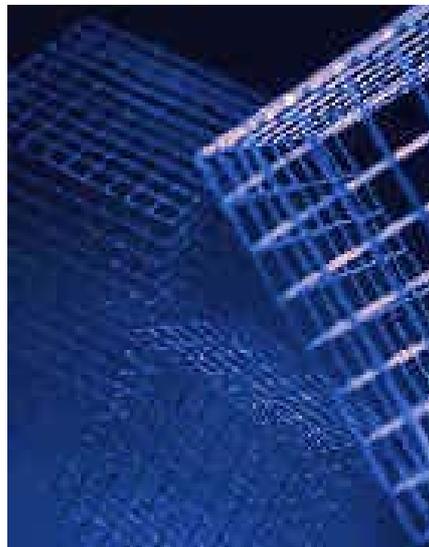


# Database Manager GUI: SAP DB



Version 7.4



## Copyright

© Copyright 2003 SAP AG.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation.

For more information on the GNU Free Documentaton License see <http://www.gnu.org/copyleft/fdl.html#SEC4>.

## Icons

Icon	Meaning
	Caution
	Example
	Note
	Recommendation
	Syntax

## Typographic Conventions

Type Style	Description
<i>Example text</i>	Words or characters that appear on the screen. These include field names, screen titles, pushbuttons as well as menu names, paths and options.  Cross-references to other documentation.
<b>Example text</b>	Emphasized words or phrases in body text, titles of graphics and tables.
EXAMPLE TEXT	Names of elements in the system. These include report names, program names, transaction codes, table names, and individual key words of a programming language, when surrounded by body text, for example, SELECT and INCLUDE.
Example text	Screen output. This includes file and directory names and their paths, messages, source code, names of variables and parameters as well as names of installation, upgrade and database tools.
EXAMPLE TEXT	Keys on the keyboard, for example, function keys (such as F2) or the ENTER key.
<b>Example text</b>	Exact user entry. These are words or characters that you enter in the system exactly as they appear in the documentation.
<Example text>	Variable user entry. Pointed brackets indicate that you replace these words and characters with appropriate entries.

Database Manager GUI: SAP DB 7.4 .....	6
Calling the Database Manager GUI .....	7
Executing Commands in the Database Manager GUI .....	7
Creating or Initializing a Database Instance .....	8
Managing Database Manager Operators .....	10
Changing Operator Passwords .....	10
Changing the Operator Authorizations .....	11
Assigning Authorizations for Operating System Users .....	12
Managing Database Instances .....	12
Terminating an Operator Session .....	13
Activating and Deactivating Indexes .....	13
Updating the Statistics Information .....	14
Changing the Log Mode .....	15
Changing the Log Mode from SINGLE to DUAL .....	15
Changing the Log Mode from DUAL to SINGLE .....	16
Changing Volume Parameters .....	16
Creating MapChar Sets .....	17
Defining Volumes .....	17
Displaying and Changing the Current Database Parameters .....	18
Displaying, Changing, and Deleting Mapchar Sets .....	19
Defining Volume Parameters .....	20
Closing a Database Session Manually .....	20
Registering Database Instances .....	21
Closing a Parameter Session .....	22
Starting the Database Instance (Restart) .....	22
Stopping the Database Instance (Shutdown) .....	23
Monitoring the Operational State .....	23
Connection to the DBM Server .....	24
Backing Up Data .....	25
Managing the Backup Media .....	25
Defining a Single Backup Medium .....	25
Defining a Group of Parallel Backup Media .....	26
Changing Backup Media .....	27
Deleting Backup Media .....	28
Backup Processes .....	29
Activating and Deactivating the Automatic Log Backup .....	29
Backing Up to a Single Backup Medium .....	30
Backing Up to a Group of Parallel Backup Media .....	31
Backups with Automatic Tape Loaders .....	32

Backing Up with External Backup Tools.....	33
Backups with Manually Changed Media.....	33
Restoring Data .....	34
Continuing an Interrupted Restore .....	35
Restoring a Backup from the Backup History .....	36
Restoring the Indexes After a Database Restore.....	37
Restoring the Last Complete Data Backup.....	37
Restoring an Existing Database Instance .....	38
Restoring a Mirrored Volume .....	39
Restoring from Automatic Tape Loaders .....	40
Restoring with External Backup Tools .....	40
Restoring Without a Backup History .....	41
Displaying Information .....	42
Displaying User Sessions.....	42
Displaying Cache Information .....	43
Displaying Data Area Information .....	43
Displaying Read and Write Operations .....	44
Displaying Log Area and Log Writer Information .....	44
Displaying the Backup History.....	45
Displaying Locks and Lock Requests .....	46
Displaying Versions.....	47
Options for Diagnosing Problems .....	47
Activating the Database Trace Function .....	47
Activating the Event Monitor.....	49
Using the Command Line Version of the Program.....	49
Checking the Logical Consistency of the Volumes .....	50
Displaying OMS Caches .....	50
Displaying OMS Class Containers .....	50
Activating the OMS Monitor.....	51
Checking the Physical Consistency of a Backup .....	51
Reading Database Files .....	52
Checking the Server of a Database Instance.....	52
Installing an Update of the Database Software .....	52
Updating the System Tables .....	53
Accessing the SQL Studio Program .....	53
Screen Areas of the Database Manager GUI.....	54



## Database Manager GUI: SAP DB 7.4

### Purpose

The Database Manager is the administration tool for SAP DB database instances. You can use this tool to create and monitor database instances on local or remote hosts. You can back up the data of database instances and, if necessary, recover it.

### Implementation Considerations

The Database Manager can be called as a commando line-oriented tool ([Database Manager CLI](#)), as a Web-based tool ([Web DBM](#)), or as a graphical user interface (Database Manager GUI). All three properties of this program offer the same functions.

The Database Manager CLI and the Web DBM can be installed on every operating system that is supported by the database system SAP DB. The Database Manager GUI runs on Microsoft Windows only.

Each property of the Database Manager can manage SAP DB database instances on every operating system that is supported by the database system SAP DB ([Operating System Platforms](#)).

You can use this version 7.4 of the Database Manager for all SAP DB database systems as of version 7.1.



To work with the Database Manager, a sound knowledge of database administration is required.



You can download the program Database Manager GUI free of charge from the Internet address [www.sapdb.org](http://www.sapdb.org). An installation wizard guides you through the installation process.

For general information about the database system SAP DB, see the documentation [SAP DB Database System](#) at the Internet address [www.sapdb.org](http://www.sapdb.org) → *Documentation*.

The **Database Manager GUI**, Version 7.4, is described in the following sections:

[Calling the Database Manager GUI \[Page 7\]](#)

[Executing Commands in the Database Manager GUI \[Page 7\]](#)

[Creating or Initializing a Database Instance \[Page 8\]](#)

[Managing Database Manager Operators \[Page 10\]](#)

[Managing Database Instances \[Page 12\]](#)

[Backing Up Data \[Page 25\]](#)

[Recovering Data \[Page 34\]](#)

[Displaying Information \[Page 42\]](#)

[Options for Diagnosing Problems \[Page 47\]](#)

[Installing an Update of the Database Software \[Page 52\]](#)

[Accessing the SQL Studio Program \[Page 53\]](#)

[Screen Areas of the Database Manager GUI \[Page 54\]](#)



## Calling the Database Manager GUI

### Procedure

To call the Database Manager GUI program, choose *Start* → *Programs* → *SAP DB* → *Database Manager*.

You can also call the Database Manager GUI (`dbmgui.exe`) from the command line. To do so, enter the following command: `dbmgui [<options>]`.

#### Options for the Command Line Call

Display all command line options	-?
Name of the <a href="#">database instance</a>	-d <database_name>
Name of the <a href="#">X Server</a> host	-n <database_server>
Opening a <a href="#">database session</a> specifying operator data (operator name and password)	-u <user_name>, [<password>]
Open the configuration file	-f <file_name>
Window for entering connection details	-prompt
Database Manager GUI version	-version
Call the Database Manager without initial screen display	-nologo
Activate <a href="#">database trace</a>	-trace



## Executing Commands in the Database Manager GUI

### Use

You can call all of the functions of the Database Manager using the menu bar at the top of the screen ([Screen Areas of the Database Manager GUI \[Page 54\]](#)).

Every time you select an option from a menu, a new window is opened in the display.

### Procedure

To switch between multiple open windows, choose *View* → *Previous Window* | *Next Window* | *List of Windows*.

A window is only closed by the Database Manager once you have made the required entries and confirmed them; that is, when the command has been executed.



