

# SAVE: Save Database

The ADASAV SAVE (database) function saves the contents of the database to a sequential data set. It saves all blocks that are in use in the database.

The SAVE (database) function may be executed with the Adabas nucleus active or inactive. If executed while the Adabas nucleus is

- *active*, the RESTONL function must be used to later restore the database.
- *inactive*, the RESTORE function must be used to later restore the database.

In both cases, it is possible to restore just one or a few files from the database saved on the SAVE data set.

If the Adabas nucleus is *inactive*, it cannot be started while the SAVE function is executing, and no utility (such as ADALOD, for example) that makes changes to the database being saved can be run during the save. The SAVE function cannot be executed offline if a nucleus session autorestart is pending, or if another offline utility (such as ADALOD or ADASAV) is currently running.

If the Adabas nucleus is *active* during the execution of the save operation, users have full access to the database being saved. They can perform read, find, update, insert, and delete commands. However, utilities that make changes to the database to be saved (such as ADALOD, ADAINV, or ADADBS REFRESH, for example) must not be running and cannot be started while the save function is executing. An online save operation is also not possible if the nucleus is running without protection logging.

In an online save operation, the database to be saved may be changed while ADASAV is performing the save operation. Therefore, the Adabas nucleus writes all changed blocks to the protection log as well. This protection log must be supplied for a subsequent restore operation (that is, a RESTONL function).

The start of an online database save is marked by a SYN1 checkpoint. At the end of the online save, the nucleus synchronizes all currently active transactions. This means that Adabas performs no more update commands for users at ET status but allows the other active users to continue until they reach ET status. This status is then marked by a SYN2 checkpoint. The SYN2 checkpoint thus marks a consistent state of the database where no transactions are in progress. This state is reproduced when the database or files are restored from the SAVE data set later on.

The maximum time required for the transaction synchronization can be limited by the TTSYN parameter.

Databases residing on several disk volumes are saved to several SAVE data sets in parallel when the DRIVES parameter is specified. This mode of operation may significantly reduce the duration of the save. The resulting SAVE data sets, when concatenated in the order of ascending drive number, are equivalent to a single SAVE data set produced without the DRIVES parameter.

The SAVE (database) function does not save files that are in invert, load, refresh, reorder, or restore status. In fact, it removes such files from the file list, prints message ADAU15, and performs the save operation for the remaining files. At the end, ADASAV terminates with return code 4.

If the Recovery Aid (RLOG) option is active, the SAVE (database) function starts a new RLOG generation.

This chapter covers the following topics:

- Syntax
  - Optional Parameters
  - Example
- 

## Syntax

```
ADASAV SAVE [BUFNO = { number-of-buffers | 1 }]  
            [DRIVES = { count | 1 }]  
            [INCREMENTAL]  
            [NOUSERABEND]  
            [PERDRIVE = disk-drive-per-tape-drive , ...]  
            [TTSYN = seconds ]  
            [TWOCOPIES]  
            [TEST]
```

## Optional Parameters

### BUFNO: Count of Buffers

The BUFNO value allocates fixed buffers for the SAVE operation. A value of 2 or 3 usually provides optimum performance; up to 255 is possible. A value greater than 5, however, provides little advantage and allocates a lot of space. The default is 1 (one buffer per drive).

### DRIVES: Tape Drives for Parallel Save Processing

DRIVES is the number of sequential output data sets (usually on tape drives) to be used for parallel SAVE operations. A maximum of 8 drives may be specified. The default is 1.

### INCREMENTAL: Save Changed Files Only

INCREMENTAL saves only those files that have been changed since the last ADASAV SAVE operation. If INCREMENTAL is not specified, the SAVE function saves all database files.

### 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.

**PERDRIVE: Disk Drives Per Tape Drive**

PERDRIVE specifies the number of disk drives to be assigned to a single DRIVES tape drive. For example, if the database is contained on seven disk drives and three tape drives are available for SAVE 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<sub>n</sub> and DD/DUAL<sub>n</sub> job control specifications, as described at the end of this document.

The total number of drives specified by PERDRIVE must equal the sum of all Associator (ASSO) and DATA disks; if both ASSO and DATA 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 operation. Only the syntax of the specified parameters can be tested; not the validity of values and variables.

**TTSYN: SYN2 Checkpoint Control**

TTSYN allows the user to decrease the ADARUN TT (maximum transaction time) of the nucleus during the synchronized checkpoint processing of the current ADASAV operation. The value specified is the approximate time in seconds (TT &#65533; 1.05 seconds), and must be less than the current ADARUN TT value. If TTSYN is not specified or if TTSYN is greater than the TT value of the nucleus, that TT value becomes the default.

If the Adabas nucleus is active while ADASAV SAVE is running, a synchronized SYN2 checkpoint is taken at the end of the SAVE operation. This ensures that there is a point in time where all users are at ET status. If a user is not at ET status, no new transactions can be started for other users; they must wait until the SYN2 checkpoint can be taken.

The ADARUN TT value controls the maximum elapsed time permitted for a logical transaction. This is the maximum wait time until the SYN2 checkpoint can be processed. The ADASAV SAVE TTSYN parameter allows the user to decrease the TT value only during the synchronized checkpoint processing. The original TT value becomes effective again when ADASAV ends the SAVE operation.

**TWOCOPIES: Create Two Copies of Output**

TWOCOPIES creates two physical copies of the ADASAV output.

**Example**

```
ADASAV SAVE DRIVES=4
```

The SAVE function is to be executed using four tape drives in parallel.