

# REORFILE: Reorder File

The REORFILE function reorders the Associator and Data Storage for a single file. Associator and Data Storage for other files are not affected.

This function requires exclusive EXU control of the database files involved in the operation. In addition, parts of the database are overwritten during ADAORD execution, so we recommend that you back up the database (or file) using the ADASAV utility first, before running ADAORD functions.

If the file specified for this function was originally loaded with ISNREUSE=YES active, this reorder function will reset the first unused ISN value in that file's control block (FCB) to the actual first unused ISN found in the address converter.

This is the syntax of the ADAORD REORFILE function:

```

ADAORD REORFILE FILE = file-number
    [ACRABN = starting-rabn ]
    [AC2RABN = starting-rabn ]
    [ALLOCATION = { FORCE | NOFORCE } ]
    [ASSOPFAC = padding-factor ]
    [ASSOVOLUME = 'Associator-extent-volume' ]
    [DATAPFAC = padding-factor ]
    [DATAVOLUME = 'Data-Storage-extent-volume' ]
    [DSDEV = device-type ]
    [DSRABN = starting-rabn ]
    [DSRELEASE]
    [DSSIZE = size ]
    [INDEXCOMPRESSION = { YES | NO } ]
    [INSIZE = { 3 | 4 } ]
    [MAXISN = highest-isn ]
    [MAXISN2 = highest-isn ]
    [MAXRECL = record-length ]
    [NIRABN = starting-rabn ]
    [NIRELEASE]
    [NISIZE = size ]
    [PASSWORD = password ]
    [SORTSEQ = { descriptor | ISN } ]
    [UIRABN = starting-rabn ]
    [UIRELEASE]
    [UIsize = size ]
    [EXCLUDE = file-list ]
    [LIP = { isn-pool-size | 16384 } ]
    [LPB = prefetch-buffer-size } ]
    [NOUSERABEND]
    [TEST]

```

This chapter covers the following topics:

- Essential Parameter
  - Optional Parameters
  - Examples
- 

## Essential Parameter

### **FILE: File Number**

FILE specifies the file to be processed, and to which the parameters that follow in the statement sequence apply. Several files and their related parameters may be specified within one REORFILE operation; see the examples at the end of this section. If a component file of an Adabas expanded file is specified, only that file's Associator and Data Storage are reordered; this has no adverse effect on the other component files.

## Optional Parameters

### **ACRABN: Starting RABN for Address Converter**

ACRABN is the beginning RABN for the file's address converter extent. If this parameter is omitted, ADAORD assigns the starting RABN. The space requested must be available in one extent.

When specifying the starting RABN for Associator extents, the space needed for the FCBs, FDTs, and DSST should also be considered.

### **AC2RABN: Starting RABN for Secondary Address Converter**

The beginning RABN for the file's secondary address converter extent. The secondary address converter is used to map the secondary ISNs of secondary spanned records to the RABNs of the Data Storage blocks where the secondary records are stored.

If this parameter is omitted, ADAORD assigns the starting RABN. The space requested must be available in one extent. If the file contains no secondary address converter extents, this parameter is ignored.

### **ALLOCATION: Action to Follow File Extent Allocation Failure**

ALLOCATION specifies the action to be taken if file extent allocations cannot be obtained according to the placement parameters ACRABN, DSRABN, NIRABN, or UIRABN.

By default (that is, ALLOCATION=FORCE), the utility terminates with error if any file extent allocation cannot be met according to RABN placement parameters.

If ALLOCATION=NOFORCE is specified and any allocation with placement parameters fails, the utility retries the allocation without the placement parameter.

### **ASSOPFAC: Associator Padding Factor**

ASSOPFAC specifies the new Associator block padding factor. The number specified represents the percentage of each Associator block not to be used during the reorder process. A value in the range 1-90 may be specified. The remaining number of bytes after padding must be greater than the largest descriptor value plus 10.

If this parameter is omitted, the current Associator padding factor in effect for the file is used.

### **ASSOVOLUME: Associator Extent Volume**

**Note:**

The value for the ASSOVOLUME parameter must be enclosed in apostrophes.

ASSOVOLUME identifies the volume on which the file's Associator space (that is, the AC, NI, and UI extents) should be allocated. If the requested number of blocks cannot be found on the specified volume, ADAORD allocates the remaining blocks on other volumes according to its default allocation rules.

If ACRABN, UIRABN, or NIRABN is specified, ADAORD ignores the ASSOVOLUME value when allocating the corresponding extent type.

If ASSOVOLUME is not specified, the file's Associator space is allocated according to ADAORD's default allocation rules.

### **DATAPFAC: Data Storage Padding Factor**

DATAPFAC specifies the new Data Storage padding factor, which is the percentage of each Data Storage block reserved for record expansion when the file is reordered. A value in the range 1-90 may be specified (see the ADALOD LOAD DATAPFAC parameter discussion for more information about setting the padding factor). If this parameter is omitted, the current padding factor for the file is used.

