

ADAULD Utility

The ADAULD utility can unload an Adabas file from a full save tape together with 1 to 8 delta save tapes and optionally a DSIM dataset. The save tapes may have been created online or offline using any version of the Delta Save Facility. Adabas files are unloaded from a combination of full or delta save tapes to reestablish the files from archive save tapes and load them into a database.

Note:

It may not be possible to do this using the ADASAV RESTORE/RESTONL function if the archived database resided on device types that are no longer in use in the data cluster.

The records are unloaded in physical sequence; that is, in the order in which they are physically positioned within Data Storage.

The unloaded record output is in compressed format. The output records have the same format as the records produced by the Adabas ADACMP utility.

When using the MODE=SHORT option, descriptor entries (which are required to create the normal index and upper index for the file) are omitted during the unload process. This reduces the time required for unloading. Note, however, that output created using MODE=SHORT has a different FDT from the same file unloaded without MODE=SHORT, since all descriptor information is removed.

In general, the Adabas nucleus does not need to be active while unloading a file from a save tape.

Note:

An interrupted ADAULD UNLOAD FILE run must be reexecuted from the beginning.

This chapter covers the following topics:

- UNLOAD FILE: Unload Specified File from a Save Tape
- Save Tape Input Processing
- ADAULD Output Processing
- ADAULD User Exit 9
- JCL/JCS Requirements and Examples

UNLOAD FILE: Unload Specified File from a Save Tape

Syntax

```

ADAULD [UNLOAD] FILE = file-number
                SAVETAPE
                [ CODE = cipher-key ]
                [ DDISN ]
                [ DSIMDEV = device-type | ADARUN-device ]
                [ MODE = SHORT ]
                [ NOUSERABEND ]
                [ NUMOUT = 1 | 2 ]
                [ NUMREC = number ]
                [ PATTERN = merge-pattern ]
                [TEST]

```

Essential Parameters

FILE

FILE specifies the number of the file to be unloaded. Neither the checkpoint file nor the security file can be unloaded.

SAVETAPE

SAVETAPE is used to unload a file from a full save tape together with 1 to 8 delta save tapes and optionally, a DSIM dataset. This is useful when moving a file from a save tape with one blocksize to a database with another, or when using a file from a save tape in different test environments.

If the file to be unloaded from the save tape is ciphered, the CODE parameter must be specified as usual.

User exit 9 can be used to select records for a particular client of a multiclient file. For more information, see the section ADAULD User Exit 9.

For more information, see the section Save Tape Input Processing.

Optional Parameters

CODE: Cipher Code

If the file to be unloaded is ciphered, CODE *must* supply the appropriate cipher code.

DDISN: Create DD/ISN Output File of Unloaded ISNs

Specifying the DDISN parameter instructs ADAULD to write the list of unloaded ISNs to the sequential output file DD/ISN. DD/ISN is structured so that it can be used as input to ADALOD UPDATE for the purpose of deleting the unloaded records.

If the DDISN keyword is specified but the DD/ISN file is missing in the JCL, ADAULD terminates with error-081.

DSIMDEV: DSIM Device Type

The DSIMDEV parameter specifies the device type of the DSIM dataset. This parameter is required only if the DSIM device type is different from that specified by the ADARUN DEVICE parameter (which is the default).

MODE=SHORT: Exclude Descriptor Information

This parameter indicates whether the descriptor information used to build the normal index and upper index are to be included in the output.

If MODE=SHORT is specified, no descriptor information will be unloaded, and all descriptor information is stripped from the field definition table (FDT) when it is written to the output dataset.

If the output is to be used as direct input to the ADALOD utility, the file will have no descriptors.

NOUSERABEND: Termination without ABEND

When an error is encountered while the function is running, the utility 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.

NUMOUT: Number of Output Files

NUMOUT specifies the number of output files to be produced. If the number is greater than one, user exit 9 must be used to control DDOUT1 and 2 or OUT1 and 2 output file selection. For additional information, see the *Adabas DBA Tasks* documentation. Permitted values are 1 (default) and 2.

NUMREC: Number of Records to Be Unloaded

NUMREC limits the number of records to be unloaded. No limit will be in effect if the parameter is omitted.

PATTERN: Merge Input Pattern

The PATTERN parameter can be specified if the save tape from which the file is to be unloaded is supplied as a full save tape plus 1 to 8 delta save tapes. The parameter serves as a cross-check between the intended save input and the actual save datasets specified by the job control.

PATTERN specifies the count and type of input datasets to the merge operation as single letters "F" (full save tape) or "D" (delta save tape). Only one "F" can be specified in the pattern field, but up to eight "D"s can be entered—one for each existing delta save tape to be merged.

The pattern starts with an "F" if a full save input is specified, and continues (or starts) with a "D" for every delta save input. For example, PATTERN=FDD specifies that the input comprises one full save dataset and two delta save datasets. No special indicator is given for a save dataset being online.

The pattern string must exactly match the input save datasets specified in the job control.

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.

Examples

Example 1:

```
ADAULD FILE=6, SAVETAPE
```

File 6 is to be unloaded from a full save tape. The records are to be unloaded in the sequence in which they are physically positioned on the tape.

Example 2:

```
ADAULD FILE=6, SAVETAPE, MODE=SHORT
```

File 6 is to be unloaded in physical sequence from a full save tape. The entries used to create the normal index and upper index are not to be unloaded. All descriptor information is removed from the field definition table (FDT) in the output.

Example 3:

```
ADAULD FILE=6, SAVETAPE, PATTERN=FDD
```

File 6 is to be unloaded in physical sequence from a full save tape and two delta save tapes. The unloaded file will have the same contents as it had at the time of the second delta save operation.

