

Introducing Adabas Delta Save Facility

Software AG's Adabas Delta Save Facility (DSF) offers significant enhancements to ADASAV utility processing by backing up and restoring only the changed (delta) portions of Adabas databases.

The Delta Save Facility

- includes logic in the Adabas nucleus for recording the RABNs of all changed database blocks. Adabas utilities that change the database contain logic to record which portions of the database have been affected.
- creates and merges interim "delta" save tapes while the database remains online, if required.
- consolidates delta save tapes with the most recent database save tape to create an up-to-date full save tape.
- restores the database from the most recent full save tape and all subsequent delta save tapes.

While maintaining the availability of Adabas SAVE and RESTORE functions, the high performance of Adabas, and the ability of Adabas to survive system failures, DSF increases database availability by

- reducing the volume of save output produced; and
- shortening the duration of save operations.

Because DSF allows more frequent save operations, database recovery time is reduced.

DSF is intended for Adabas sites with one or more large, heavily updated databases that need to be available most of the time. Especially for sites where the volume of data changed on a day-to-day basis is considerably smaller than the total database volume, DSF provides for

- more frequent saves without interrupting database availability;
- enhanced "24 x 7 x 52" operation;
- full offline saving in parallel with the active database; and
- shorter REGENERATE duration during recovery.

This document tells you how DSF operates and how it manages Adabas database save operations more effectively.

This chapter covers the following topics:

- How the Adabas Delta Save Facility Works
 - Delta Save Functions
-

How the Adabas Delta Save Facility Works

During normal Adabas database operation, the Adabas Delta Save Facility keeps track of changed Associator and Data Storage blocks. The RABNs of blocks that are altered are recorded in a portion of the Associator called the Delta Save logging (DLOG) area, allocated for the purpose. During a delta save operation, images of these changed blocks are copied to create a delta save tape.

Delta save and normal Adabas nucleus operations can occur in parallel. Any changes made by the nucleus during the delta save operation are later retrieved from the Adabas protection log (PLOG). In most cases, the resulting delta save output is written to tape for later use.

Merging the most recent ADASAV SAVE result (SAVE database) with all following delta save output results in a new full database save that reflects the database status following the last delta save operation. Alternatively, the latest SAVE database and all subsequent delta save results can be supplied as input to a database restore (ADASAV RESTORE DELTA) operation.

The specific stages in this process are as follows:

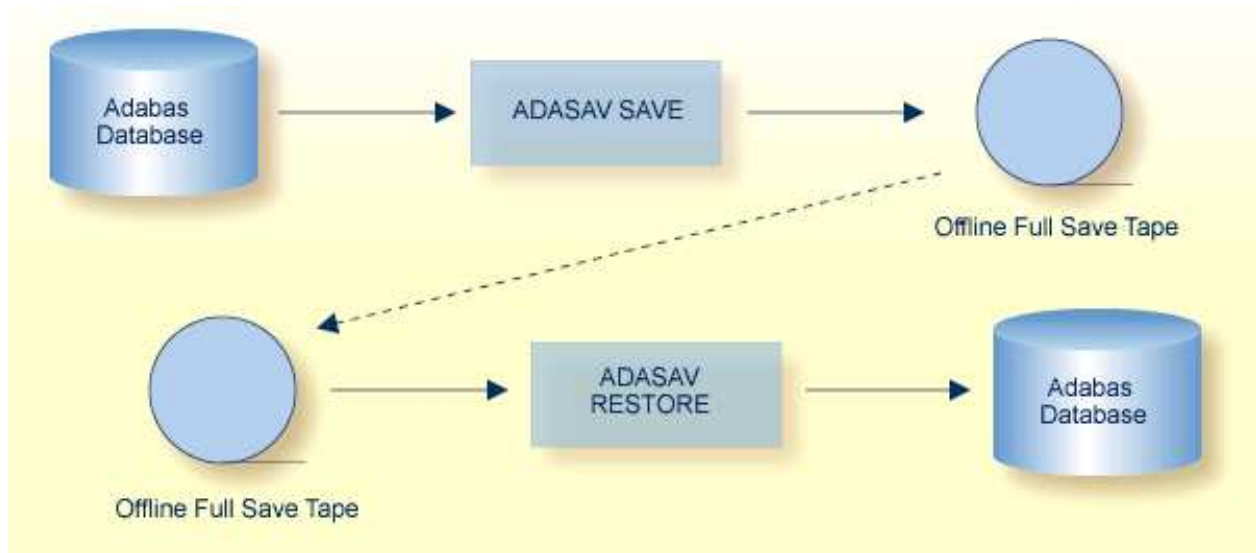
1. Enable Delta Save operation by performing a normal full save (ADASAV SAVE database) that contains the entire database.
2. At some point following step 1, perform a delta save (ADASAV SAVE DELTA) operation, creating output that contains only the blocks that were changed since the last full or delta save operation was performed; this can be done multiple times.
3. Optionally, merge several previously created delta save outputs to produce a "consolidated" delta save tape comprising all changes since the last full save tape creation.
4. Merge the previously created full and delta save output to produce an up-to-date full save tape. Note that the old full or delta save inputs can also be the result of previous merge operations.
5. If the need arises, restore the database by supplying either a fully merged full save tape, or the most recent full save tape and all subsequent delta save tapes.

Conventional Database Save Operation

Without the Delta Save Facility, a conventional database save operation can be performed with the Adabas nucleus inactive (offline) or active (online).

Offline Save

An offline save is performed with the Adabas nucleus inactive. No concurrent access or update of the database is allowed so that the resultant save tape contains all database Associator and Data Storage blocks consistent with the time of the save operation.

