

ADASAV MERGE

This chapter covers the following topics:

- Consolidating Delta Saves or Updating Full Saves
 - Merging Online Saves and DSIM Datasets
 - Unloading DSIM Datasets
 - Using Unloaded DSIM Datasets as "Online" MERGE Input
 - Concatenated Datasets in MERGE Input
 - Syntax
 - Optional Parameters
 - Examples
-

Consolidating Delta Saves or Updating Full Saves

The ADASAV MERGE function takes as input zero or one full save dataset and zero to eight delta save datasets:

DD/DEL1...DD/DELn	up to eight consecutive delta save tapes covering a contiguous time span with no gaps or overlaps. Do not omit a delta save tape.
DD/FULL	the optional full save input tape.

The input datasets may originate from SAVE database, SAVE DELTA, or MERGE functions. They must cover a continuous sequence of full and delta save operations.

ADASAV MERGE combines the input datasets into a single dataset: either a consolidated delta save tape or an updated full save tape:

DD/SAVE1	the new full or delta save output tape.
----------	---

The output dataset may subsequently be used as input for MERGE, RESTORE, or RESTORE DELTA functions.

If the DRIVES parameter is specified, ADASAV MERGE cuts the *full* save output into as many pieces as specified by DRIVES:

DD/SAVE1...DD/SAVE _n	up to eight consecutive partial save tapes comprising together the new full save output tape.
---------------------------------	---

The output datasets created can be used as input to a RESTORE function with the same value of DRIVES.

If only delta save input datasets are specified, the MERGE function creates a new delta save dataset that contains every database block from each input save dataset. If a block is present on more than one input save dataset, its latest occurrence is taken. The result is a "consolidated" delta save dataset that contains all changes since the last full save or since the last delta save not included in the MERGE function.

If a full save dataset is specified as input, the MERGE function creates a new full save dataset that contains the latest image of every used block of the database. This full save dataset is equivalent to the database at the time of the last SAVE DELTA execution.

Merging Online Saves and DSIM Datasets

The "online" variation of the MERGE function takes as input a full or delta save dataset that originates from an online save operation *and* the accompanying DSIM dataset that was specified for the online save operation:

DD/DSIMR1	interim dataset from an online SAVE operation.
-----------	--

An online full or delta save dataset with its corresponding DSIM dataset can also be specified in a merge with "offline" save datasets. The output save dataset from the "online" merge operation is always like an "offline" save dataset and may be supplied together with another online save dataset for a subsequent MERGE operation.

The DSIM dataset is automatically built while ADARES PLCOPY copies the dual or multiple protection logs (PLOGs). Building the DSIM dataset is completed when the PLCOPY function encounters the online save operation's SYN2 checkpoint on the PLOG.

If the MERGE function is performed directly after the online save operation, the DSIM dataset will normally not be completely built when the MERGE function starts. The DSIM dataset will not be ready for merge before the next PLCOPY operation has been performed. The DSIMWAIT parameter tells the MERGE function how long to wait for the DSIM dataset to become ready for merge.

Upon completion of the MERGE function, the DSIM dataset is reset for use in another online save/merge cycle.

Unloading DSIM Datasets

The "unload" variation of the MERGE function transfers the contents of the DSIM dataset

DD/DSIMR1	interim dataset from the last online SAVE operation.
-----------	--

-to a sequential delta save dataset

DD/SAVE1	delta save dataset.
----------	---------------------

-that can be used in MERGE or RESTORE DELTA operations later on as needed.

The "unload" frees the DSIM dataset for the next online SAVE operation.

Unlike the "merge" of the DSIM dataset, the "unload" does not process the associated online full or delta save dataset again immediately after it is created. Instead, more than one online save dataset can be processed later on in a single MERGE or RESTORE DELTA operation by sequencing for each the delta save dataset created from its associated DSIM dataset.

The DSIM dataset is unloaded by running ADASAV MERGE with only the DSIM dataset as input. The PATTERN parameter must be omitted or blank, and the DRIVES parameter must not be set to a value greater than one. The DSIMWAIT parameter functions normally.

Whereas the DSIM dataset is direct access and contains changed blocks from the database in no particular order, the output dataset is sequential and contains the blocks in ascending RABN sequence without duplicates, like any delta save dataset.

An unloaded DSIM dataset is identified by a flag in the save dataset header. The header specifies

- the delta save ID (DSID) of the associated online save; and
- the positions of the online save's SYN1 and SYN2 checkpoints on the PLOG.

The fields mentioned in the header record are filled from the corresponding fields in the DSIM header block.

If the DSIM dataset was (re)built by an ADARES COPY execution, the DSID is not known and will be zero. The associated online save is still correctly identified by the positions of its SYN1 and SYN2 checkpoints on the PLOG.

When a DSIM dataset is unloaded, the MERGE function still generates a SYN0 checkpoint but with an indication that the created save dataset is an unloaded DSIM dataset.

Using Unloaded DSIM Datasets as "Online" MERGE Input

When an unloaded DSIM dataset is supplied as input for a MERGE operation, it must be specified

- as one of the DD/DELn input datasets. More than one unloaded DSIM dataset may be supplied for a single MERGE.
- directly after its associated online delta save input; that is, with the next higher DD/DELn index number. For example, if an online delta save input is supplied as DD/DEL2, its associated unloaded DSIM dataset must be given as DD/DEL3. If an unloaded DSIM dataset is associated with an online full save, it must be specified as DD/DEL1.