Save Tape Input Processing

ADAULD is used to unload an Adabas file from a full save tape together with 1 to 8 delta save tapes and optionally, a DSIM dataset:

- the full save is expected as a DD/FULL sequential input file;
- the delta saves are expected as DD/DEL1-DEL8 sequential input files; and
- the DSIM dataset as DD/DSIMR1. The DSIM dataset is required if any delta save tape specified is an online tape that is not accompanied by its unloaded DSIM dataset.

Save tapes created using any version of the Delta Save Facility can be used.

The ADARUN DBID specified for the ADAULD run must match the DBID found on the save tapes. The ADARUN parameter DSF=YES must be specified for the ADAULD run.

If the file has hyperdescriptors defined, the corresponding hyperexits must be specified in the ADARUN parameters for ADAULD. If the hyperexit routines are no longer available, the file must be unloaded with MODE=SHORT specified. See the *Adabas DBA Tasks* documentation for more information about hyperexits.

The ADAULD utility protocol on DD/DRUCK displays a short header indicating the kind of save tapes encountered, when each was created, the version of ADASAV used to create it, the database ID found on the save tape, and possibly the delta save ID of the save tapes when merged:

```

A D A U L D  v7.1   SM1   DBID = 00200   STARTED   YYYY-MM-DD  13:33:25

PARAMETERS:
-----

ADAULD UNLOAD FILE=3; SAVETAPE

*****
*
* UNLOAD FROM          ONLINE DATABASE SAVE      *
*
* CREATED AT          1997-04-23                 *
*
* BY ADASAV VERSION   v71                       *
*
* DBID                200                       *
*
* DSID                1 / 0 / 1997-04-23  13:32:19 *
*
* PLOG SESSION NR     17                        *
*
* SYN1/4 BLOCK NR     137                      *
*
*****

```

ADAULD first reads the file control block (FCB) and file definition table (FDT) from the save tape. Then:

- for full save tapes and merged full save tapes, ADAULD scans the tape to find the file's Data Storage RABNs, extracts the Data Storage records, and for each Data Storage record, generates the descriptor values according to the FDT.
- for full save tapes plus delta save tapes and optional DSIM dataset, ADAULD calls the Delta Save merge facility to receive the latest version of each affected block. ADARUN DSF=YES must be specified so that the ADADSFS module containing the merge routines is loaded.

DSIM Dataset or the Protection Log

In order to unload a full save tape and one or more delta save tapes, one of which is online, you *must* supply the DSIM dataset associated with the remaining online save. If the DSIM dataset no longer exists, it is necessary to rebuild it from the protection log (PLOG) using ADARES COPY.

If an online full save tape *without* delta save tapes is unloaded, you have the option of providing either the associated DSIM dataset or the protection log (PLOG) as of the time of the online save. If the full save is not the most recent online save, it is necessary to rebuild the DSIM dataset from the PLOG using ADARES COPY. In this case, it is preferable to supply the PLOG directly to the unload. See the document *ADAULD* in the *Adabas Utilities* documentation for more information.

You can ensure that you have the DSIM datasets you need by unloading the original DSIM dataset to a sequential file using the ADASAV MERGE function. The unloaded DSIM dataset can be used wherever the original DSIM would be used. In addition, several unloaded DSIM datasets can be used in a merge, restore, or unload operation, whereas only one original DSIM can be supplied to any of these operations.

ADAULD Output Processing

ADAULD unloads the records in the specified sequence. The unloaded records are written to one of two sequential datasets: DD/OUT1 and DD/OUT2. Writing to these output datasets is controlled by user exit 9.

The records output are identical in format to the output produced by the ADACMP utility unless the MODE=SHORT option is used, in which case the descriptor entries required for the normal index and upper index are omitted and the descriptor information is removed from the Adabas FDT. The ISN of the record immediately precedes the compressed data record, and is provided as a four-byte binary number.

Specifying the DDISN parameter instructs ADAULD to write the list of unloaded ISNs to a sequential output file DD/ISN. Only one DD/ISN file is created, containing the superset of ISNs written to either or both DD/OUT1 and DD/OUT2. ISNs that are rejected by userexit 9 are not written to DD/ISN.

DD/ISN is structured so that it can be used as input to ADALOD UPDATE for the purpose of deleting the unloaded records.

The number of ISNs written to DD/ISN is displayed in the ADAULD statistics on the DD/DRUCK utility protocol:

```

A D A U L D   S T A T I S T I C S
-----

NUMBER OF OUTPUT DATA SETS      = 1
NUMBER OF REQUESTED RECORDS     = 16777215
STARTISN                         = 0
OPTIONS                          = DVT
UNLOAD SEQUENCE                  = PHYS SEQ

NUMBER OF RECORDS READ           =          1000
NUMBER OF RECORDS WRITTEN        =          1000
RECORDS WRITTEN TO DDOUT1        =          1000
RECORDS WRITTEN TO DDOUT2        =           0
RECORDS REJECTED BY USEREXIT-9  =           0
NUMBER OF ISNS WRITTEN TO DDISN =          1000
    
```

The number of ISNs written to DD/ISN should always be the number of records read minus the number of records rejected by user exit 9.

ADAULD User Exit 9

User exit 9 is called (when present) for each record selected before writing the record to the output dataset. The user exit is supplied with the record address, and returns an action code as follows:

1	write record to DD/OUT1
2	write record to DD/OUT2
3	write record to DD/OUT1 and DD/OUT2
I	ignore this record

The above datasets must have the same blocksize. See the *Adabas DBA Tasks* documentation for more information about user exits.