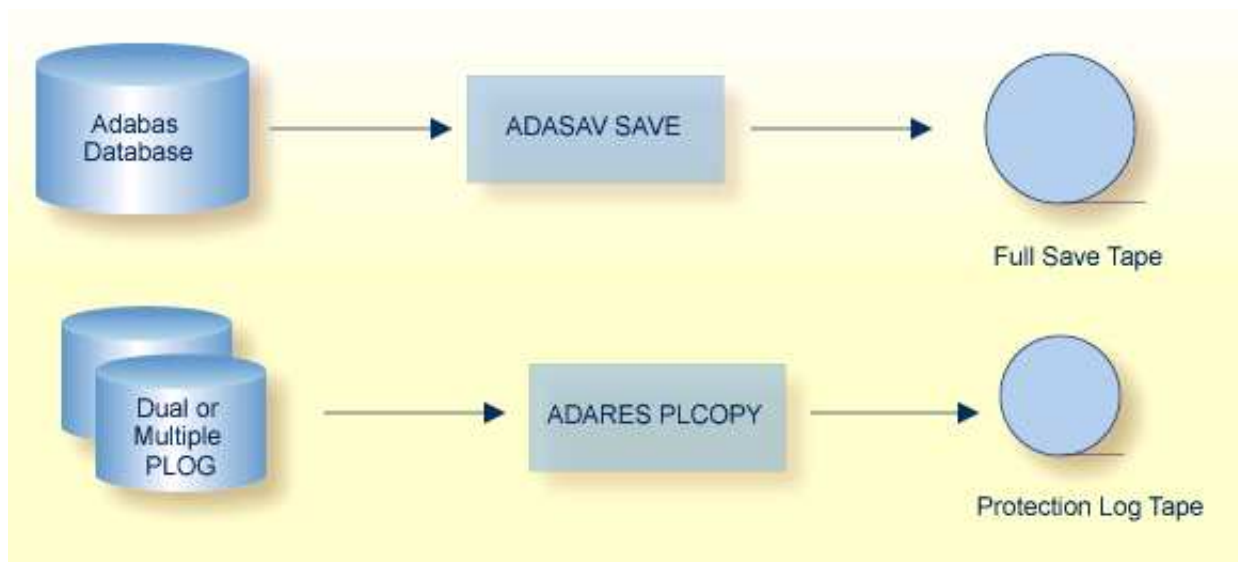


The graphic shows the general sequence of a conventional Adabas database offline save operation without Delta Save. First, the ADASAV utility’s SAVE function is used to create a full database save (SAVE database), usually on tape. If the need arises, this full database save can be used for restoring either the entire database or specific files to their state at the time of the save operation (using the ADASAV RESTORE function).


Online Save

An online save is performed while the Adabas nucleus is active and the database is available for normal use. During the online save, users may access and update records in the database. Concurrent utility updates are not allowed.

At the end of the save, ET synchronization ensures that all partial transactions are completed before new transactions are started so that the database is in a consistent state. Since the database can be modified during the save operation, the save tape alone may not be consistent with the database. Therefore, the images of any blocks in the database that were changed while the save operation is in progress are recorded in the protection log (PLOG), as shown in Online Adabas Save Sequence.



If required, the entire database or specific files can be restored to the state at the end of the online save operation by using the full database save tape together with the associated protection log as input to the ADASAV RESTONL function, as shown below.



Save Operation with Delta Save

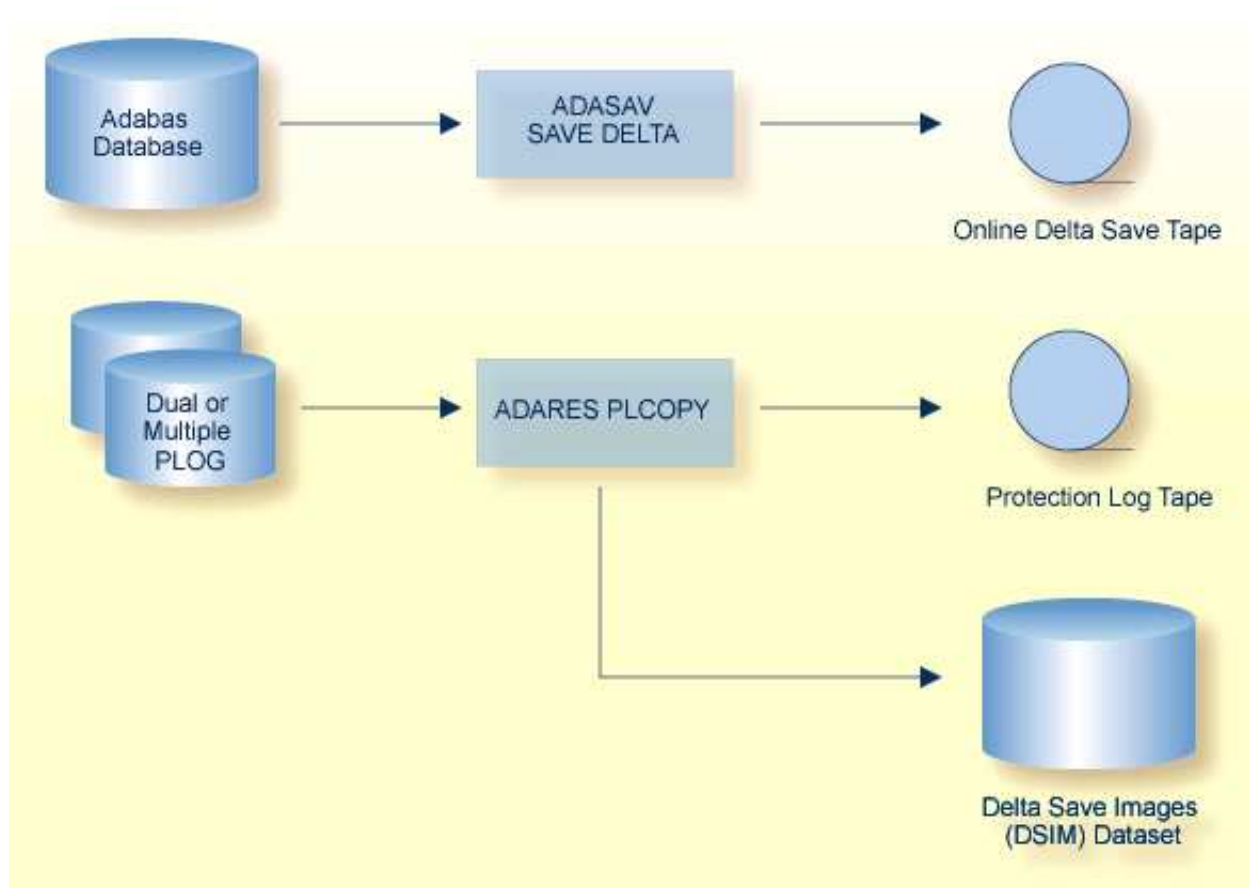
The Delta Save Facility provides full save and delta save operations. The full save saves the whole database, similar to a conventional database save, and is invoked by the ADASAV SAVE function. The delta save operation saves only blocks changed since the last full or delta save operation, and is invoked by the ADASAV SAVE DELTA function. For full save and delta save operations, the database administration blocks, FCBs, FDTs, and DSST are always saved.