These rules for specifying unloaded DSIM datasets ensure the correct output: each block is taken from the last input in sequence that contains a block image with the same RABN.

The MERGE function checks whether the online save fits the contents of the DSIM dataset. If so, the online save is associated with the DSIM dataset. If not, the next input delta save in logical sequence must be a matching unloaded DSIM dataset.

Concatenated Datasets in MERGE Input

The full save input to the ADASAV MERGE function must be supplied as a single entity. Full saves that are spread over several datasets (by means of the DRIVES parameter in ADASAV), must be concatenated as follows:

- For BS2000/OSD systems, by using the DDFULL01, DDFULL02, ... link names for the second, third, ... datasets;
- For z/OS, and z/VM systems, by using operating system concatenation capabilities;
- For VSE/ESA and z/VSE systems, by using the FULL01, FULL02, ... symbolic names for the second, third, ... datasets.

For a MERGE operation, a delta save dataset may not be concatenated to another full or delta save dataset. If the MERGE operation detects a concatenated dataset at the end of any other dataset in the logical sequence, it terminates with an error message.

Note:

In cases where the MERGE function can be completed without reading all input save datasets to their end, wrongly concatenated input save datasets may still go undetected and be ignored.

See the JCL examples for the ADASAV MERGE function in section JCL/JCS Requirements and Examples. for more information on dataset concatenation on BS2000 and VSE/ESA systems, see the *Adabas Utilities* documentation.

Syntax

```
ADASAV MERGE [ DRIVES = count | 1 ]  
              [ DSIMDEV = device-type | ADARUN-device ]  
              [ DSIMWAIT = seconds | 0 ]  
              [ NOUSERABEND ]  
              [ PATTERN = merge-pattern ]  
              [ PERDRIVE = disks-per-tape ]  
              [ TEST ]  
              [ TWOCOPIES ]
```

Optional Parameters

DRIVES: Number of Output Datasets

DRIVES can only be specified if the output save dataset is a full save.

The DRIVES parameter specifies the number of output datasets to be created, all of which together build the full save output. The MERGE function cuts the output full save into as many pieces as specified by DRIVES.

Note:

The full save input must always be specified using a single DD/link name (DD/FULL). Concatenation can be used, if necessary. See Concatenated Datasets in MERGE Input.

The output datasets created can be used as input to a RESTORE function with the same value of DRIVES.

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).

DSIMWAIT: DSIM Dataset Wait Time in Seconds

The DSIMWAIT parameter specifies how long the MERGE function should wait for the DSIM dataset to become ready for merge.

It can be used if a merge step is to be executed directly after an online save operation but must wait for the ADARES PLCOPY function to copy the nucleus' dual or multiple protection log.

The maximum time to wait is specified in seconds. If the DSIM dataset does not become ready for merge during this interval, and one of the input save datasets is an online save dataset, the MERGE function fails.

If DSIMWAIT is not specified, the MERGE function will not wait for the DSIM dataset to become ready for merge, but instead will begin operation immediately or fail if the DSIM dataset is not ready for merge (the default).

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.

PATTERN: Merge Input Pattern

PATTERN is an optional parameter that provides a cross-check between the intended merge input and the actual 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 (additional delta save tapes can be merged in subsequent merge operations with the result of this merge operation).

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.

PERDRIVE: Disk Drives Per Tape Drive

PERDRIVE specifies the number of disk drives to be assigned to a single output tape drive. For example, if the database is contained on seven disk drives and three tape drives are available for MERGE processing, PERDRIVE=3,2,2 would cause the first three disk drives to be written to tape drive 1, the next two disk drives to be written to tape drive 2, and the next two disk drives to be written to tape drive 3. The drive sequence corresponds to the DD/SAVE_n or DD/DUAL_n job control specifications, as described in the section JCL/JCS Requirements and Examples.

The total number of drives specified by PERDRIVE must equal the sum of all Associator (DD/ASSOR_n) and Data Storage (DD/DATAR_n) disks; if both Associator and Data Storage are on a single disk, this counts as two separate disks. If the DRIVES parameter is used and the PERDRIVE parameter is omitted, ADASAV determines the most efficient utilization of the tape drives.

TEST: Test Syntax

The TEST parameter tests the operation syntax without actually performing the merge operation. Note that the validity of values *cannot* be tested; only the syntax of the specified parameters can be tested.

TWOCOPIES: Create Two Copies of Output

TWOCOPIES creates two physical copies of the ADASAV output.

Examples

Example 1:

```
ADASAV MERGE
```

```
DSIMDEV=8381, DSIMWAIT=300, PATTERN=DD
```

Two delta save tapes, one created during an online delta save, are to be merged into a single delta save tape. The merge operation will wait up to five minutes (300 seconds) for the DSIM dataset to become available. The DSIM dataset's device type is 8381.

Example 2:

```
ADASAV MERGE PATTERN=FD
```

```
ADASAV          DRIVES=4
```

```
ADASAV          TWOCOPIES
```

One full save and one delta save dataset are to be merged. The full save output is to be distributed over four datasets, making them suitable for a restore using four drives. Two copies of each output dataset are to be made.

Example 3:

```
ADASAV MERGE PATTERN=F
```

Only one (online) full save and its corresponding DSIM dataset are to be merged. The output is equivalent to an offline full save taken at the end of the online save operation.

Example 4:

```
ADASAV MERGE PATTERN='FDDDDDDDD'
```

A full save and eight delta saves (the maximum number allowed) are to be merged.