**Special Cases:**  
**Changing the Log Mode and Changing Volume Parameters**

As changing parameters in a parameter session and changing volume parameters or the log mode are mutually exclusive, if there is an active parameter session, you must first close the parameter session and can then change volume parameters or the log mode.

To do this, choose *View* → *All Windows*. Select the entry for the parameter session and choose *Close Window*.

Proceed as described under [Changing Volume Parameters \[Page 16\]](#) and [Changing the Log Mode \[Page 15\]](#).



## Creating or Initializing a Database Instance

### Prerequisites

If you want to create the [database instance](#) locally and the Microsoft Windows operating system is installed on your host, your operating system user must have administration rights.

If you want to create the database instance on a remote host and the operating system Microsoft Windows is installed, then the operating system user must possess both administration rights and also the *log on as batch job* authorization. (You can find information about maintaining operating system users in your operating system documentation.)

If you want to initialize an existing database instance, you must possess administration rights for this database instance and the database server.

### Procedure

Select *Instance* → *Install* to open the *Database Wizard*. This installation wizard now guides you through the whole process for creating a database instance.

1. Enter the name of the database instance ([Database name](#)). This can also be the name of an existing database instance, if you want to initialize this existing instance. In this case, choose *Next* to go directly to step 3.  
If the new database instance is to be created on a remote server, you must also specify the name of this server and your operating system user name and password for this server.

Choose *Next*.

2. There may be more than one installation of the database software on the same server. The different installations may also be of different versions. Of the installations present on the server, select the one you want.

Choose *Next*.

3. Enter the [database manager operator DBM](#) and the [database system administrator SYSDBA](#), and assign passwords for *DBM* and *SYSDBA*.  
When specifying *SYSDBA* and when later using it in the SQL environment, note that the Database Manager automatically converts all letters to uppercase.

If you are planning to copy the parameter file from a backup, enter the name and password of the *SYSDBA* saved in this parameter file.

Choose *Next*.

4. Select one of the following options:

*Initialize parameters with default values*

When the parameter file for the database instance is created, the standard configuration stored when the database software was installed is used.

*Use current parameters*

You only have this option if you are initializing an existing database instance. You use its parameters.

*Copy parameters from existing database*

You want to use the configuration of a database instance on the server.

Select the database instance from which the parameter file is to be copied and choose *Next*.

*Restore parameters from a backup medium*

You want to use the parameter file in a complete [data backup](#) to create the new database instance. Select the desired complete data backup.

5. Choose an [instance type](#) for the database.  
If you chose *Use current parameters* in the previous steps, the system uses the instance type of the existing database instance.
6. Adjust the parameters to your requirements. Select the required parameter and then choose *Edit*.  
It now appears in the bottom part of the window. Next to the parameter you will see an explanation of what the parameter does and, in some cases, how it is calculated.  
  
In the *New Value* field, enter the new value for the parameter and confirm your entry. The new value appears in the *New Value* column. It is stored in the internal data structures and becomes effective when the database instance is next restarted.  
  
Choose *Next*.  
When you exit the input screen, the parameters are checked in accordance with the rules stored on the server. You may be asked to make and confirm changes before you can leave the input screen.

7. Define the properties of the [log volumes](#) and the [data volumes](#). When doing so, bear in mind the parameters set in the previous step.  
  
Select a [volume](#) and then choose *Edit*. Enter the size and ID or the absolute path of the volume and confirm your entries. Repeat this process for each volume.

Choose *Next*.

8. Choose between the following options:

*Install and start instance*

A completely new database instance is created with the parameters in point 6., and is switched to the `ONLINE` operational state.

*Install instance for recovery*

The database instance is reinitialized with the parameters in point 6., and is ready to import backups. The database instance can then be recovered ([Recovery \[Page 34\]](#)), or a homogeneous system copy can be created ([Creating a Homogeneous System Copy](#)).

Choose *Next*.

9. Choose *Install* to start the creation of the new database instance or the initialization of the specified instance.

You can see how the generation or initialization of the database instance is progressing. If any errors occur, the process is stopped immediately. However, the state that had been reached when the error or fault occurred is preserved.

Before you complete this step, the system prompts you to enter the name under which the database instance is to be registered in the Database Manager GUI. Confirm your entries.

10. Select *Close* to close the *Database Wizard*.  
You are taken back to the start screen of the Database Manager GUI.

## Result

If you created a new database instance, its [operational state](#) is `ONLINE`.

If you initialized an existing database instance, its operational state is `ADMIN`. Now you can recover a database instance or create a homogeneous system copy of a database instance.



## Managing Database Manager Operators

The SAP DB database system differentiates between two main user classes:

- [Database Manager operators \(DBM operators\)](#)
- [Database Users](#)

When creating new database instances, you must first create the first user of each of these classes for each instance. These are the **first Database Manager operator** and the **database system administrator (SYSDBA)**.

Each of these users then has the authorization to create additional users of the specified user classes.

**See also:** *User Manual: SAP DB* → [User Concept](#)



Note that a Database Manager operator can only manage a database instance but cannot work with the database objects.

A database system administrator, on the other hand, can work with the database objects and also has the authorizations of the first Database Manager operator. He or she can also perform all Database Manager functions.

[Changing User Passwords \[Page 10\]](#)

[Changing Server Authorizations \[Page 11\]](#)

[Assigning Authorizations for Operating System Users \[Page 12\]](#)



## Changing Operator Passwords

### Use

You can change the passwords of all DBM operators and the SYSDBA (first database user) using the Database Manager in various [operational states](#) of the [database instance](#):

User Class	User	Operational State of Database Instance
<a href="#">Database Manager Operator</a>	First Database Manager operator All other DBM operators	OFFLINE Operator password can be changed in any operational state
Database User	<a href="#">SYSDBA</a>	ONLINE

### Prerequisites

You have transferred the database instance into the appropriate operational state (see [Starting the Database Instance \[Page 22\]](#), [Stopping the Database Instance \[Page 23\]](#)).

## Procedure

1. Choose *Instance* → *Configuration* → *Users*.
2. Select the required operator and then choose *Users* → *Edit*.
3. Select the tab page *General* and then choose *Change Password*.
4. Enter the old and new passwords and save your entries.



## Changing the Operator Authorizations

### Use

An operator authorization is an authorization to execute certain commands in the Database Manager.

An authorization may cover more than one command and one command may have more than one authorization assigned to it.

Database Manager has the following operator authorizations:

<i>Request status data</i>	Requesting status information
<i>Execute the LOAD program</i>	Running the LOAD program
<i>Execute operating system commands</i>	Executing operating system commands
<i>User management</i>	User Management
<i>File access (read and write)</i>	File access (read and write)
<i>File access (read only)</i>	File access (read only)
<i>Database file access (read only)</i>	Database file access (read only)
<i>Saving backups</i>	Carrying out backups
<i>Installation management</i>	Installation management
<i>Load the system tables</i>	Loading system tables
<i>Parameter access (checked write)</i>	Parameter access (checked write)
<i>Parameter access (read and write)</i>	Parameter access (read and write)
<i>Parameter access (read only)</i>	Parameter access (read only)
<i>Start database instance</i>	Starting the database instance
<i>Stop database instance</i>	Stopping the database instance
<i>Restoring backups</i>	Restoring backups
<i>Access to SQL session</i>	Accessing SQL session
<i>Access to utility session</i>	Accessing a utility session

### Prerequisites

You have the *User Management* operator authorization.

### Procedure

1. Choose *Instance* → *Configuration* → *Users*. Select the operator whose operator authorizations you want to change and then choose *Edit*.

2. On the *User Rights* tab page, you can now extend or limit the operator authorizations of the selected operator by selecting or deselecting the desired authorizations.
3. Choose *OK*.



## Assigning Authorizations for Operating System Users

### Use

To allow [Database Manager operators \(DBM operators\)](#) to manage [database instances](#) remotely, you can also grant them authorizations as operating system users for the servers of these database instances.

### Prerequisites

You are an operating system user on the database server and also have DBM operator authorizations for the database instance on this server.

### Procedure

1. Choose *Instance* → *Configuration* → *Users*.
2. Select the required operator and choose *Users* → *Edit*.  
On the *System Account* tab page, enter the name and password of the operating system user. Enter the password again.
3. Choose *OK*.

### Result

The selected operator also has operating system user authorization on the server of the current database instance.



