

Utility Changes

All Adabas utilities have been updated to support the new extended features of Adabas 8. Some of this support appears in the form of new or modified utility parameters. In other cases, support was added internally that does not affect your use of the utility at all.

Note:

Be sure to read about the utility restrictions in *Limitations and Restrictions* for a brief description of what the limitations and restrictions are for utility use in Adabas 8.

This chapter describes the Adabas utilities whose user interface has changed in this release. For complete information about the functions of any Adabas utility, read *Utilities*.

- ADAACK Utility Changes
- ADACDC Utility Changes
- ADACMP Utility Changes
- ADADBS Utility Changes
- ADADCK Utility Changes
- ADAFRM Utility Changes
- ADAICK Utility Changes
- ADALOD Utility Changes
- ADAORD Utility Changes
- ADAREP Utility Changes
- ADASAV Utility Changes
- ADASEL Utility Changes
- ADAULD Utility Changes

ADAACK Utility Changes

Spanned record support is provided in the ADAACK utility. However, ADAACK assumes that an ISN passed to it is a primary ISN or is the only ISN for a record. If an ISN is the primary ISN for a spanned record, all of the associated segment records of the spanned record are automatically checked in the secondary address converter.

When printing error information about a particular ISN, the ADAACK utility will now indicate whether the problem is with a primary or secondary ISN, if the record is spanned.

ADACDC Utility Changes

Spanned records are not supported by the ADACDC utility at this time. However, when the IGNORESPANNED parameter is specified in an ADACDC run, ADACDC processing ignores spanned records, issues a warning message, and continues its processing. A return code of "4" is returned.

ADACMP Utility Changes

- General Changes
- ADACMP COMPRESS Changes
- ADACMP DECOMPRESS Changes

General Changes

Traditionally, the DD/FEHL error data set produced for ADACMP errors has truncated rejected records that exceeded the FEHL physical record length. In Version 8, the rejected records are segmented instead of truncated. Because of this change, the DD/FEHL LRECL setting must be at least 500 bytes.

ADACMP COMPRESS Changes

The following *new* parameters have been added to the ADACMP COMPRESS utility to support the MU/PE extension, spanned records, and LB fields:

- The DATADEVICE parameter specifies the Data Storage device type to be used for the segmentation of spanned records. If the SPAN parameter is specified, ADACMP will break long, spanned, compressed records into segments that are just a bit smaller than the Data Storage block size implied by the DATADEVICE parameter.

If DATADEVICE is specified without the SPAN parameter, it is used to derive the maximum compressed record length that will be accepted. Longer compressed records will be considered in error and will be written to the DD/FEHL data set.

- The HEADER parameter indicates whether or not the ADACMP compression logic should expect segmented ADACMP record headers in the uncompressed input records. The default is NO. (Contrast this with the HEADER parameter introduced in this release for ADACMP DECOMPRESS.)
- The LOBDEVICE parameter identifies the device type that will be used for loading the *LOB file* produced by the COMPRESS function.
- The LOBVALUES parameter indicates whether the uncompressed input data may contain long LB values (larger than 253 bytes).
- The MAXLOGRECLLEN parameter can be used to specify the size, in bytes, of a buffer used by ADACMP to assemble the physical uncompressed segmented records into logical compressed records. This buffer is allocated and used only if HEADER=YES is also specified.

- The MUPEX parameter indicates whether the extended MU/PE limits are allowed for the file. If this option is not specified, the maximum number of MU fields and the maximum number of PE fields that can be specified is 191.
- The MUPECOUNT parameter specifies the size of the value count field in the input record for the COMPRESS function. Valid values are "1" or "2". If "1" is specified, each value count field preceding the MU or PE values in the input data must be one byte with a value of no more than "191". If "2" is specified, each value count field preceding the MU or PE values in the input data must be two bytes. A value count may exceed 191 only if the MUPEX parameter is also specified.
- The SPAN parameter allows the record to be spanned after it is compressed if the compressed record exceeds the data storage block size of the device.

