex)	Description
0C	PART=NAMED failed. Following are the list of reason codes from DFSPSEL0.

Reason Code (Hex)	Description
8031	OPEN failure.
C030	Invalid AUTH reason code.

Return Code (Hex)	Description
10	HALDB online reorganization internal error.

Reason Code (Hex)	Description
abend code	PART=NAMED failed. User abend returned.

System Action: IMS continues processing.

System Programmer Response: Since HALDB online

reorganization must complete it can be resumed with the INIT OLREORG command, or offline reorganization will need to be run.

Problem Determination: 1, 6, 27

Module: DFSORP10

Problem Determination:

Return Code (Hex)	Description
14	HALDB online reorganization internal error.

Reason Code (Hex)	Description
01	Terminate OLR was initiated.
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 call failure during Cursor update.
07	Notify failure during Cursor update.

Return Code (Hex)	Description
18	HALDB online reorganization internal error.

Reason Code (Hex)	Description
abend code	Failure in DFSORP20. Pseudo abend was set.

System Action: IMS continues processing.

System Programmer Response: Since HALDB online reorganization must complete it can be resumed with the INIT OLREORG command, or offline reorganization will need to be run.

Problem Determination: 1, 6, 27

Module: Issued by DFSORP10 for errors detected in DFSORP20.

Problem Determination:

Return Code (Hex)	Description
1C	HALDB online reorganization internal error.

Reason Code (Hex)	Description
01	Backout failure. No pseudo abend.
02	Terminate force was initiated.
03	All 'F's key found, or failure doing chain maintenance.

20   HALDB online reorganization internal error.

Reason Code (Hex)	Description
ABEND CODE	Failure in DFSORP30. Pseudo abend was set.

24   HALDB online reorganization internal error.

Reason Code (Hex)	Description
STATUS CODE	Bad status returned on call.

System Action: IMS continues processing.

System Programmer Response: Since HALDB online reorganization must complete it can be resumed with the `INIT OLREORG` command, or offline reorganization will need to be run.

Problem Determination: 1, 6, 27

Module: Issued by DFSORP10 for errors detected in DFSORP30.

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.
- Recovering a data set created by a HALDB online reorganization.

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-thru-J and X or the M-thru-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline. If this error occurred recovering a data set, the recovery is not performed.

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, refer to the appropriate publication in the following table to determine the cause of the error, and follow any recommended actions:

Macro	Publication
DEVTYPE	*DFSMS/MVS DFSMSdfp Advanced Services*
GETDSAB	*z/OS MVS Programming: Authorized Assembler Services Reference, Vol 2*
LOCATE	*DFSMS/MVS DFSMSdfp Advanced Services*

Macro	Publication
OBTAIN	*DFSMS/MVS DFSMSdfp Advanced Services*
SWAREQ	*z/OS MVS Programming: Authorized Assembler Services Reference, Vol 4*
TRKCALC	*DFSMS/MVS DFSMSdfp Advanced Services*

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.

If the error was detected in recovering a data set, reissue the `/RECOVER` command for recovery of the data set.

Module: DFSORA80

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 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 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.

Case 3: In recovering a data set that was the output of a HALDB online reorganization for the HALDB partition *partname*, IMS found that the data set into which the data was to be recovered either:

- Was not of a type that could used for the applicable data set.
- Had certain characteristics that did not match those of the corresponding input data set that was to be recovered.
- 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, the tracking of the reorganization at the RSR tracking site, or the recovery of the output data set from a reorganization.

Table 20. Reason Codes, Environments, and Errors Detected for Message DFS2991I

Reason Code	Environment	Errors Detected for Pre-Existing Data Set
1	All	The data set is not a DASD data set.
2	All	The data set is a VSAM data set, but the	database is defined to use OSAM.
3	All	The data set is a PDS, PDSE, or a member of a PDS or PDSE.
4	All	The data set is not a VSAM ESDS, but the database is defined to use VSAM.
5	All	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 ESDS 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.
9	All	The primary index data set of a PHIDAM database is not a VSAM KSDS.
10	Reorganization	The primary index data set of a PHIDAM database, that is, the KSDS, is defined with a different key offset or key length than that of the corresponding input data set.