## Managing Database Instances

[Activating and Deactivating Indexes \[Page 13\]](#)

[Updating the Statistics Information \[Page 14\]](#)

[Changing the Log Mode \[Page 15\]](#)

[Changing Volume Parameters \[Page 16\]](#)

[Creating Mapchar Sets \[Page 17\]](#)

[Defining Volumes \[Page 17\]](#)

[Displaying and Changing the Current Database Parameters \[Page 18\]](#)

[Displaying, Changing, and Deleting Mapchar Sets \[Page 19\]](#)

[Defining Volume Parameters \[Page 20\]](#)

[Manually Ending a Database Session \[Page 20\]](#)

[Registering Database Instances \[Page 21\]](#)

[Closing a Parameter Session \[Page 22\]](#)

[Starting the Database Instance \(Restart\) \[Page 22\]](#)

[Stopping the Database Instance \(Shutdown\) \[Page 23\]](#)

[Monitoring the Operational State \[Page 23\]](#)

[Connection to the DBM Server \[Page 24\]](#)



## Terminating an Operator Session

### Use

In special cases, it can be necessary to terminate an operator session.

### Prerequisites

The [database instance](#) is in the `ONLINE` operational state.

### Procedure

Choose *Instance* → *Sessions*.

Select the operator session(s) that you want to terminate.

Choose *Sessions* → *Kill Selected Sessions*.



## Activating and Deactivating Indexes

### Use

You can determine which of the existing indexes of a [database instance](#) are used.

You can also activate or deactivate selected indexes and thereby change the access times.

### Prerequisites

The database instance is in the `ONLINE` [operational state](#).

### Procedure

1. Choose *Instance* → *Tuning* → *Index Use*.
2. You can choose which indexes are to be displayed. To do so, enter appropriate search arguments under `Owner` and `Table Name` or `Index Name`.

Owner	Owner
Table Name	Name of the table
Index Name	Name of the index

You can also choose between the following options:

Disabled indexes only	Only indexes that have been deactivated
Unused indexes only	Only indexes that have not been used

3. Then choose *Indexes* → *Search*. The system displays a list of the indexes that match your criteria.

4. Select the indexes you want to activate in the list. If you want to activate all of the indexes displayed, choose *Indexes* → *Mark All*.  
Deselect the indexes you want to deactivate in the list.
5. Choose *Indexes* and the desired action.



## Updating the Statistics Information

### Use

Various statistics are compiled for each [database instance](#), such as the number of table entries, the size of the tables and indexes, and the value distribution (various values) of indexes and columns. This information is required by the [SAP DB Optimizer](#) to define the best strategy for executing complex SQL statements.

The statistics information is stored in the [database catalog](#).

If the database size or the values it contains have changed considerably, you will have to update the statistics. This should be carried out roughly once a week.

The statistics values can be updated for certain tables, columns, or for all base tables.

### Prerequisites

The database instance is in the ONLINE [operational state](#).

### Procedure

Choose *Instance* → *Tuning* → *Optimizer Statistics*.

#### Updating statistics entries for all tables

If you then choose *OptimizerStatistics* → *Execute*, all of the statistics entries for **all** of the tables are updated.

#### Updating statistics entries for selected tables or columns

You can also choose to update **selected** statistics only:

To do so, enter appropriate search arguments under *Owner*, *Table Name* and *Column Name*.

You can specify the size of sample by entering appropriate values in the *Estimate Rows* and *Estimate Percent* fields. The update is then only carried out across the number of rows or the specified percentage, based on the total number of database tables. This option, therefore, is faster.



The significance of the random sample depends on the size of the tables and the physical position of the data. Measurements have shown that for large tables, random samples 1000 and 5000 rows have short response times and produce high-quality results.

<i>Owner</i>	Owner
<i>Table Name</i>	Table name
<i>Column Name</i>	Column name
<i>Estimate Rows</i>	Number of rows checked
<i>Estimate Percent</i>	Percentage of rows checked

Then choose *OptimizerStatistics* → *Search*. A list of the tables that match your criteria is then displayed. From the `Update Statistics Date` and `Update Statistics Time`, you can tell when the statistics for this table were last updated.

If you want to update the statistics for **one of the tables** in the list, simply select the relevant entry. If you want to update the statistics for **all** of the tables, choose *OptimizerStatistics* → *Mark All*.

Choose *OptimizerStatistics* → *Execute* to update the statistics for the selected tables.



## Changing the Log Mode

[Changing the Log Mode from SINGLE to DUAL \[Page 15\]](#)

[Changing the Log Mode from DUAL to SINGLE \[Page 16\]](#)



For SAP DB database system up to version 7.3, you can also use the functions:  
Activate the DEMO variant of the log mode  
Deactivate the DEMO variant of the log mode  
**See also: *Database Manager GUI: SAP DB 7.3***



## Changing the Log Mode from SINGLE to DUAL

### Use

You can change the [log mode](#) for the current [database instance](#).

### Prerequisites

There are no open parameter sessions ([Closing a Parameter Session \[Page 22\]](#)).

The database instance can be in any [operational state](#).

The current log mode is *SINGLE*.

### Procedure

Under *New LogMode* select option *DUAL* to mirror the [log volume](#).

If no mirror volume has yet been defined for the log volume, you must do this now:

- Select the log volume for which you want to define a mirror volume, and choose *Edit*.
- Under *Mirrored*, define the path and the type of the mirror volume.

To integrate the mirror volume in the database instance, you must restart the database instance. If you are not in the *OFFLINE* operational state, select the option *Restart Database and reintegrate mirror volumes implicitly after LogMode change*.

Confirm your entries.

Choose *LogMode* → *Change LogMode*.

### Result

The new log mode *DUAL* is displayed as the current log mode.



## Changing the Log Mode from DUAL to SINGLE

### Use

You can change the [log mode](#) for the current [database instance](#).

### Prerequisites

There are no open parameter sessions ([Closing a Parameter Session \[Page 22\]](#)).

The database instance can be in any [operational state](#).

The current log mode is *DUAL*.

### Procedure

Under *New LogMode*, select the option *SINGLE*.

To change the log mode in the parameter file of the database instance, you must restart the database instance. If you are not in the *OFFLINE* operational state, select the option *Restart Database implicitly after LogMode change*.

Confirm your entries.

Choose *LogMode* → *Change LogMode*.

### Result

The new log mode *SINGLE* is displayed as the current log mode.



## Changing Volume Parameters

### Prerequisites

You have copied the [log volume](#) or [data volume](#) to its new path and you wish to inform the Database Manager GUI of this new path.

The [database instance](#) is in the *OFFLINE* [operational state](#).

There are no open parameter sessions ([Closing a Parameter Session \[Page 22\]](#)).

### Procedure

1. Choose *Instance* → *Configuration* → *Data Volumes* or *Log Volumes*.  
You are shown all the existing volumes of the selected type.

Name	Name of the volume
Size	Size of the volume
Type	File, raw device, or symbolic link
Location	Path of the volume
M. Type	File, raw device, or symbolic link (mirrored volume)
M. Location	Path of the mirrored volume

2. Choose *Volumes* → *Edit*.

3. Select the log or data volume to be changed. It now appears in the bottom part of the window where it can be edited.  
Enter the new path of the volume under `Location`.
4. Save your entries.  
When you now [restart the database instance \[Page 22\]](#) the database system knows the new path.



## Creating MapChar Sets

### Use

You can use this command to create MapChar sets. **See also:** [Language Support \(MapChar Sets\)](#)

You cannot use this command for [Unicode](#) databases.

### Procedure

1. Choose *Instance* → *Configuration* → *MapChar Sets*.  
You are shown the MapChar sets that are already available.
2. Choose *MapChar Sets* → *New*.

Define the name of the new [MapChar set](#).

Select the code type used for the source files (ASCII code in accordance with ISO 8859/1.2 or EBCDIC code CCSID 500, code page 500).

Define the conversions for the characters required. Each one-byte character to be converted must be specified in its original form and in the target form with a maximum length of two bytes.

Name	Name of the MapChar set, a maximum of 18 bytes in length
Code type	The code underlying the MapChar set Valid codes are ASCII code in accordance with ISO 8859/1.2 or EBCDIC code CCSID 500, codepage 500.
Internal	Original form in hexadecimal format
External Hex	Target form in hexadecimal format
Ext. ASCII	Target form as a printable character

3. Choose *Save* to save your entries.



## Defining Volumes

### Use

You want to extend the storage capacity of your [data](#) or [log area](#).

Volume Type	Database Parameter
<a href="#">Data Volume</a>	MAXDATAVOLUMES

[Log Volume](#)

MAXARCHIVELOGS

## Prerequisites

The [database instance](#) is in the `ONLINE` [operational state](#).