### **DATAVOLUME: Data Storage Extent Volume**

**Note:**

The value for the DATAVOLUME parameter must be enclosed in apostrophes.

DATAVOLUME specifies the volume on which the file's Data Storage space (DS extents) are allocated. If the number of blocks requested with DSSIZE cannot be found on the specified volume, ADAORD allocates the remaining blocks on other volumes according to its default allocation rules.

If DSRABN is specified, DATAVOLUME is ignored for the related file.

If DATAVOLUME is not specified, the Data Storage space is allocated according to ADAORD's default allocation rules.

### **DSDEV: Data Storage Device Type**

DSDEV specifies the device type to be used for the file's Data Storage. The specified device type must already be defined to Adabas, normally when the database was created or by the ADADBS utility's ADD function.

If this parameter is not specified, ADAORD attempts to allocate the file on the device type used before reordering.

**DSRABN: Data Storage Starting RABN**

DSRABN specifies the beginning RABN for the file's Data Storage extent. If the DSRABN parameter is omitted, ADAORD assigns the starting RABN.

**DSRELEASE: Release Unused Data Storage Blocks**

Specifying DSRELEASE releases unused Data Storage (DS) blocks belonging to the file. If DSRELEASE is not specified, ADAORD allocates *at least* the number of DS blocks that were allocated before the file was reordered.

**Note:**

Adabas calculates the file extent size using any changed padding factor or block size values *before* the file is reordered.

**DSSIZE: Data Storage Size**

DSSIZE specifies the number of blocks or cylinders to be allocated for the Data Storage. A block count must be followed by a "B" (for example, "2000B").

If this parameter is omitted, ADAORD calculates the file extent size in proportion to any increase or decrease in the DATAPFAC padding factor.

**EXCLUDE: Exclude Specified Files from Reorder**

EXCLUDE lists the numbers of the files to be excluded from REORDER processing; that is, the files that are not to be reordered.

The parameter is optional: if not specified, no files are excluded. A file number may be listed only once.

Files specified in the EXCLUDE parameter must also be specified in the FILE parameter.

The EXCLUDE parameter is provided for use in recovery jobs built by the Adabas Recovery Aid (ADARAI).

**INDEXCOMPRESSION: Compress File Index**

INDEXCOMPRESSION indicates whether the index for the file is rebuilt in compressed or uncompressed form. A compressed index usually requires less index space and improves the efficiency of index operations in the Adabas nucleus.

If INDEXCOMPRESSION is not specified, the default is the current form of the file.

**ISNSIZE: 3- or 4-Byte ISN**

ISNSIZE specifies whether ISNs in the file are to be 3 or 4 bytes long. The default is the value currently used for the file; this value is stored in the file control block (FCB).

**Note:**

It is not possible to change the ISNSIZE of a physically coupled file using ADAORD.

**LIP: ISN Buffer Pool Size**

The LIP parameter can be used to decrease the number of Associator I/O operations when recreating the address converter. For best performance, specify a size that accepts all ISNs of the largest file to be processed.

LIP specifies the size of the ISN pool for containing ISNs and their assigned Data Storage RABNs. The value may be specified in bytes as a numeric value ("2048") or in kilobytes as a value followed by a "K" ("2K"). The default for LIP is 16384 bytes (or 16K).

The length of one input record is ISNSIZE + RABNSIZE. Thus the entry length is at least 6 bytes (the ISNSIZE of the file is 3 and the RABNSIZE of the database is 3) and at most 8 bytes (the ISNSIZE is 4 and the RABNSIZE is 4).

**Note:**

When ADAORD is processing files that contain spanned records with secondary ISNs, a second LIP will be allocated to contain these ISNs.

**LPB: Prefetch Buffer Size**

LPB specifies the size, in bytes, of the internal prefetch buffer. The maximum value is 32760 bytes. The default depends on the ADARUN LU parameter. ADAORD may reduce a specified LPB value if the LU value is too small.

**MAXISN: Highest ISN Permitted for the File**

MAXISN is the highest ISN which may be allocated for the file. This value must be greater than the current TOPISN value displayed in the ADAREP database report.

ADAORD uses the specified value to calculate the address converter space required. If this parameter is omitted, the current MAXISN value for the file is retained.

**MAXISN2: Highest Secondary ISN Permitted for the File**

MAXISN2 specifies the desired size of the secondary address converter (AC2) in ISNs. This value must be greater than the current TOP AC2 ISN value displayed in the ADAREP database report. The secondary address converter is used to map the secondary ISNs of secondary spanned records to the RABNs of the Data Storage blocks where the secondary records are stored.

ADAORD uses the specified value to calculate the space required in the secondary address converter for the file. If this parameter is omitted, the current MAXISN2 value for the file is retained. If the file contains no secondary address converter extents, this parameter is ignored.

**MAXRECL: Maximum Compressed Record Length**

Use the MAXRECL parameter to change the maximum record length, after compression, permitted in the file. Specifying MAXRECL has two effects:

- The DATA data set for the file can be allocated only to devices that support the specified length.
- If the file contains Data Storage records that exceed the specified length, ADAORD abends and prints ERROR-126 (Data Storage record too long).