The following *changes* have been made to existing ADACMP COMPRESS parameters:

- The report produced by the DEVICE parameter now includes an indication of whether or not the MUPEX parameter has been set for a file.
- The syntax of the FNDEF parameter has been changed, to allow you to specify the number of occurrences of the MU and PE options. You can also specify the NB and LB field options.
- The MAXPE191 parameter is no longer supported for the ADACMP COMPRESS function in Adabas 8. When specified, a warning message is issued and processing continues.
- If USERISN is specified with HEADER=YES, the ISN immediately follows the ADAH header as part of the logical record.

When running ADACMP COMPRESS for a file that includes LB fields, a second sequential output data set (identified by the DDAUSB1 JCL DD control statement) is produced for the *LOB file* containing the LB field values.

ADACMP DECOMPRESS Changes

ADACMP DECOMPRESS processing supports the extended MU and PE limits and spanned records as input. The size of the value count preceding MU or PE values in each decompressed output record depends upon the extended MU and PE support for a file. A two-byte value count is given for files with extended MU and PE support. A one-byte value count is given for files without extended MU and PE support. In addition, the following functionality is supported.

- The decompression of LB fields. A new LOBVALUES parameter has been added for this support.
- A new HEADER parameter has been added to the ADACMP DECOMPRESS utility to indicate whether or not the ADACMP decompression logic should produce the ADACMP segmented record headers (ADAH and ADAC) as part of the decompressed output. The default is NO.
- A new MAXLOGRECLLEN parameter can be used to specify the size, in bytes, of a buffer used by ADACMP to assemble logical records that span one or more physical records of uncompressed output data. This buffer is allocated and used only if HEADER=YES is also specified.
- The ISN parameter has been altered so that if it is specified with HEADER=YES, the ISN immediately follows the ADAH header as part of the logical record.

ADADBS Utility Changes

Three new database services have been added to the ADADBS utility to support the extended MU/PE field counts and record spanning:

- The MUPEX function allows you to specify the maximum number of occurrences allowed for MU or PE fields in a file.
- The RECORDSPANNING function activates the use of spanned records in a file.
- The SPANCOUNT function counts the spanned records in a file.

In addition, a new RESETPPT function resets the PPT blocks on the ASSO data set.

ADADCK Utility Changes

The ADADCK utility checks the header of a spanned record for plausibility, as follows:

- The ISNs listed in the header are verified. The header contains the ISN of the primary record in the spanned record chain, the ISN of the previous spanned record in the chain, and the ISN of the next spanned record in the chain.
- The spanned record identifier bits are checked. The header of a spanned record includes bits that indicate whether the record is a primary or secondary spanned record. Only one of these bits can be turned on in the header of any spanned record.

A new MAXPISN parameter has also been introduced that allows you to set the maximum number of primary ISNs that will be checked for a spanned Data Storage file. The default is 1000.

ADAFRM Utility Changes

The ADAFRM utility can now be used to clear multiple PLOG headers from the PLOG, without requiring that you reformat the entire PLOG. To do this, you should specify the NUMBER parameter in conjunction with the FROMRABN parameter and set the SIZE parameter to "1".

In addition, this utility can now handle the increased number of physical Associator and Data Storage extents (99) allowed in Adabas 8.

ADAICK Utility Changes

If you run the ADAICK DSCHECK function, the primary and secondary ISNs are now identified in the output.

ADALOD Utility Changes

The following new parameters are introduced for the ADALOD LOAD and UPDATE functions in support of spanned records and their associated secondary address converter:

- The AC2RABN parameter allows you to specify or update space allocation for the secondary address converter. 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.
- The optional MAXISN2 parameter allows you to specify the desired size of the secondary address converter (AC2) in ISNs. The secondary address converter is used to map secondary ISNs of secondary spanned records to the RABNs of the Data Storage blocks where the secondary records are stored.

The following new parameters and parameter values are introduced for the ADALOD LOAD functions in support of large object (LB) fields and their associated *LOB files*:

- The optional LOBFILE parameter allows you to specify the file number of the *LOB file* associated with a *base file*. This parameter is used when loading base files.
- The optional BASEFILE parameter allows you to specify the file number of the *base file* associated with a *LOB file*. This parameter is used when loading LOB files.
- A new file type, LOB, specified on the FILE parameter, allows you to indicate that you are loading an Adabas LOB file with a predefined FDT.