The value of the database parameter for the volume type is equal at least to the total of existing and planned [volumes](#).



The system checks that there is enough space available for at least one volume of each type before each start of the database instance. If there is not, it automatically increases the setting of the appropriate database parameter by 1.

If you wish to enlarge the database instance by more than one volume of a type, adjust the relevant parameter accordingly ([Defining Volume Parameters \[Page 20\]](#)). Later on, this enables you to add a volume of the given type a number of times in succession with the database instance in the `ONLINE` operational state without having to stop the database instance.

## Procedure

1. Choose *Instance* → *Configuration* → *Data Volumes* or *Log Volumes*.  
The system displays all of the existing volumes of the selected type and their data, as well as the number of volumes still possible in accordance with the parameter value that was set.

Name	Name of the volume
Size	Size of the volume
Type	File, raw device, or symbolic link
Location	Path of the volume
M. Type	File, raw device, or symbolic link (mirrored volume)
M. Location	Path of the mirrored volume

2. Choose one of the proposed volumes.
3. Choose *Volumes* → *Edit*. Define the new volume by entering the size, type, and path.
4. Save the entries you made.



## Displaying and Changing the Current Database Parameters

### Use

You can display and change the [database parameters](#) of the current [database instance](#). This is possible in any [operational state](#) of the database instance.

The changes are stored in the parameter file but only become effective after the database has been restarted ([Starting the Database Instance \(Restart\) \[Page 22\]](#)).

## Procedure



To change the parameters [MAXDATAVOLUMES](#) and [MAXLOGVOLUMES](#), close the parameter session.

Proceed as described in [Changing Volume Parameters \[Page 16\]](#).

For all other parameters, follow the procedure below:

1. Choose *Instance* → *Configuration* → *Parameters*. The system displays the general database parameters (*General*).
2. If you also want to display the special and support database parameters, choose *Parameters* → *Show Extended and Support Parameters*.
3. Select the desired tab page and the parameter that you want to change.

Name	Name of database parameter
Value	Value of the database parameter currently being used in the kernel of the database instance
New value	<p>No entry: The database parameter has not been changed since the database instance was last started.</p> <p>New value entered: The database parameter has been changed since the database instance was last started. However, no changes are saved to the parameter file and none of them become active until the next time the database instance is restarted (<a href="#">Starting the Database Instance (Restart) [Page 22]</a>).</p>
Description	Help text on the selected parameter

4. Choose *Parameter* → *Edit*  
The parameter appears in the bottom part of the window again. Next to the parameter you will see an explanation of what the parameter does and, in some cases, how it is calculated.
5. Enter the new value for the parameter in the `New Value` field.  
Confirm your entries.  
The new value appears in the *New Value* column. It is stored in the internal data structures and becomes effective when the database instance is next restarted.



## Displaying, Changing, and Deleting Mapchar Sets

### Procedure

1. Choose *Instance* → *Configuration* → *Mapchar Sets*.  
The system displays the [Mapchar sets](#) that are already available.
2. If you want to change an available Mapchar set, select it and then choose *Mapchar Sets* → *Edit*.  
If you want to delete it, choose *Mapchar Sets* → *Delete*.

If you want to create a new Mapchar set, choose *Mapchar Sets* → *New*.

Name	Name of the Mapchar set, a maximum of 18 bytes in length
Code type	The code on which the Mapchar set is based. The valid codes are

	the ASCII code in accordance with ISO 8859/1.2 or the EBCDIC code CCSID 500, codepage 500.
Internal	Original form in hexadecimal format
External Hex	Target form in hexadecimal format
Ext. ASCII	Target form as a printable character

3. Make the required changes or entries and choose **Save** to save these.



## Defining Volume Parameters

### Use

You define the maximum possible number of volumes for a database instance using the relevant [database parameters](#).

Volume Type	Parameter
<a href="#">Data Volume</a>	MAXDATAVOLUMES
<a href="#">Log Volume</a>	MAXARCHIVELOGS

You can change the settings of both parameters at a later date regardless of which [operational state](#) the database instance is in. Like all parameter changes, these changes take effect the next time the database instance is restarted.

### Procedure

1. First, follow the procedure described in [Displaying and Changing Current Database Parameters \[Page 18\]](#).
2. Then transfer the database instance to the `ADMIN` operational state.
3. Restart the database instance.  
The Database Manager GUI transfers the database instance first into `OFFLINE` operational state and then into `ONLINE` operational state and saves the new value in the parameter file. It reserves the amount of memory space required for this number plus a reserve equal in size to the largest volume of this type. This reserve allows you to create an additional volume of this type while the database is running if the number originally planned is insufficient.



How to create volumes is described in [Defining Volumes \[Page 17\]](#).



## Closing a Database Session Manually

### Use

In certain cases, it may be necessary to stop a specific database session.

### Procedure

Choose *Instance* → *Sessions*. A list of the current [database sessions](#) is displayed.

Select the database session that you want to close.

Choose *Sessions* → *Kill Selected Sessions* to close the database session.



## Registering Database Instances

### Use

If [database instances](#) are to be managed by the Database Manager GUI, they have to be registered.

### Prerequisites

If you want to access a remote database instance with the Database Manager, you have to start the **X Server** service on the relevant database server (`Status: started`).

### Procedure

1. Choose *File* → *Register Database*  
Enter the name of the server from which you wish to request a list of the database instances installed on it.  
Choose *Enter* or the *exclamation mark (!)* next to the input field.
2. On the list of the installations present on the server that is displayed, select the one that is to be registered.  
Choose *Register*.
3. Enter a name for the database instance you want to register and the DBM operator of the database instance, which comprises the DBM operator name and the DBM password.



You can choose any name as the registered name of the database instance. It only needs to be unique within the Database Manager GUI. Generally speaking it should be the same as the name defined for the database instance during installation. If this has already been assigned you must use a different registered name.

4. Choose *OK*.

### Result

The database instance has been registered and from now on it will appear on the list of registered databases.



To organize your directory of registered database instances, choose *File* → *New Folder* | *Delete* | *Rename*.



## Closing a Parameter Session

### Use

Every call of a function to change parameters by choosing *Instance* → *Configuration* → *Parameters* opens a parameter session. This is only closed by the system once the command has been transferred to the Database Manager; that is, once you have made and confirmed the required entries.

If you do not want to complete the entries for a command to change parameters, you must end the parameter session manually.

### Procedure

To do this, choose *View* → *All Windows*. Select the entry for the parameter session and choose *Close Window*.



## Starting the Database Instance (Restart)

### Prerequisites

You have selected the [database instance](#) you wish to start.

### Procedure

#### Transferring the Database Instance to the ONLINE Operational State

To start the database instance, that is to transfer it to a running state, choose *Instance* → *Set Online*.

This function can be performed in both the ADMIN [operational state](#) and the deactivated operational state (OFFLINE).



If this command fails to set the database to the ONLINE operational state, you can look for the cause of the error in the log file `knldiag`.

#### Transferring the Database Instance to the ADMIN Operational State

Choose *Instance* → *Admin* to transfer the database instance from the deactivated operational state (OFFLINE) to the ADMIN operational state.



Each time you start the database instance by choosing *Instance* → *Online* or *Instance* → *Admin*, a copy of the console log file (file `knldiag` in the [run directory](#) of the database instance) is created with the name `knldiag.old`, and file `knldiag` is overwritten.

See also: [Database Files](#)



## Stopping the Database Instance (Shutdown)

### Prerequisites

You have selected the [database instance](#) that you wish to stop.

The database instance is in the `ONLINE` [operational state](#).

### Procedure

Transferring the Database Instance to the <code>ADMIN</code> Operational State	Transferring the Database Instance to the <code>OFFLINE</code> Operational State
Choose <i>Instance</i> → <i>Admin</i> , if you want to transfer the database instance from the <code>ONLINE</code> operational state to the <code>ADMIN</code> operational state.	Choose <i>Instance</i> → <i>Offline</i> if you wish to shut the database instance down.



Both commands fail, if a [data backup](#) or [log backup](#) is running or if the [automatic log backup](#) is active. An error message tells you that the backup action has to be completed first or the automatic log backup disabled before the database instance can be stopped.

### Options on Stopping Database Instance

<i>OK</i>	First the system waits until all transactions have been completed correctly, then it stops the database instance.
<i>Cancel</i>	You cancel the command given.



## Monitoring the Operational State

### Prerequisites

You have selected the [database instance](#) of which you want to check the [operational state](#).

### Operational State `ONLINE` | `ADMIN` | `OFFLINE`

There are various ways in which you can find out the current operational state of the database instance.