Although the procedure for creating a delta save tape is much the same as for a full database save tape, the actual time required is much less because only the changed areas of the database are saved. Secondly, by performing less frequent full saves but more frequent delta save operations, the overall time needed to run the save operations is normally considerably less than performing the same number of conventional database save operations.

Further processing of full save and delta save tapes is provided by the ADASAV MERGE function. The merge function can be performed offline and therefore does not affect database processing.

When an ADASAV SAVE DELTA operation is performed, the logged Delta Save information is read and compiled into a save tape containing only those changes made since the last full or delta save operation, as shown in the following Offline Delta Save Tape and Online Delta Save Tape diagrams.





Online Delta Save Tape shows that the online delta save-like a full save-results in protection log (PLOG) output for any activities performed while the delta save was in process.

In addition, however, the PLOG copy (ADARES PLCOPY) operation after the delta save produces a delta save images (DSIM) data set. This new data set holds copies of all Associator and Data Storage blocks changed while the online delta save operation was in progress, and provides them for a subsequent merge operation.

The save tapes created by the ADASAV SAVE FILE or ADASAV SAVE INCREMENTAL functions cannot be used as delta save tapes and therefore cannot be used in merging operations.

Tracking Database Changes

When the Delta Save Facility is used, the nucleus logs the RABNs of changed blocks in the Delta Save logging area (DLOG). The DLOG area is allocated within the Associator by the Adabas basic online services facility. The delta save operation retrieves RABNs of blocks to be saved from the DLOG area. When a full save or delta save operation completes successfully, the DLOG area is reset.

The DLOG area is continuously updated as more blocks are changed by normal nucleus operation. In addition, Associator and Data Storage changes made by utilities are recorded, although not in the DLOG area but in the FCBs of the affected files.

When a utility modifies a file, information is recorded in the related FCBs that the file's Data Storage, address converter, and/or index was changed. The changed RABNs themselves are not logged. Thus, for example, when the ADAORD REORFDATA function reorders the Data Storage of a file, it sets marks in the FCB of the file to indicate that the entire Data Storage and address converter of the file are to be saved

in the next delta save operation.

If the utility function is ADALOD UPDATE, the RABN ranges of changed Data Storage and address converter blocks are stored in a file-specific extension of the DLOG area allocated by ADALOD. Only the blocks in the extension are marked for saving during the next delta save operation. (ADALOD UPDATE always changes, and marks as changed, the entire index of the updated file.) If the area reserved for logging RABNs changed by ADALOD UPDATE overflows, the logging information is compressed and more blocks than were actually modified will be saved in the next delta save operation.


Merging the SAVE Output

The ADASAV MERGE function combines the delta save results and, where applicable, the original or unloaded DSIM data set from one or more ADASAV SAVE DELTA operations to create a new delta or full save tape. This is illustrated in Merge Operation - Making a Consolidated Delta Save Tape.

When you "unload" the DSIM data set, you transfer its contents to a sequential delta save data set without specifying and processing the online full or delta save data set associated with the DSIM data set. The unloaded DSIM data set can be used for MERGE or RESTORE DELTA operations.

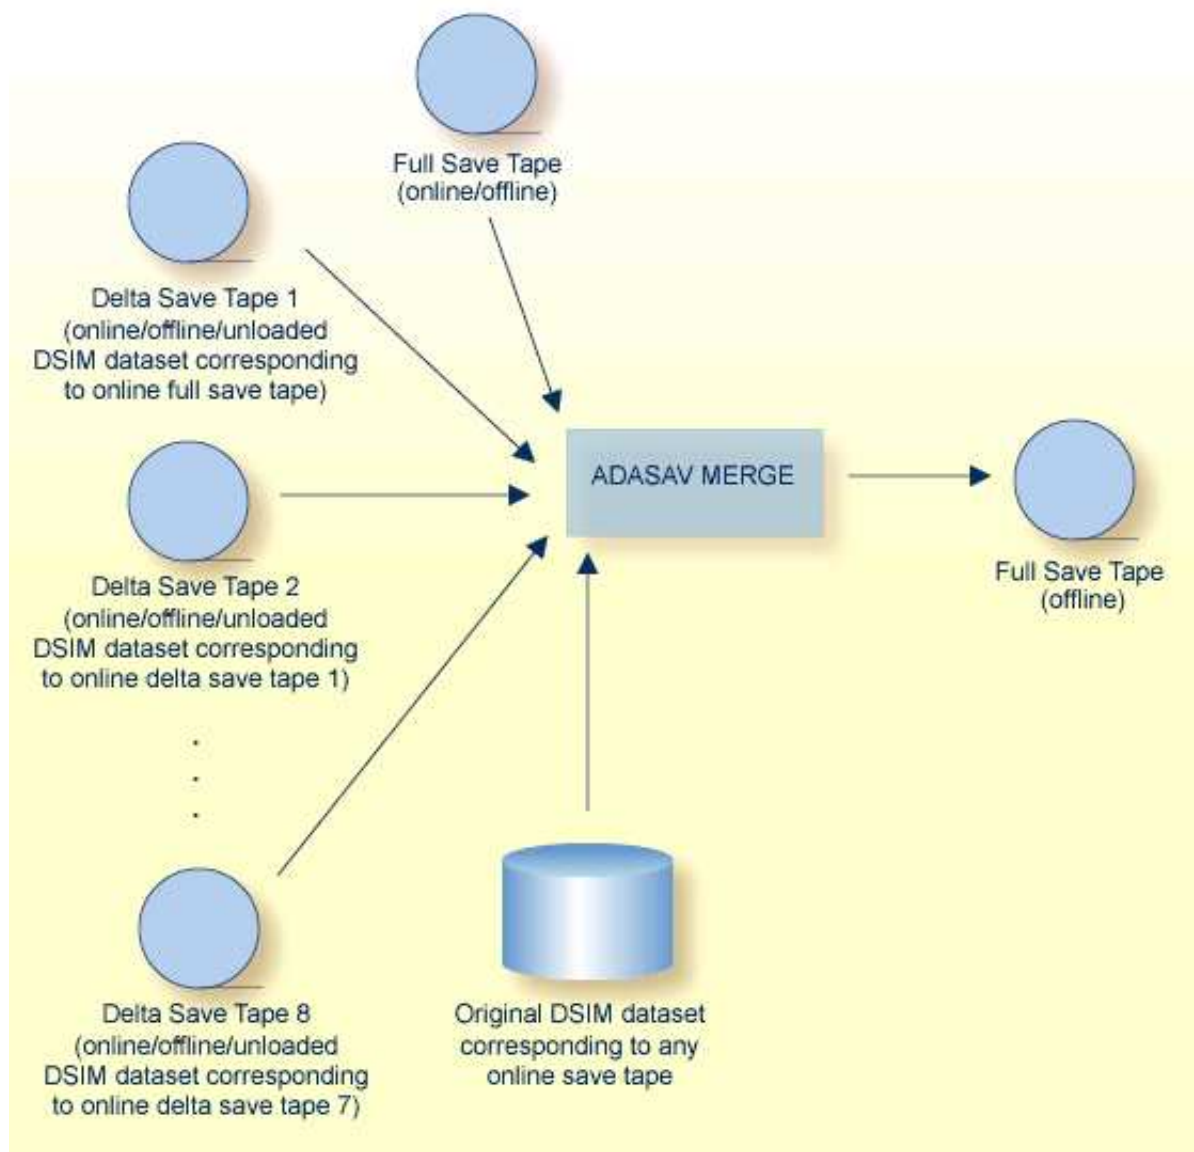
If only delta saves (that is, no full saves) are merged, the output is a single, consolidated delta save tape. Such a save tape comprises all database changes that occurred in the time period covered by the input delta saves.