ADAORD Utility Changes

To support spanned records, the following two new parameters have been added to the REORASSO, REORDB, REORFASSO, REORFILE, and STORE functions of ADAORD.

- The AC2RABN parameter allows you to specify the beginning RABN for the file's secondary address converter extent.
- The optional MAXISN2 parameter allows you to specify the desired size of the secondary address converter (AC2) in ISNs.

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.

You can restructure databases and files from Adabas 5.1 - 5.3, 6.1, 6.2, 7.1, 7.2, or 7.4 and store them in Version 8 using ADAORD STORE. However, you cannot store the restructure output of an Adabas 8 database or file in a database running with any prior Adabas version (for example, Adabas 7).

You can restructure databases and files from an Adabas version prior to Adabas 8 and store them in an Adabas 8 database using ADAORD STORE. However, you cannot store the restructured output of an Adabas 8 database or file in a database running with any prior Adabas version (for example, Adabas 7). If you attempt this, the following warning will be generated and ADAORD will end with a CC=4:

```
*** Warning: The input dataset is from V8 and will not be processed
```

ADAREP Utility Changes

The report produced by ADAREP now lists whether the MUPEX and spanned record options have been set for the database.

In the "Contents of Database" section of the report, the following changes have been made:

- The padding factor has been removed from this table to make room for the larger extent values. However, it is still shown in the individual file details also provided in the report and it appears again if LAYOUT=1 is specified.
- When LAYOUT=1 is specified during report creation, the larger extent values, supported by Adabas 8, appear for each file.
- If a file is unable to build at least ten further file extents, ADAREP marks the file with an asterisk (*) on the right.

In the "File Options" section of the report, a "T" indicates that two-byte MU/PE indexes are active for the file and an "S" indicates that use of spanned records has been activated for the file. In addition, the "Contains LOB Fields" column indicates whether the file contains one or more LB fields (an "L" appears if it does) and the "LOB File" column (the last column) indicates whether or not the file is a LOB file (an "L" appears if it is). Note that these two LB field columns are mutually exclusive; only one of them will be marked.

In the "Physical Layout of the Database" section of the report, secondary address converter extents (for spanned records) are identified as "AC2" in the "Table File Type" column.

In the "File Information" section of the report, a new field called "Two Byte MU/PE" indicates whether two-byte MU/PE indexes are active for the file. In the same section, the highest, maximum expected, and minimum secondary ISN are given as well as a new field called "Spanned Rec Supp", which indicates whether or not spanned records are activated for the file. In addition, the "Contain LOB Fields" field indicates whether or not the file contains one or more LB fields and the "LOB File" field indicates whether or not the file is a LOB file.

In the Space Allocation section of the report, secondary address converter extents (for spanned records) are identified as "AC2" in the "List Type" column.

Finally, three new checkpoints may be written by the Adabas nucleus: 75, 76, and 77.

ADASAV Utility Changes

The following new parameters are introduced for the ADASAV RESTONL FMOVE and ADASAV RESTORE FMOVE functions in support of spanned records and their associated secondary address converter:

- The AC2RABN parameter allows you to specify the starting secondary address converter RABN for each file specified by FMOVE.
- The MAXISN2 parameter allows you to specify the new number of secondary ISNs to be allocated for each file specified by FMOVE.

ADASEL Utility Changes

While there are no new parameters, ADASEL has been updated to support the MU/PE extension. In particular, when specifying indexes for MU or PE fields in an ADASEL SELECT IF statement, the indexes can now range from "1" through "65,534". Prior releases of Adabas restricted these indexes to

values ranging from "1" through "191".

ADASEL recognizes spanned records in its processing, but it cannot process files containing spanned records.

ADAULD Utility Changes

While there are no new parameters, ADAULD has been updated to support the MU/PE extension and spanned records. In particular, two new statistics listing the number of record *segments* that were read and written during the run are produced by the ADAULD utility.