- Check the indicator that is assigned to the general icon for the database instance in the display of the registration name.
- Choose *Instance* → *Details* and check the indicator that is assigned to the general icon for the database instance in the display.
- Choose *View* → *Large Icons* | *Small Icons* | *List* and check the indicator that is assigned to the general icon for the database instance in the list of registered database instances

	General icon for a database instance
	Database instance is in the <code>OFFLINE</code> operational state

	Database instance is in the <code>ADMIN</code> operational state
	Database instance is in the <code>ONLINE</code> operational state

- Choose *View* → *Details* and read the operational state of the current database instance from the *State* column of the list of registered database instances.



## Connection to the DBM Server

Before you can manage [database instances](#) using the Database Manager GUI program, a connection must be set up between the database instance to be managed and your [DBM Server](#).

The system is configured in such a way that the connections between the database instances displayed and the DBM Server are established automatically when the Database Manager GUI is launched.



You can also change this default setting if necessary. Choose *View* → *Options*. If you deselect the `Autoconnect instances` option on the *General* tab page, the system no longer creates connections between the database instances and the DBM Server when the Database Manager GUI is called.

## Use

You can selectively set up the connection to the DBM Server for one, multiple, or all registered database instances, depending on the performance of your system. Similarly, you can terminate the connection between individual database instances and your DBM Server in order to reduce the load on your system.

## Procedure

### Setting up a connection between a database instance and the DBM Server

The entry for the state of the database instance `State` is `Not connected`.

Choose the relevant database instance with the cursor. The system then sets up the connection to the DBM Server for exactly this database instance. The Database Manager can now also determine the current [operational state](#) of the selected database instance.

### Terminating a connection between a database instance and the DBM Server

To terminate the connection between a database instance and the DBM Server, select the desired instance and choose *Instance* → *Disconnect*.

The entry for the state of this database instance `State` is now `Not connected`. The connection to the DBM Server has been terminated and the Database Manager is also no longer able to determine the current operational state of the database instance.



## Backing Up Data

### [Database Manager CLI: SAP DB 7.4](#)

With the Database Manager, you can perform complete and incremental [data backups](#) and [log backups](#). Both types of backup can be activated either interactively using the Database Manager GUI or by a background call using the Database Manager CLI.

For log backups, we recommend that you activate automatic log backup ([Activating and Deactivating the Automatic Log Backup \[Page 29\]](#)). However, you can also perform an [interactive log backup](#) to a file.

For SAP applications, the backup function of the Database Manager can be optimally scheduled and automated using the DBA planning calendar of the Computing Center Management System (CCMS). Log files are created for the performed backups in the SAP system logs and these give you an accurate picture of what happened during the backups.

#### See also:

### [Database Manager CLI: SAP DB 7.4](#)

*Computing Center Management System*, section [DBA Planning Calendar](#)



## Managing the Backup Media

### [Defining a Single Backup Medium \[Page 25\]](#)

### [Defining a Group of Parallel Backup Media \[Page 26\]](#)

### [Changing Backup Media Definitions \[Page 27\]](#)

### [Deleting Backup Media \[Page 28\]](#)



## Defining a Single Backup Medium

### Procedure

1. Choose *Instance* → *Configuration* → *Backup Media*
2. Choose *BackupMedia* → *New* → *Medium*.
3. Select *General* and make the following entries:

Name :	Freely definable name to identify the medium  If you want to perform a <a href="#">backup [Page 25]</a> or a <a href="#">restore [Page 34]</a> using one of the external backup tools ADSM/TSM, NetWorker, Backint for Oracle, or Backint for SAP DB, start the name of the <a href="#">backup medium</a> with ADSM, NSR, BACK, or BACK ( <a href="#">Backing Up with External Backup Tools [Page 33]</a> ).
Location:	Specify the location to which backups are to be made.  With files and pipes, we recommend specifying the absolute path to the backup medium.
Device Type:	Select the type of the backup medium: File, tape, or pipe.

	<p>If you back up the log to a file, the Database Manager GUI automatically generates version files. An automatic <a href="#">log backup</a> cannot be made to tape.</p> <p>If you are saving with external backup tools, enter <code>pipe</code> as the <code>Device Type</code>.</p> <p>If you are backing up to tape, you need to address the device driver correctly under the UNIX operating system. Use rewind tapes for backups. On Microsoft Windows NT/Windows 2000 operating systems, the system sets the device driver automatically.</p>
Backup Type:	Select the type of backup you want (full or incremental <a href="#">data backup</a> , interactive log backup, or automatic log backup)

4. Choose *Extended* and make the following entries:

Size:	The size of the backup medium (in KB)
Overwrite:	Select the <code>Overwrite</code> field if you want to use the overwrite mode for the medium.
Autoloader:	Select <code>Autoloader</code> to use a tape device with autoloader facilities ( <a href="#">Backing Up with Automatic Tape Loaders [Page 32]</a> ).
OS command:	You can enter any operating system commands here that are appropriate for backing up to tape.



## Defining a Group of Parallel Backup Media

### Use

To allow you to [backup \[Page 25\]](#) at a higher speed, you can back up to a number of media in parallel. To do this, define a [group of parallel backup media](#) under a single medium name. A backup that is to be made to a number of media simultaneously should have precisely this medium name assigned to it.

The maximum number of parallel backup media in a group is determined by the value set in the parameter `MAXBACKUPVOLS`.

When you [restore \[Page 34\]](#), you can import the data from this type of backup in parallel. The number of backup media from which the import can be performed in parallel is also dependent on the value of the `MAXBACKUPVOLS` parameter.

### Procedure

1. Choose *Instance* → *Configuration* → *Backup Media*.
2. Select *BackupMedia* → *New* → *Media Group*.  
In the `Configuration - Backup Media` window, the system creates the media group `New Group` in the form of a folder whose media type is `Parallel`.
3. Assign any name to this media group. You can change this name at a later stage using *BackupMedia* → *Rename*.
4. Double click the folder to open it.
5. Now define each individual medium in the media group:  
Choose *BackupMedia* → *New* → *Medium*.

Select **General** in the **Configuration - Backup Media** window and make the following entries:

Name:	<p>Freely definable name to identify the medium</p> <p>If you want to perform a <a href="#">backup [Page 25]</a> or a <a href="#">restore [Page 34]</a> using one of the external backup tools AD SM, NetWorker, Backint for Oracle, or Backint for SAP DB, start the name of the medium with AD SM, NSR, BACK, or BACK (<a href="#">Backing Up with External Backup Tools [Page 33]</a>).</p>
Location:	<p>Specify the location to which backups are to be made.</p> <p>With files and pipes, we recommend specifying the absolute path to the backup medium.</p>
Device Type:	<p>Select the type of the backup medium: File, tape, or pipe.</p> <p>If you back up the log to a file, the Database Manager GUI automatically generates version files. An automatic <a href="#">log backup</a> cannot be made to tape.</p> <p>If you are saving with external backup tools, enter <b>Pipe</b> as the Device Type.</p> <p>If you are backing up to tape, you need to address the device driver correctly under the UNIX operating system. Use rewind tapes for backups. On Microsoft Windows NT/Windows 2000 operating systems, the system sets the device driver automatically.</p>
Backup Type:	<p>Select the type of backup you want (full or incremental <a href="#">data backup</a>, interactive log backup, or automatic log backup)</p>

Select **Extended** in the **Configuration - Backup Media** window and make the following entries:

Size:	Size of the backup medium
Overwrite:	Select the <b>Overwrite</b> field when you want to use the overwrite mode for the medium (only possible with media type <b>File</b> ).
Autoloader:	Select <b>Autoloader</b> to use a tape device with autoloader facilities ( <a href="#">Backing Up with Automatic Tape Loaders [Page 32]</a> ).
OS command:	You can enter any operating system commands here that are appropriate for backing up to tape.



## Changing Backup Media

### Procedure

1. Choose **Instance** → **Configuration** → **Backup Media**  
The system displays a list of the [backup media](#) already created.
2. Select the desired backup medium.
3. Follow the appropriate description below:

## Change Media Name

Choose *BackupMedia* → *Rename* and make your changes.

Choose *Enter* to save your changes.

## Change the Properties of a Medium

Select *BackupMedia* → *Edit* and make the desired changes.