If the input contains a full save tape, the merge operation results in a single, merged full save output. A merged full save tape contains the whole database as it was at the time when the last input delta save was written.



The original DSIM data set corresponds to any online save input. Any online delta save input to which the original DSIM, if supplied, does not correspond must be followed (with the next higher sequence number) by the corresponding unloaded DSIM.

If the full save input is an online save tape, it must be accompanied either by its original DSIM data set or followed by the delta save data set created when the original DSIM data set was unloaded. This is illustrated in Merge Operation - Making a New Full Save Tape.



The unload frees the DSIM data set for the next online save operation. Unlike a MERGE of the DSIM data set, its "unload" eliminates the need to process the associated online save data set-which may be very large in the case of a full save-again just after it has been created.

The original rather than the unloaded DSIM data set related to any online save operation can be supplied for the merge operation. Following the merge operation, the DSIM data set is reset and can then be used for the next online save operation.

The possible inputs for merging are

- zero or one full save tapes;
- zero to eight consecutive delta save tapes including unloaded DSIM data sets for each online save tape;

- the DSIM data set for any online save unless the unloaded DSIM data set is provided.

It is a requirement for all input save tapes to be from a continuous sequence of save operations on the same database. No save tapes can be omitted; otherwise, the result would be inconsistent. Therefore, a non-continuous sequence of input tapes is rejected by the ADASAV MERGE function.

Each input save tape can be the result of earlier merge operations.

The output tape may be used in subsequent merge operations.

Any online save input to the merge operation must be accompanied by either the original or the unloaded DSIM data set related to that online save operation. Only one original DSIM data set can be specified.

The output of a merge operation is always the same as an offline save tape; it can be specified with a new online delta save tape in a subsequent merge operation.

If more than eight consecutive delta save tapes have been created, you must merge the first eight, and then the result of that merge with the next delta save tape(s), and so on, in chronological order.

Delta Save Identifier (DSID)

If the Delta Save Facility is used, every save tape (full or delta) is assigned a Delta Save identifier (DSID), which is used for consistency checking in merge operations, for example, to prevent invalid merge operations. The DSID comprises the full save number, the delta save number, and a time stamp.

The full save number counts the number of full save operations since the Delta Save Facility was installed. The delta save number counts the number of delta save operations since the last full save. The time stamp records the date and time of the last full or delta save.

Each full save operation increments the full save number by one, resets the delta save number to zero, and assigns a new time stamp. Each delta save operation leaves the current full save number unchanged, increments the delta save number by one and assigns new time stamp.

When save tapes (full or delta) are subsequently input for merging, they must have the same full save number, consecutive (or overlapping) delta save numbers, and ascending time stamps. Overlapping delta save numbers in data sets used as input to the merge are allowed.

During a merge operation, the DSID of the merged save tape (full or delta) is derived from the input DSIDs as follows:

- full save number;
- delta save number range (that is, first delta save number - last delta save number);
- last time stamp.

Example:

Full save input DSID:	4/0/t1	(Full save number 4, delta save number 0, date/time t1)
Delta save input 1 DSID:	4/1/t2	
Delta save input 2 DSID:	4/2/t3	
Merged full save output DSID:	4/0-2/t3	

Invalid Merges

Before a merge operation starts, the Delta Save Facility checks that the DSIDs of the full and delta save tapes are a valid combination. In almost all cases, this check reliably prevents any invalid merge operation.

There are, however, some unusual cases in which the check is insufficient, namely when DSIDs are "reused" after

- removal and reinstallation of the DLOG area, or
- restore of the database to an old status.

Example:

Full save	DSID 1/0/t1	(t1 is a time stamp)
... updates ...		
Delta save	DSID 1/1/t2	
Restore full save, restores	DSID 1/0/t1	
... updates ...		
Delta save	DSID 1/1/t3	
... updates ...		
Delta save	DSID 1/2/t4	

It is now possible to merge full save 1/0/t1, delta save 1/1/t2, and delta save 1/2/t4, even though this is an invalid combination. Such a mistake can be avoided by observing the following rules:

Important Rules	
1.	When the database is restored to an old status, all delta save tapes produced after the time represented by that status must no longer be used and should be discarded.
2.	When the DLOG area is removed and reinstalled, all existing delta save tapes should be discarded.

In the example above, delta save tape 1/1/t2 should have been discarded after the restore.

Protection Log (PLOG) and DSIM Data Set Processing

When the Delta Save Facility (DSF) performs an "online" save operation, special processing is required for the PLOG and the DSIM data set that is not necessary for an offline save operation.

The nucleus should run with a dual or multiple PLOG when performing an online save with DSF. PLOG and DSIM data set processing proceeds as follows when dual or multiple protection logging is in effect:

1. The DSIM data set must be free (not in use) initially.
2. The online save acquires the DSIM data set for use.
3. The online save operation starts (full or delta).
4. The nucleus writes images of changed blocks to the dual or multiple PLOG.