Table 20. Reason Codes, Environments, and Errors Detected for Message DFS2991I (continued)

Reason Code	Environ- ment	Errors Detected for Pre-Existing 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 ESDS 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, that is, the KSDS, is defined with a different key offset or key length than that of the tracked output data set at the active site.
15	Recovery	The VSAM data set has a different record length than that of the data set that was to be recovered.
16	Recovery	The VSAM ESDS has a different control interval size than that of the data set that was to be recovered.
17	Recovery	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 data set that was to be recovered.
18	All	Either: • The data set is cataloged but it is not on the volume to which the catalog entry refers. or • The data set was allocated with a DD statement that had a VOLUME parameter, but the data set was not on the referenced volume.
19	All	The data set was allocated with a DD statement, but the data set name in the DSNAME parameter does not follow the data set naming convention defined for this partition's data sets.
20	Tracking, Reorg- anization	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.
21	Recovery	The data set being recovered does not exist, but the original data set resided on multiple volumes and was 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-thru-J and X or the M-thru-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline. If this error occurred recovering a data set, the recovery is not performed.

In the first two 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, as well as 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 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-thru-J or the M-thru-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-thru-J or the M-thru-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, as well as 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-thru-J or the M-thru-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-thru-J or the M-thru-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.

Case 3: Before restarting the recovery of the data set, either:

- Delete the preexisting data set into which the data was being recovered, 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, as well as 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 recovering the output of the reorganization:

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

- Is a VSAM key-sequenced data set (KSDS).
- Has the REUSE attribute.
- Is neither a PDS, PDSE, nor a member of a PDS or PDSE or that for a VSAM-defined database the data set for the database segments (one of the A-thru-J or the M-thru-V data sets):
  - 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 data set that was the output of the reorganization (as shown in the X'2930' log record) or that for a PHIDAM database the data set for the primary index (the X or Y data set):
    - Is a VSAM key-sequenced data set (KSDS).
    - Has the REUSE attribute.
    - Has the same fixed-length record length as the data set that was the output of the reorganization (as shown in the X'2930' log record).
    - Has the same key offset and key length as the data set that was the output of the reorganization (as shown in the X'2930' log record).

Then reissue the /RECOVER command for recovery of the data set.

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.
- Recovering a data set created by a HALDB online reorganization.

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 IDC3009I.

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-thru-J and X or the M-thru-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline. If this error occurred recovering a data set, the recovery is not performed.

In the first two 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.

If the error was detected in recovering a data set, reissue the `/RECOVER` command for recovery of the data set.

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: DFSORA80

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.
- Recovering a data set created by a HALDB online reorganization.

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-thru-J and X or the M-thru-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline. If this error occurred recovering a data set, the recovery is not performed.

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

Programmer Response: Refer to 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.

If the error was detected in recovering a data set, reissue the `/RECOVER` command for recovery of the data set.

Module: DFSORA80

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-thru-J and X or the M-thru-V and Y data sets.

Programmer Response: Refer to *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: DFSORA80

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.
• Recovering a data set created by a HALDB online reorganization.

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, refer to *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, refer to *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-thru-J and X or the M-thru-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline. If this error occurred recovering a data set, the recovery is not performed.

In the first 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. In there is 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 Online Forward Recovery to continue updating the partition.

If the error was detected in recovering a data set, reissue the `/RECOVER` command for recovery of the data set.

Module: DFSORA80

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.
*   Recovering a data set created by a HALDB online reorganization.

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; this caused 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-thru-J and X or the M-thru-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline. If this error occurred recovering a data set, the recovery is not performed.

In the first two 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 isn't 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

Online Forward Recovery to continue updating the partition.

If the error was detected in recovering a data set, reissue the `/RECOVER` command for recovery of the data set.