Name :	<p>Freely definable name to identify the medium</p> <p>If you want to perform a <a href="#">backup [Page 25]</a> or a <a href="#">restore [Page 34]</a> using one of the external backup tools ADSM, NetWorker, Backint for Oracle, or Backint for SAP DB, begin the name of the medium with ADSM, NSR, BACK, or BACK (<a href="#">Backing Up with External Backup Tools [Page 33]</a>).</p>
Location:	<p>Specify the location to which backups are to be made.</p> <p>With files and pipes, we recommend specifying the absolute path to the backup medium.</p>
Device Type:	<p>Select the type of the backup medium: File, tape, or pipe.</p> <p>If you are backing up to tape, you need to address the device driver correctly under the UNIX operating system. Use rewind tapes for backups. On Microsoft Windows NT/Windows 2000 operating systems, the system sets the device driver automatically.</p> <p>If you back up the log to a file, the Database Manager GUI automatically generates version files. An automatic <a href="#">log backup</a> cannot be made to tape.</p> <p>If you are saving with external backup tools, enter <b>Pipe</b> as the Device Type.</p>
Backup Type:	<p>Select the type of backup you want (full or incremental <a href="#">data backup</a>, interactive or automatic <a href="#">log backup</a>).</p>

Select *Save* to save your entries.

## Change the Properties of a Group of Parallel Backup Media

Open the desired group of parallel backup media by double clicking it.

Choose the backup medium to be changed and proceed as described above.



## Deleting Backup Media

### Procedure

1. Choose *Instance* → *Configuration* → *Backup Media*.  
The system displays a list of the [backup media](#) available.
2. Select the backup medium that you wish to delete.
3. Select *BackupMedia* → *Delete* and then *YES* to delete the backup medium.



## Backup Processes

Depending on the technical requirements, the following processes are available for backing up data:

- [Saving to a Single Backup Medium \[Page 30\]](#)
- [Saving to a Group of Parallel Backup Media \[Page 31\]](#)

If the capacity of the selected backup medium is sufficient for the total quantity of data that needs to be backed up, then the database administrator does not need to intervene during the backup process.

If the capacity of the selected backup medium is not sufficient for the backup, succeeding media must be used, meaning that when one backup medium is full, it must be automatically or manually replaced by an empty one:

- [Backups with Automatic Tape Loaders \[Page 32\]](#)
- [Backups with Manually Changed Media \[Page 33\]](#)

The Database Manager GUI requests additional backup media when a medium has been completely filled but the backup has not yet been completed. To allow it to do this, the media size must be stated when defining the medium, or the tape device must detect the end of the tape.

You also have the option of using an external backup tool:

- [Backing Up with External Backup Tools \[Page 33\]](#)

There is an automatic backup function for backing up the log entries:

- [Activating and Deactivating the Automatic Log Backup \[Page 29\]](#)



## Activating and Deactivating the Automatic Log Backup

### Use

[Automatic log backup](#) is used in production systems to ensure data security.

### Procedure

1. In the menu bar of the Database Manager GUI, choose *Instance* → *Backup* → *AutoLog on/off*.  
A list of all of the [backup media](#) available for log backups is displayed. If you have already used the automatic log backup function, the medium used is already selected. You can choose a different medium or create a new one ([Defining a Single Backup Medium \[Page 25\]](#)).
2. Choose *AutoLog* → *Set Auto Log on* to activate the automatic log backup. It remains activated until you deactivate it again by choosing *Instance* → *Backup* → *Auto Log on/off* and then *Auto Log* → *Auto Log off*, even if you have restarted the database instance in the meantime ([Starting the Database Instance/Restart \[Page 22\]](#)).



## Backing Up to a Single Backup Medium

### Use

You can back up the data to a single [backup medium](#). If the capacity of the backup medium selected is too small for the backup, succeeding media have to be used, meaning that when one backup medium is full, it must be automatically or manually replaced by an empty medium when prompted by the Database Manager GUI:



If possible, it is always preferable to [back up to a group of parallel backup media \[Page 31\]](#) rather than backing up to a single medium plus succeeding media. Parallel backup of this kind is not however possible for log backups.



We recommend that you check the logical consistency of the volumes before every complete data backup ([Checking the Logical Consistency of the Volumes \[Page 50\]](#)).

### Prerequisites

Only one backup device is available.

You have defined the backup medium to which backup is to take place ([Defining a Single Backup Medium \[Page 25\]](#)).

If you want to perform a backup using an automatic tape loader ([Backups to Automatic Tape Loader \[Page 32\]](#)), or with the help of external backup tools ([Backups with External Backup Tools \[Page 33\]](#)), then you must state this specifically when defining the backup media.



Under normal circumstances, you should backup [database instances](#) in the `ONLINE` [operational state](#).

You can also backup database instances in the `ADMIN` operational state, but only in exceptional circumstances.

### Procedure

1. Choose *Instance* → *Backup* and the required backup type:
  - *Complete Data* ([complete data backup](#))
  - *Incremental* ([incremental data backup](#))
  - *Log* ([interactive log backup](#))
  - *Auto Log on/off* ([automatic log backup](#))

The system displays the list of the media already defined for the selected backup type. If no media have been defined so far, they must be defined now. Select the backup medium to be used for the backup.

2. Check the details of the medium. Change them if required ([Changing Backup Media \[Page 27\]](#)).
3. Choose the menu entry for the selected backup type in the menu bar at the top of the screen → *Next Step*.

Only for database systems up to version 7.3 (inclusive):

If you want to perform a [data backup](#), now choose whether the backup is to be performed for **Recovery** or **Migration**.

**Recovery** : Data backup for restoring the database instance if errors occur

**Migration:** Data backup with checkpoint; that is, to migrate the database instance  
For more information about checkpoints, see the *User Manual: SAP DB 7.3*.

- If you have selected a **data backup or interactive log backup**, choose the menu entry for the selected backup type → *Start*. The backup action now starts.  
If the capacity of the backup medium specified is not large enough for the backup and you are not using an automatic tape loader, the Database Manager GUI asks you to specify and insert a new medium as soon as the previous medium has been filled (see [Backups to Manually Changed Media \[Page 33\]](#)).  
Once the backup procedure has been successfully completed, you see the message Backup completed.

If you want to activate the **automatic log backup**, choose *AutoLog* → *Auto Log on*.



## Backing Up to a Group of Parallel Backup Media

### Use

To ensure maximum data throughput, the [database instance](#) can be backed up to a several different backup media at the same time. You can set the maximum number of tape devices which can be used at the same time using parameter [MAXBACKUPDEVS](#). The use of up to 32 tape devices allows you to reduce backup times considerably.

Parallel backing up is only possible for [data backups](#).



We recommend that you check the logical consistency of the volumes before every complete data backup ([Checking the Logical Consistency of the Volumes \[Page 50\]](#)).

### Prerequisites

There are multiple backup devices available.

You have defined the [backup medium](#) to which backup is to take place ([Defining a Single Backup Medium \[Page 26\]](#)).

If you want to perform a backup using an automatic tape loader ([Backups to Automatic Tape Loader \[Page 32\]](#)), or with the help of external backup tools ([Backups with External Backup Tools \[Page 33\]](#)), then you must state this specifically when defining the media.



Under normal circumstances, you should backup [database instances](#) in the ONLINE [operational status](#).

You can also backup database instances in the ADMIN operational state, but only in exceptional circumstances.

### Procedure

- Choose *Instance* → *Backup* and the selected backup type:  
→ *Complete* ([complete data backup](#))  
→ *Incremental* ([incremental data backup](#))  
**A group of parallel media cannot be used for log backups.**

A list of the backup media already defined for the selected backup type appears on the screen.

If no backup media have been defined yet, you must define them now ([Defining a Group of Parallel Backup Media \[Page 26\]](#)).

2. Select the backup medium to be used for the backup.
3. Check the details of the medium. Change them if required ([Changing Backup Media \[Page 27\]](#)).
4. Choose the menu entry for the selected backup type in the menu bar at the top of the screen → *Next Step*.

Only for database systems up to version 7.3 (inclusive):

If you want to perform a [data backup](#), now choose whether the backup is to be performed for **Recovery** or **Migration**.

**Recovery** : Data backup for restoring the database instance if errors occur

**Migration**: Data backup with checkpoint; that is, to migrate the database instance

For more information about checkpoints, see the *User Manual: SAP DB 7.3*.

5. Choose the menu entry for the selected backup type → *Start*.  
The backup process begins.  
If the capacity of the group of parallel media is not large enough for the backup, Database Manager GUI asks you to specify and insert a new medium as soon as any of the media has been filled (see [Backups to Manually Changed Media \[Page 33\]](#)).

Once the backup procedure has been successfully completed, you see the message `Backup completed`.