During a single online save operation, several ADARES PLCOPY operations may be performed to copy the dual or multiple PLOG information to the sequential PLOG and to extract images of the changed Associator and Data Storage blocks and write them to the DSIM data set.

ADARES can be invoked documentionly or by user exit 2 (dual logging) or user exit 12 (multiple logging) when it is programmed as recommended (see the *Adabas DBA Tasks* documentation for information about the user exits). The user exit invokes ADARES when a PLOG is switched, which happens either when one PLOG data set runs full or when the FEOFPL function is issued from ADADBS OPERCOM, AOS, or the operator console.

Creating an online save tape and building a corresponding DSIM data set are related but separate operations. The DSIM data set can be built step by step during the several PLCOPY operations that may be performed. However, when the online save operation ends, a PLCOPY operation for completing the DSIM data set is still outstanding.

5. The online save operation terminates. At this point there is information on the PLOG that must be copied to the DSIM data set.
6. ADARES PLCOPY is run again to complete the DSIM data set.
7.
 1. A subsequent merge operation combines the online save tape and the DSIM data set.

For the merge operation:

- zero or one (1) full save data sets and zero to eight (8) delta save data sets can be merged.
- both offline and online save data sets can be supplied to a single merge operation.
- online full or delta save data sets must be supplied in sequence with the appropriate DSIM data set. Unloaded DSIM data sets count as delta save data sets. Any online save data set must be accompanied by its associated original or unloaded DSIM data set.

The result of merging an online save data set and its DSIM data set is an "offline" save tape.


The merge function waits for completion of the DSIM data set if the DSIMWAIT parameter is specified.

2. Alternatively, you can "unload" the DSIM data set without immediately merging it with the online save, deferring this merge until a later time (see the DSIMWAIT parameter).
8. On completion of the merge or unload operation, the DSIM data set is freed (no longer in use).

Example

Delta Save - Save/ Merge Sequence (four steps) shows several Delta Save save and merge operations.

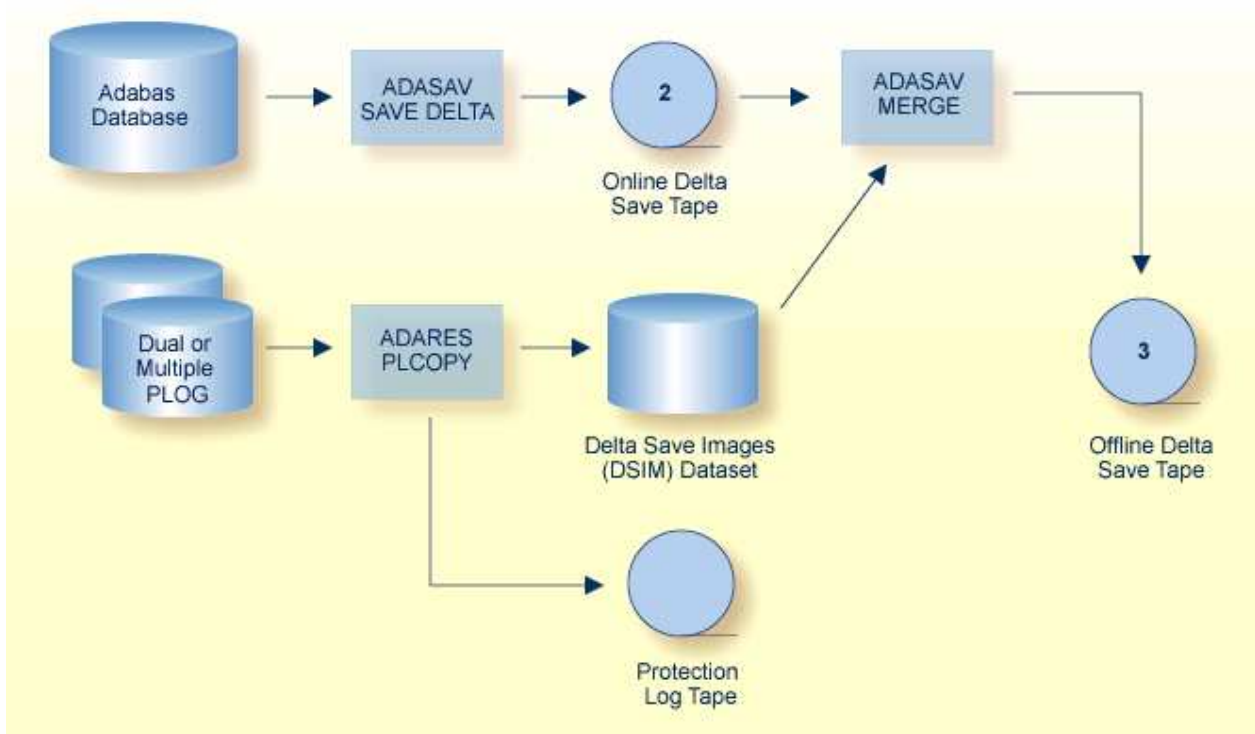
Step 1: Online Full Save



- full save operation is performed online;
- ADASAV SAVE creates an online full save tape;
- Assuming the save is the first after installation of the DLOG area, the DSID of the save tape is 1/0/t1 (that is, full save number 1, delta save number 0, and date/time of the save, t1);
- During the online save, the nucleus writes images of changed database blocks to the dual or multiple PLOG;
- ADARES PLCOPY copies the dual or multiple PLOG to a sequential PLOG; it also extracts block images from the PLOG and writes them to the DSIM data set;
- When the save operation is complete and the DSIM data set has been built, the DSIM data set is unloaded to a sequential file for later use. Its DSID is 1/0/t1;