Module: DFSORA80

DFS2998E   [DATA SET CREATION ERROR. | 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.
*   Recovering a data set created by a HALDB online reorganization.

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, refer to *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, refer to *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-thru-J and X or the M-thru-V and Y data sets. If this error occurred at the RSR tracking site, then the shadow partition is taken offline. If this error occurred recovering a data set, the recovery is not performed.

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.

If the error was detected in recovering a data set, reissue the `/RECOVER` command for recovery of the data set.

Module: DFSORA80

HALDB OLR DSP Series Messages

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

DSP0197I   ONLINE REORG NOTIFICATION NOT VALID FOR DATABASE ORGANIZATION

Explanation: A command to update RECON regarding an online reorg 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 additional assistance is required.

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 reorg 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(AIM) is not allowed if online reorg is active.

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

Programmer Response: Contact the IBM Support Center if additional assistance is required.

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 datasets are active.

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

Programmer Response: Contact the IBM Support Center if additional assistance is required.

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 additional assistance is required.

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 reorg cannot be used as input to recovery or the recovery-to-time falls within an online reorganization and timestamp recoveries are not allowed.

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

Programmer Response: Contact the IBM Support Center if additional assistance is required.

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 additional assistance is required.

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 database data set (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 dataset 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 additional assistance is required.

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 additional assistance is required.

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:

System Action: Processing is terminated. The

RECON data sets are not updated.

Programmer Response: Contact the IBM Support Center if additional assistance is required.

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 additional assistance is required.

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 new Error Codes for Abend U381.

Abend U0880 There are some new user abend reason codes for Abend U0880.

Related Reading: For detailed information on these or other abends, see the *IMS Messages and Codes, Volume 1* and the *IMS Failure Analysis and Structure Tables (FAST) for Dump Analysis*.

Chapter 8. Log Records of HALDB OLR

The following are the log records changes that have occurred as a result of HALDB OLR.

Related Reading:

- See the section entitled,"Logging Effects of Reorganization" on page 12 for more important information on performance and logging of HALDB OLR.
- See the *IMS Diagnosis Guide and Reference* for more details on these log records.

DFSDLOG - 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

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 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.

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.

Additional Log Records: x'29' OLR Event Log Records

Several additional x'29' log records reflect the progress of a HALDB OLR.

The following new log records are created at the indicated points in the online reorganization process.

Table 21. X'29' OLR Event Log Record Subtypes

TYPE or SUB-TYPE	WHY WRITTEN	ISSUING 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 DFSRDBL0
x'2940'	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'	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.

Part 4. Appendixes

RECON Listing HALDB OLR Support

The following are the LIST command output (RECON) record modifications for HALDB OLR support:

HALDB MASTER RECORD

Here is an example of a HALDB Master record listing. 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.

```
-----------------------------------------------------------------------
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
   OLREORG 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 OLR 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
   OLREORG CAPABLE      =YES       DATA SET GROUP MEMBERS  =1


Partition DB record:

DB
```

91

```
          DBD=POHIDJA   MASTER DB=DBHIDJ03  IRLMID=*NULL  CHANGE#=2     TYPE=PART
          USID=0000000003  AUTHORIZED USID=0000000003  HARD USID=0000000003
          RECEIVE USID=0000000003  RECEIVE NEEDED USID=0000000000
          DBRCVGRP=**NULL**
          DSN PREFIX=IMSTESTS.DBHIDJ03                       PARTITION ID=00001
          PREVIOUS PARTITION=**NULL**    NEXT PARTITION=POHIDJB
          OLRIMSID=**NULL**    ACTIVE DBDS=A-J

        FREE SPACE:
          FREE BLOCK FREQ FACTOR=0    FREE SPACE PERCENTAGE=0

        PARTITION HIGH KEY/STRING (CHAR):            (LENGTH=5  )
           J2000
        PARTITION HIGH KEY/STRING (HEX):
           D1F2F0F0F04040404040404040404040404040404040404040404040404040404040

        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   =0
                                              HELD AUTHORIZATION STATE=0
                                              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
          OLREORG CAPABLE       =YES
```

The RANDOMIZER values are only listed for PHDAM databases.