## Backups with Automatic Tape Loaders

The Database Manager GUI supports automatic tape loaders. Alternatively, you can backup using an external backup tool ([Backing Up with External Backup Tools \[Page 33\]](#)).

If you want to use an automatic tape loader for [backups \[Page 25\]](#), select *Extended* when defining the medium ([Defining a Single Backup Medium \[Page 25\]](#), [Defining a Group of Parallel Backup Media \[Page 26\]](#) or [Changing Backup Media \[Page 27\]](#)) and select the `Autoloader` option.

- On the Microsoft Windows NT/Windows 2000 operating system, this defines the automatic loading of tapes.
- If you are using a different operating system, under media definition enter the command the operating system requires for automatic tape loading; to do this, select *Extended* and enter the operating system command in the `OS Command` field.

After inserting the tape cartridge in the autoloader, the first tape must be selected manually. When the end of this tape is reached, the autoloader will then take the next tape available.

For a backup to be successful, you will need to ensure an adequate supply of tapes.

At the end of the backup, the tape device shows the number of tapes written. The tapes should be marked with the [backup ID](#) that was displayed and confirmed at the start of the backup.



Before a backup you must check that none of the tapes is write-protected because this will terminate the entire backup process.



## Backing Up with External Backup Tools

### Use

You can use [external backup tools](#) to back up to tapes and succeeding tapes. The Database Manager GUI currently supports the following backup tools:

- ADSM/TSM (IBM/Tivoli)
- Backint for Oracle
- Backint for SAP DB
- NetWorker (Legato)



Contact [SAP DB Support](#) if you would like to use external backup tools that are not included in this list.

### Prerequisites

The [database instance](#) is in the `ONLINE` or the `ADMIN` [operational state](#).

You have created a [single backup medium](#) or a [group of parallel backup media](#), taking the naming conventions into account ([Name of an External Backup Medium](#)) with the `Device Type` `Pipe` and specified the operating system path for this pipe under `Location`. ([Defining a Single Backup Medium \[Page 25\]](#), [Defining a Group of Parallel Backup Media \[Page 26\]](#))

### Procedure

Proceed as described under [Backing Up to a Single Backup Medium \[Page 30\]](#) or [Backing Up to a Group of Parallel Backup Media \[Page 31\]](#).



## Backups with Manually Changed Media

### Prerequisites

You have started a [backup to a single backup medium \[Page 30\]](#) or a [backup to a group of parallel backup media \[Page 31\]](#).

The capacity of the [backup medium](#) or backup media inserted is not sufficient for the selected backup. The advancing bar showing the progress of the backup stops before reaching 100%.

You are prompted to insert the succeeding medium.

### Procedure

#### Succeeding Medium is Tape

1. Insert the next tape.
2. Ensure that the tape that has just been written is stored in a safe place and that the right tape has been inserted.

3. Choose *Start*.



We recommend you specify the exact capacity of the medium when defining the medium. Provided the tape device driver gives a reliable indication when the end of the tape is reached, the Database Manager GUI reacts as described even when `Size = 0`.

If the end of the tape is not recognized correctly, the message `Writing Error` is displayed. Even when this happens you can continue as described in steps 1-3.

## Succeeding Medium is a File

1. Under `New Location`, enter the name and full path of the file to which the backup should be made.
2. Choose the menu entry for the selected backup type → *Continue*.

Repeat the procedure appropriate to the media type until the backup process has been completed.

As long as the backup process is being performed, no other backup can be started. A backup process is complete when the message `Backup completed` appears.



## Restoring Data

### Use

If a hardware error occurs, the Database Manager supports the recovery of the [database instance](#) by importing [backups](#). You can recover either the last database state or an earlier one.



If the automatic log backup function was not activated before the recovery of the database instance started, this function is **not** automatically reactivated after recovery has taken place.

To reactivate the log backup function, execute the command for [activating the automatic log backup function \[Page 29\]](#).

## Prerequisites

All the [backup history](#) data is available.

## Procedure

To obtain the highest possible throughput of data, complete and incremental [data backups](#) can be retrieved from a number of [backup media](#) simultaneously. The number of parallel backup media in this case does not depend on the number of parallel backup media used to make the backup originally. Even a backup made to single medium plus succeeding media can be restored in parallel.

You can set the maximum number of backup devices from which it is possible to restore simultaneously with the database parameter [MAXBACKUPDEVS](#). The use of up to 32 tape devices allows you to reduce restore times considerably.

The Database Manager GUI always suggests the quickest way of effecting a restore. Complete data backup(s) are restored first in this case and these are followed firstly by incremental backups, if required, and then by [log backups](#).



Log backups cannot be read in parallel.

## Restore Process

[Continuing an Interrupted Restore \[Page 35\]](#)

[Restoring a Backup from the Backup History \[Page 36\]](#)

[Restoring the Indexes After a Database Restore \[Page 37\]](#)

[Restoring the Last Complete Data Backup \[Page 37\]](#)

[Restoring an Existing Database Instance \[Page 38\]](#)

[Restoring a Mirrored Volume \[Page 39\]](#)

[Restoring with an Automatic Tape Loader \[Page 40\]](#)

[Restoring with External Backup Tools \[Page 40\]](#)

[Restoring Without a Backup History \[Page 41\]](#)



## Continuing an Interrupted Restore

### Prerequisites

The [database instance](#) is in [operational state](#) ADMIN.

One of several possible restore procedures was started but was canceled before completion.

- [Restoring the Last Complete Data Backup \[Page 37\]](#)
- [Restoring a Backup from the Backup History \[Page 36\]](#)
- [Restoring Without a Backup History \[Page 41\]](#)
- [Restoring an Existing Database Instance \[Page 38\]](#)

### Procedure

1. Choose *Instance* → *Recovery* → *Database*.
2. Choose the option  
*Continue restoring pages/log*
3. Choose *Database* → *Next Step*.  
The system displays a list of all relevant [log backups](#) for [data backups](#) that have already been restored.
4. Choose *Database* → *Replace* to continue the restore operation. You can follow the progress of the restore operation and its completion on the [screen \[Page 54\]](#) display.
5. *Close* closes the restore operation.



## Restoring a Backup from the Backup History

### Prerequisites

The [log files of the Database Manager](#) `dbm.knl` and `dbm.mdf` in the [run directory](#) of the [database instance](#) have not been corrupted, and the [backup history](#) is therefore complete. If you are performing a [restore using external backup tools \[Page 40\]](#), the log file `dbm.ebf` must also have remained uncorrupted.

The database instance is in the ADMIN [operational state](#).

### Procedure

1. Choose *Instance* → *Recovery* → *Database*.
2. Select the option  
*Restore a specified backup from history.*  
If you select the option  
*Restore until a specific time*  
on the same screen, you can define a time to which the system should perform the restore. The system default is the current time. If you do not change this system default, the backups are restored completely and with all the changes they contain.
3. Choose *Database* → *Next Step*.  
The system displays the backup history of [complete data backups](#).
4. Select the desired complete data backup.  
Choose *Database* → *Next Step*.  
The system displays the selected complete data backup and the associated incremental data and [log backups](#). You can deselect in the list any [incremental data backups](#) that are not available or that have been damaged. The Database Manager GUI then shows the appropriate log backups in place of the page backups.
5. Choose *Database* → *Next Step*. The system prompts you to insert the specified [backup medium](#).
6. Choose *Database* → *Start*.  
If you are performing a restore using external backup tools, the [external backup IDs](#) are displayed at this point.  
The complete data backup is restored.  
If you then need to restore any incremental data backups, the Database Manager prompts you to confirm each one.  
If you then need to make log backups, the Database Manager prompts you to confirm that you want to restore all displayed backups.

Once the [restore \[Page 34\]](#) is complete, the database instance is in the ONLINE operational state.



When the consecutive backups are restored, the Database Manager GUI prompts you to insert the next backup medium each time, unless you are [restoring from an automatic tape loader \[Page 40\]](#).

You can if you wish stop the restore process and resume it later on (see [Continuing an Interrupted Restore \[Page 35\]](#)).



## Restoring the Indexes After a Database Restore

### Use

During a [restore \[Page 34\]](#) of a [database instance](#), the indexes contained in the [log area](#) are not restored for optimization reasons. Indexes that have not been restored are marked `BAD` and can be restored explicitly by following the procedure described here.

### Prerequisites

The database instance is in the `ONLINE` [operational state](#).

### Procedure

Choose *Instance* → *Recovery* → *Indexes*.

You can choose the tables for which the indexes marked `BAD` are to be displayed. To do so, enter appropriate search arguments under `Owner`, `Table Name`, and `Index Name`.