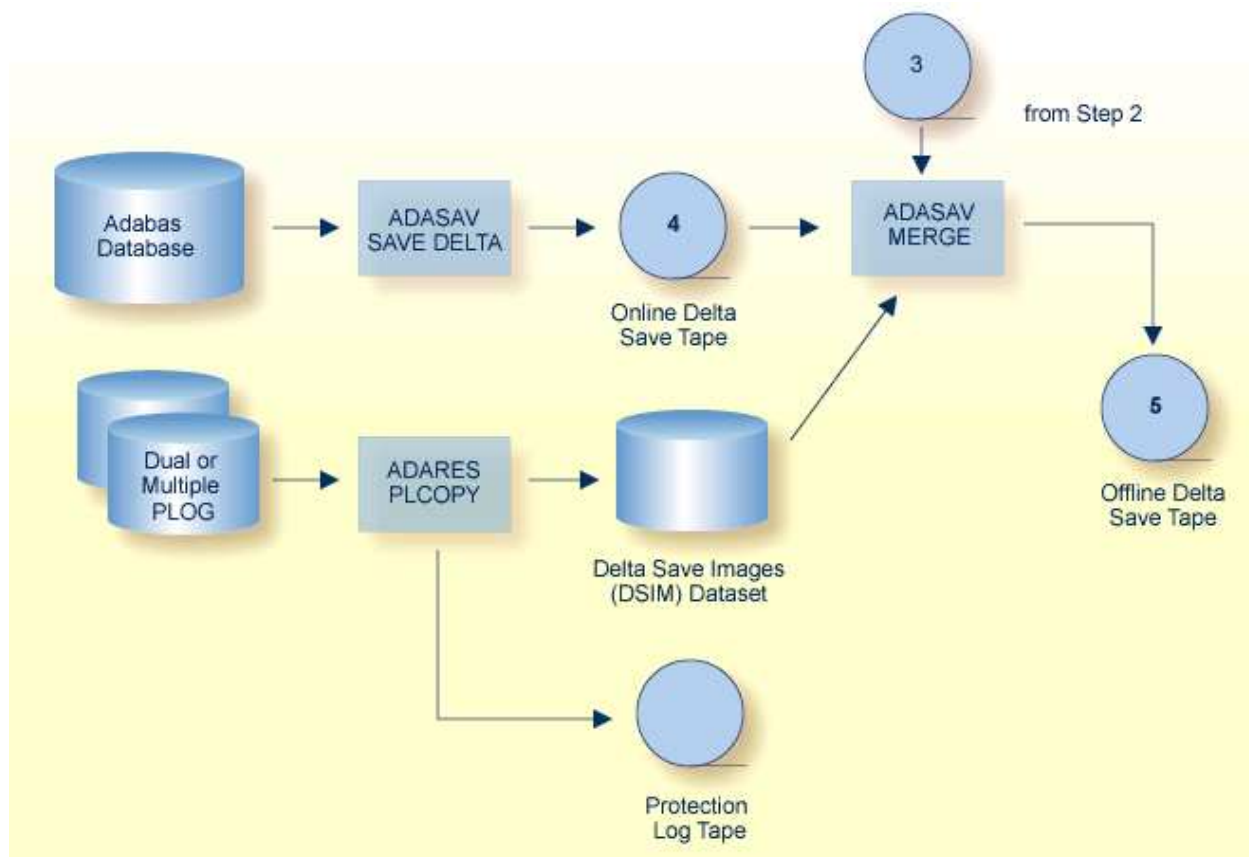
- It is possible to merge the online full save with the DSIM data set to produce an offline full save. In this case, the DSID of the merged save tape would still be 1/0/t1.
- The DSIM data set is now free for the next online save.

Step 2: Online Delta Save



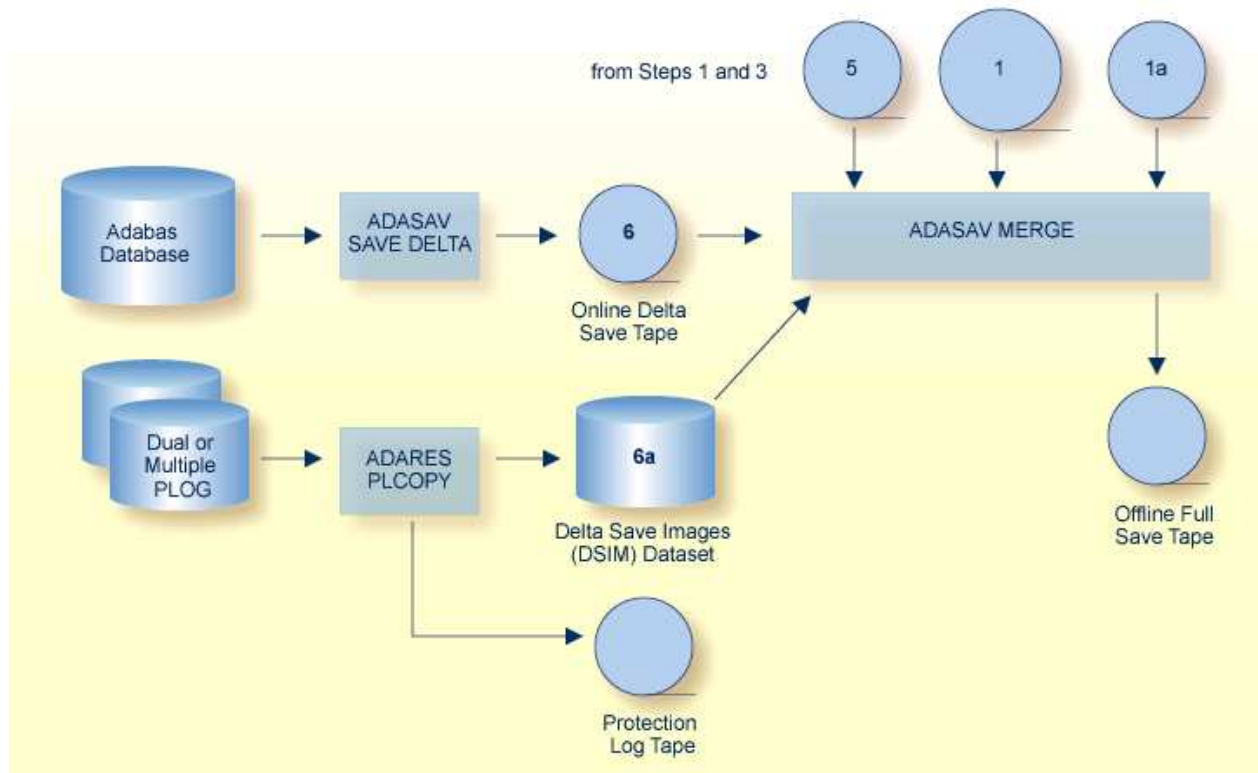
- A delta save operation is performed online;
- ADASAV SAVE DELTA creates an online delta save tape;
- The DSID of the delta save tape is 1/1/t2 (where t2 denotes the date/time of the save operation);
- In the same way as in step 1, the images of database blocks changed during the online save operation are copied to the DSIM data set;
- When the save operation has completed and the DSIM data set has been built, either the DSIM data set can be unloaded to a sequential file for later use or ADASAV MERGE can be used to combine the online delta save and the block images on the DSIM data set to an offline delta save;
- If ADASAV MERGE is used, the DSID of the merged save tape is still 1/1/t2.