If this HALDB uses a partition selection routine: label "PARTITION HIGH KEY:" is listed as "PARTITION SELECTION STRING:". Also, if the high key or selection string contains any unprintable characters, the entire string is listed in hexadecimal format. The key or string is formatted into as many 64-character lines as are necessary. "PARTITION HIGH KEY (HEX):" is listed if the key or string contains hex values.

The OSAM block size information section is only listed if the HALDB is an OSAM database. Sizes are listed for the defined data set group members only.

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    = 1999.364 11:59:12.0 -08:00    *  RUN USID     = 0000000037
     STOP   = 2000.001 12:34:56.7 -08:00    ONLINE  RECOV=YES
    ---------------------------------------------------------------------
```

Index

Special characters

R

RANDOMIZER values   92
RATE option
   UPDATE OLREORG   56
RATE parameter   33, 42
reason codes
   INITIATE OLREORG   35
   QUERY OLREORG   43
reclustering of data   3
RECON records   91
RECOV parameter   39
REORG record listing   92
reserved keywords   32, 41, 51, 57
resume online reorganization   31
return codes
   INITIATE OLREORG   35
   QUERY OLREORG   43
RUNNING parameter   43

S

secondary space allocation   29
SET keyword   33
SHAREPLEX, IMS   xii, 3
sharing system acknowledgement   30
SHOW keyword   42
single volume data sets   28
slash commands   31
SMS-managed data sets   28
space allocation
   primary   29
   secondary   29
STATUS keyword   42
STATUS parameter   42
status, cursor-active   xii, 4
STOP=   92
STOPTIME   39
storage class   29
storage volumes   28
syntax
   INITIATE OLREORG command   32
   TERMINATE OLREORG command   51
   UPDATE OLREORG command   57

T

TERMINATE OLREORG
   completion codes   53, 54
TERMINATE OLREORG command   49
TERMINATE OLREORG command ABORT option   85
type=HALDB   63

U

unit of reorganization   xiii
units of reorganization   4
UNKNOWN parameter   43
UPDATE OLREORG
   completion codes   60
UPDATE OLREORG command   55

UPDATE OLREORG command processing   56
usage
   INITIATE OLREORG command   32
   QUERY OLREORG command   42

V

values, RANDOMIZER   92
volumes, storage   28
VSAM data sets   27

X

x'20' log record   87
x'21' log record   87
x'29' log record   87
x'2910' log record   88
x'2920' log record   88
x'2930' log record   88
x'2940' log record   88
x'2950' log record   88
x'2970' log record   88
x'2990' log record   88

Z

zoikels   31

Bibliography

This bibliography includes all the publications cited in this book, including the publications in the IMS library.

IMS Version 9 Library

ZES1-2330   ADB
ZES1-2331   AS
ZES1-2332   ATM
ZES1-2333   APDB
ZES1-2334   APDG
ZES1-2335   APCICS
ZES1-2336   APTM
ZES1-2337   BPE
ZES1-2356   CR
ZES1-2338   CQS
ZES1-2339   CSL
ZES1-2340   CG
ZES1-2341   DBRC
ZES1-2342   DGR
ZES1-2343   FAST
ZES1-2345   ORG
ZES1-2346   JGR
ZES1-2347   IIV
ZES1-2348   ISDT
ZES1-2349   INTRO
ZES1-2351   MIG
ZES1-2352   MC1
ZES1-2353   MC2
ZES1-2354   OTMA
ZES1-2355   OG
ZES1-2357   RPG

ZES1-2359   URDBTM
ZES1-2360   URS

Supplementary Publications

GC17-7825   LPS
ZES1-2357   SOC

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
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

IBM®

Program Number: 5655-J38

ZES1-2346-00