If MAXRECL is not specified, the maximum compressed record length does not change.

#### **NIRABN: Starting RABN for Normal Index**

NIRABN is the beginning RABN for the normal index extent. If this parameter is omitted, ADAORD assigns the starting RABN.

#### **NIRELEASE: Release Unused Normal Index Blocks**

Specifying NIRELEASE releases unused normal index (NI) blocks belonging to the file. If NIRELEASE is not specified, ADAORD allocates *at least* the number of NI blocks that were allocated before the file was reordered.

##### **Note:**

Adabas calculates the file extent size using any changed padding factor or block size values *before* the file is reordered.

#### **NISIZE: Normal Index Size**

NISIZE specifies the number of blocks or cylinders to be allocated for the file's normal index. A block count must be followed by a "B" (for example, "2000B").

If this parameter is omitted, ADAORD computes the file extent size in proportion to any increase or decrease in the ASSOPFAC padding factor.

#### **NOUSERABEND: Termination without Abend**

When a parameter error or a functional error occurs while this utility function is running, the utility ordinarily prints an error message and terminates with user abend 34 (with a dump) or user abend 35 (without a dump). If NOUSERABEND is specified, the utility will *not* abend after printing the error message. Instead, the message "*utility* TERMINATED DUE TO ERROR CONDITION" is displayed and the utility terminates with condition code 20.

##### **Note:**

When NOUSERABEND is specified, we recommend that it be specified as the first parameter of the utility function (before all other parameters). This is necessary to ensure that its parameter error processing occurs properly.

#### **PASSWORD: File Password**

If the file is password-protected, use this parameter to specify the password.

#### **SORTSEQ: Reorder Sequence**

SORTSEQ determines the sequence in which the file is to be processed. If this parameter is omitted, the records are processed in physical sequence.

If a descriptor is specified, the file is processed in the logical sequence of the descriptor values. *Do not* use a null-suppressed descriptor field, a hyperdescriptor, a phonetic descriptor, a multiple-value field, or a descriptor contained in a periodic group.

**Note:**

Even when the descriptor field is not null suppressed, the record is *not* represented in the inverted list if the descriptor field or a field following it has never been initialized (held a value). Therefore, the record will be dropped when the utility is executed.

If ISN is specified, the file is processed in ascending ISN sequence. For the Adabas checkpoint or security file, only SORTSEQ=ISN is allowed.

**TEST: Test Syntax**

This parameter tests the operation syntax without actually performing the operation. Only the syntax of the specified parameters can be tested; not the validity of values and variables.

**UIRABN: Starting RABN for Upper Index**

UIRABN is the beginning RABN for the file's upper index extent. If this parameter is omitted, ADAORD assigns the starting RABN.

**UIRELEASE: Release Unused Upper Index Blocks**

Specifying UIRELEASE releases unused upper index (UI) blocks belonging to the file. If UIRELEASE is not specified, ADAORD allocates *at least* the number of UI blocks that were allocated before the file was reordered.

**Note:**

Adabas calculates the file extent size using any changed padding factor or block size values *before* the file is reordered.

**UI SIZE: Index Size for Upper Index**

UI SIZE specifies the number of blocks or cylinders to be allocated for the upper index. A block count must be followed by a "B" (for example, "2000B").

If this parameter is omitted, ADAORD computes the file extent size in proportion to any increase or decrease in the ASSOPFAC padding factor.

## Examples

**Example 1:**

```
ADAORD REORFILE FILE=16
```

Associator and Data Storage for file 16 are to be reordered. All current extent sizes and padding factors for the file are to be retained. No other files are reordered.

**Example 2:**

ADAORD	REORFILE	FILE=246
ADAORD		DATAFFAC=5 ,DSSIZE=10B ,SORTSEQ=MZ
ADAORD		ASSOPFAC=20 ,MAXISN=5000
ADAORD		FILE=20

File 246 is to be reordered; the Associator and Data Storage for all other files remain unchanged. The Data Storage padding factor is to be set to 5. A new Data Storage size of 10 blocks is to be used. The Data Storage is to be reordered in the logical sequence of descriptor MZ. The new Associator padding factor is 20. The highest ISN which may be assigned is 5000. File 20 is to be reordered with no changes to its current parameters.

**Example 3:**

ADAORD	REORFILE	FILE=9
ADAORD		ASSOPFAC=5 ,DATAFFAC=15
ADAORD		FILE=23
ADAORD		DSRABN=24032
ADAORD		UIRABN=3151 ,UISIZE=50B
ADAORD		NIRABN=3201

File 9 is to be reordered. The Associator padding factor is set to 5% and the Data Storage padding factor to 15%.

File 23 will also be reordered, with a new starting Data Storage RABN of 24032. In addition, the following Associator changes are being made for file 23: the new upper index starting RABN is 3151, with a new upper index size of 50 blocks. The new normal index starting RABN is 3201; the size remains the same.