Step 3: Online Delta Save, Merge Consolidated Delta Save



- A delta save operation is performed online;
- ADASAV SAVE DELTA creates an online delta save tape,
- The DSID of the delta save tape is 1/2/t3;
- ADASAV MERGE combines the merged delta save tape from step 2, the online delta save tape just created, and its corresponding DSIM data set to a new, consolidated offline delta save tape;
- The DSID of the merged delta save tape is 1/1-2/t3;
- The merged delta save contains all blocks changed since the full save in step 1.

Step 4: Online Delta Save, Merge Full Save



- A delta save operation is performed online;
- ADASAV SAVE DELTA creates an online delta save tape;
- The DSID of the delta save tape is 1/3/t4;
- ADASAV MERGE combines the online full save tape (1) and its unloaded DSIM data set (1a) from step 1, the consolidated delta save tape (5) from step 3, the online delta save tape just created (6), and its corresponding original DSIM data set (6a) to a new, merged full save tape;
- The DSID of the merged full save tape is 1/0-3/t4;
- The merged full save tape contains the database as it was at the end of the last online save;
- This full save tape may be supplied as input for a ADASAV RESTORE database function;
- The merged full save tape can also be supplied as input for ADASAV RESTORE FILE or RESTORE FMOVE functions.

Sample Schedule for Save and Merge Operations

1. Start a cycle with a full save.
2. Perform a delta save operation every day.
3. Maintain a consolidated delta save tape that contains all blocks changed since the start of the cycle.

4. To achieve this, merge the consolidated delta save tape with each new delta save tape to get a new consolidated delta save tape.
5. Once a week, merge the full save tape and the consolidated delta save tape to get a new full save tape.
6. This completes the cycle and starts a new one.

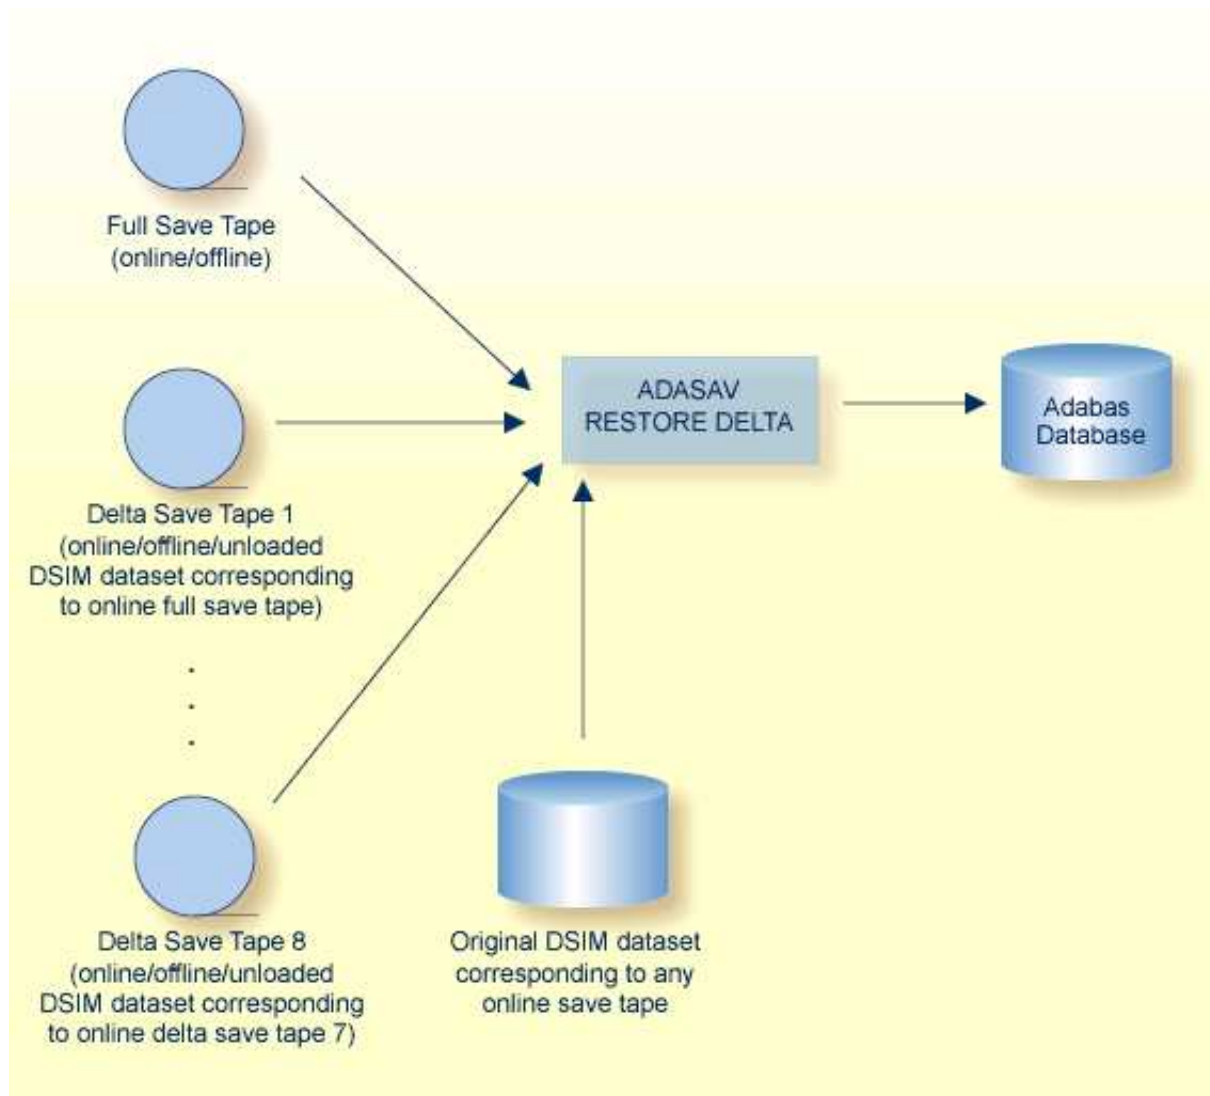
Restoring Delta Save Save Output

Merged full save tapes can be used for any ADASAV RESTORE function in the same way as conventional database save tapes. If delta saves have been taken since the production of the last full save tape, the full save and the delta saves can be merged to produce a new full save suitable for RESTORE.

To save time when recreating the database or files, full and delta saves can be merged and restored directly by using the ADASAV RESTORE DELTA function. The input for RESTORE DELTA is

- one full save tape;
- zero to eight (8) consecutive delta save tapes including unloaded DSIM data sets for each online save tape;
- the DSIM data set for the last online save unless the unloaded DSIM data set is provided.

-as shown in Delta Restore Option.



Consecutive saves must be specified for RESTORE DELTA: no save tape may be missing. Each save tape may itself be the output of an earlier merge operation. For tapes that originate from online save operations, the corresponding original or unloaded DSIM data set must be specified. An original DSIM data set is *not* reset at the end of the RESTORE DELTA operation but remains available for further merge operations.

If no delta save tape is supplied, the full save tape must be online and the corresponding DSIM data set, original or unloaded, must be specified. If only an offline full save tape is supplied, the conventional RESTORE function must be used.

If save and merge operations are performed as in the schedule described in the section Sample Schedule for Save and Merge Operations, at any time only the last (merged) full save tape, the consolidated delta save tape and, if present, the most recent yet unmerged delta save tape are required to restore the entire database to the status of the last delta save operation. This can be done with RESTORE DELTA.

Although the delta restore operation can restore the entire database, it is also possible to specify the GCB, FILE, or FMOVE parameters to restrict the delta restore operation, similar to a conventional restore.

On the input side, RESTORE DELTA is like MERGE since the input consists of a number of full save and delta save tapes. On the output side, RESTORE DELTA is like RESTORE since data is restored to the physical database.

If a delta restore operation is interrupted, it can be restarted, provided the full save has already been restored. This status is indicated by a message in the ADASAV protocol. In this case, the interrupted delta restore operation may be restarted by resubmitting the job, but without the job control specification for the full save tape. Then, ADASAV will continue the operation and restore the delta save tapes. If a restart is not possible, the whole delta restore operation must be repeated.


Operating a Shadow Database with Delta Save

The Delta Save Facility can be used to set up and maintain a "shadow" database; that is, a copy of the original database. Such a shadow database can

- accelerate reconstruction of the original database in case it is physically destroyed; and
- run read-only analysis programs to reduce the load on the original database.

The Delta Save Facility maintains a shadow database by

- restoring a full save of the original database into the shadow database's container files (see diagram Maintaining a Shadow Database, Step 1);
- performing delta save operations on the original database at regular intervals; and
- performing a delta restore into the shadow database with each delta save tape, without supplying a full save tape (see diagram Maintaining a Shadow Database, Step 2).



Such a shadow database has the following characteristics:

- The shadow database is always a slightly outdated copy of the original database. The size of the discrepancy depends on the time elapsed since the last delta save/delta restore operation.

- At any time, only an ADARES REGENERATE operation is needed to bring the shadow database to the latest status of the original database.
- This regenerate operation processes the PLOG starting from the point of the most recent delta save.
- Since an ET-synchronization is performed at the end of every online save (and when the Adabas nucleus is shut down), all full or delta save tapes, and hence the shadow database, always represent a synchronized status of the original database: all transactions are completed, there are no partial, incomplete transactions.
- Between each two delta restore operations into the shadow database, Adabas nuclei may be active on the shadow database in read-only mode to allow analysis programs to execute.
- During a delta restore operation, no read-only nucleus can be active.
- At no time may an update nucleus or an update utility (ADAINV, ADALOD, ADAORD, ADASAV) be active on the shadow database.
- Once an update nucleus or an update utility is run on the shadow database, it is no longer possible to restore a delta save tape into the shadow database without supplying a full save input.

The shadow database is kept up to date with the RESTORE DELTA function. The input requirements for RESTORE DELTA are as described in the section Restoring Delta Save - Save Output, except that the full save tape is omitted.

When the shadow database is updated by the sequence of delta restore operations, consecutive delta save tapes must be supplied; that is, no delta save may be omitted. This is checked using the DSIDs of the save tapes.

The delta save operations on the original database may be performed offline or online. When using online delta saves to update then shadow database, the DSIM data set processing applies as for normal RESTORE DELTA operations. However, an original DSIM data set is not freed at the end of the RESTORE DELTA operation, so that it can also be used for a merge operation with the online save tape.

Delta Save Functions

The Delta Save Facility provides the following additional or changed Adabas utility functions:

ADASAV

The following ADASAV functions have been added for use with the Delta Save Facility:

- MERGE to combine the latest full save result with one or more subsequent delta save results, in order to create a new full save tape, or to combine multiple delta save results into a single "consolidated" delta save tape;
- RESTORE DELTA to combine the latest full save result with one or more subsequent delta save results, in order to restore the combination directly into the database;
- SAVE DELTA to save only those blocks changed since the last full or delta save operation.

As the result of "online" SAVE DELTA operations, the images of changed database blocks are written to the Delta Save images (DSIM) data set, a direct access data set that must be located on disk. The DSIM data set is specified as

- DDDSIMR1 on BS2000, z/OS, and z/VM systems, or
- DSIMR1 on VSE/ESA systems.

New data sets used for the MERGE and RESTORE DELTA functions are specified as

- DDFULL and DDDEL1, DDDEL2, ... DDDEL8 on BS2000, z/OS, and z/VM systems, or
- FULL and DEL1, DEL2, ... DEL8 on VSE/ESA systems.

ADARES

The following ADARES functions have been revised for use with the Delta Save Facility:

- COPY to extract database block images written during an online save operation;
- PLCOPY to extract database block images written during an online save operation.

ADAREP

ADAREP utility operation has been revised to include Delta Save usage, save (full and delta), and logging information.

ADAFRM

The following ADAFRM functions have been added for use with the Delta Save Facility:

- DSIMFRM to format the Delta Save images (DSIM) data set, which is used to hold images of changed Adabas blocks during online save/merge operation;
- DSIMRESET to reset the DSIM data set, making it usable for later online save/merge operation.

ADAPRI

The ADAPRI DSIMPRI function prints one or more specified DSIM data set blocks.

ADAULD

The ADAULD utility has been modified to unload an Adabas file from a full save tape together with 1 to 8 delta save tapes and optionally a DSIM data set. The save tapes may have been created online or offline using any version of the Delta Save Facility.