<code>Owner</code>	<code>Owner</code>
<code>Table Name</code>	<code>Name of the table</code>
<code>Index Name</code>	<code>Name of the index</code>

Then choose *Indexes* → *Select*. The system displays a list of the indexes that match your criteria.

Select the indexes you want to restore in the list. If you want to restore all of the indexes displayed, choose *Indexes* → *Mark All*.

Choose *Indexes* → *Execute* to perform the action.



It is advisable to restore the indexes during periods where the workload on the system is lower.



## Restoring the Last Complete Data Backup

### Prerequisites

The [log files of the Database Manager](#) `dbm.knl` and `dbm.mdf` in the [run directory](#) of the [database instance](#) have not been corrupted, and the [backup history](#) is therefore complete. If you are performing a [restore using external backup tools \[Page 40\]](#), the log file `dbm.ebf` must also have remained uncorrupted.

The database instance is in the `ADMIN` [operational state](#).

### Procedure

1. Choose *Instance* → *Recovery* → *Database*.
2. Select the option  
*Restore last backup*.  
 If you select the option  
*Restore until a specific time*  
 on the same screen, you can define a time to which the system should perform the restore. The system default is the current time. If you do not change the time proposed

by the system, the backups are restored completely and with all the changes they contain.

3. Choose *Database* → *Next Step*.  
The system displays the last complete [data backup](#) in accordance with the time you specified earlier; that is including any incremental data backups and [log backups](#) if required. You can deselect in the list any incremental data backups that are not available or that have been damaged. The Database Manager GUI then shows the appropriate log backups in place of the page backups.
4. Choose *Database* → *Next Step*. The system prompts you to insert the specified [backup medium](#).
5. Choose *Database* → *Start*.  
If you are performing a restore using external backup tools, the [external backup IDs](#) are displayed at this point.  
The complete data backup is restored.  
If you then need to restore any incremental data backups, the Database Manager GUI prompts you to confirm each one.  
If you then need to make log backups, the Database Manager GUI prompts you to confirm that you want to restore all displayed backups.  
Choose *Database* → *Restart* to restart the database instance.



When the consecutive backups are restored, the Database Manager prompts you to insert the next backup medium each time, unless you are [restoring from an automatic tape loader \[Page 40\]](#).

You can if you wish stop the restore process and resume it later on (see [Continuing an Interrupted Restore \[Page 35\]](#)).



## Restoring an Existing Database Instance

### Prerequisites

The [log files of the Database Manager](#) `dbm.knl` and `dbm.mdf` in the [run directory](#) of the [database instance](#) have not been corrupted, and the [backup history](#) is therefore complete.

The database instance is in the ADMIN [operational state](#).

### Procedure

Select *Instance* → *Install* to launch the Database Wizard. This Database Wizard guides you through the entire procedure for restoring an existing database instance.

1. Specify under which name the existing database instance is registered in the Database Manager GUI.
2. Choose *Next*.  
The Database Manager GUI displays a message saying that a database instance with this name already exists.
6. Choose *Reinstall*. This particular database instance is then restored. The Database Manager uses the database parameters of the old version of the database instance as the default values for the new version. **All data from the old version is lost.**
7. Enter the **first Database Manager operator** and the first database user ([SYSDBA](#)) for this database instance, enter their passwords, and choose *Next*.  
The Database Manager GUI provides the `Use current parameters` option for creating the initial parameter file.

8. Choose *Next*.  
The system automatically proposes the parameter values from the old version of the database instance. You can adjust these values to your requirements. Select the required parameter and then choose *Edit*.  
It now appears in the bottom part of the window. Next to the parameter you will see an explanation of what the parameter does and, in some cases, how it is calculated.  
  
Enter the new value for the parameter in the `New Value` field.  
Save your entries.  
The new value appears in the *New Value* column. It is stored in the internal data structures and becomes effective once the [database instance is started/restarted \[Page 22\]](#).
9. Choose *Next*.  
When you exit the input screen, the parameters are checked following the rules stored on the [DBM Server](#). You may be asked to make and confirm changes before you can leave the input screen.  
For the [log volumes](#) and the [data volumes](#), the system automatically proposes the configuration of the old version of the database instance. You can adjust these values to your requirements. Bear in mind the parameters set in the previous step. Select a volume and then choose *Edit*. Enter the size and ID or the absolute path of the volume and confirm your entries. Repeat this process for each volume.
10. Choose *Next* and *Restore instance*.
3. Choose *Install*.  
A new database instance is installed with the parameters set in step 5 using a complete [data backup](#).  
The system prompts you to give the database instance a new registration name in the Database Manager GUI. This name must be unique within the Database Manager.
4. Confirm your entries.
5. Select *Close* to close the Database Wizard.  
This takes you back to the initial screen of the Database Manager GUI ([Screen \[Page 54\]](#)).

Once the restore procedure is complete, the database instance is in the `ONLINE` operational state.



When the consecutive backups are restored, the Database Manager GUI explicitly prompts you to insert the next [backup medium](#) each time, unless you are [restoring from an automatic tape loader \[Page 40\]](#).

You can if you wish stop the restore process and resume it later on (see [Continuing an Interrupted Restore \[Page 35\]](#)).



## Restoring a Mirrored Volume

### Use

If the [volumes](#) of your [database instance](#) are mirrored, one of the volume mirrors may fail. In this case, the database continues to operate but only the part of the volume pair that is intact can be accessed. The faulty volume mirror is displayed as a `Bad` volume.

You use this function to recover a damaged mirror [log volume](#).



To ensure the downtime security of your system, it is advisable to mirror log volumes.

We recommend RAID 5 configurations to protect your data volumes.

## Recovering a damaged log volume mirror

If you trigger the recovery of a log volume, it is only reinitialized. In other words, the data is not copied from the intact volume to the recovered volume. The system then writes to both log volumes sequentially. The volume originally marked `BAD` must be filled once before the contents of both volumes are identical again and both volumes can be read.

The defective volume remains marked `BAD` during the entire reintegration phase.

### Prerequisites

The database instance is the `ONLINE` or the `ADMIN` operational state.

### Procedure

1. Choose *Instance* → *Recovery* → *Volumes*.  
Select the `BAD` volumes that are to be recovered.
2. Choose *Volumes* → *Execute* to start initializing the mirror volume.



## Restoring from Automatic Tape Loaders

### Prerequisites

For the retrieval operation to be successful, all the tapes which form part of the backup must be present in the automatic tape loader.

### Procedure

Use the same media definition as for backing up ([Defining a Single Backup Medium \[Page 25\]](#)). Follow the individual steps as described for the relevant recovery procedure:

- [Restoring the Last Complete Data Backup \[Page 37\]](#)
- [Restoring a Backup from the Backup History \[Page 36\]](#)
- [Restoring an Existing Database Instance \[Page 38\]](#)
- [Restoring Without a Backup History \[Page 41\]](#)



## Restoring with External Backup Tools

The Database Manager GUI supports the use of the following [external backup tools](#) for restores:

- ADISM/TSM (IBM/Tivoli)
- Backint for Oracle
- Backint SAP DB
- NetWorker (Legato)



If you want to use an external backup tool that is not included on this list, contact [SAP DB Support](#).

## Using the backup tools

Start the restore operation from the Database Manager GUI. The naming conventions for [external backup media](#) enable the program to recognize the external backup tool and start it.

## Procedure

When you use an external backup tool to restore backups, the [external backup IDs](#) of the required backups are automatically defined by the Database Manager GUI, displayed so that the operator can check them, and used during the restore process. The operator does not need to explicitly determine or specify the external backup IDs.



The number of media in a [group of parallel backup media](#) must correspond to the number of media used to create the backup.

Follow the procedure for the relevant restore process:

- [Restoring the Last Complete Data Backup \[Page 37\]](#)
- [Restoring a Backup from the Backup History \[Page 36\]](#)
- [Restoring Without a Backup History \[Page 41\]](#)



## Restoring Without a Backup History

### Prerequisites

The [database instance](#) is in [operational state](#) ADMIN.

The [backup history](#) contains errors or is no longer complete.

### Procedure

1. Choose *Instance* → *Recovery* → *Database*.
2. Choose the option  
*Restore a medium*.
3. If you select the option  
*Restore until a specific time*  
on the same screen, you can define a time up to which a [backup medium](#) is to be read back in. The system default is the current time. If you do not change the time proposed by the system, the backups are restored completely and with all the changes they contain.
4. Choose *Database* → *Next Step*.  
A list of all of the available backup media is displayed.
5. Select the backup medium you want to restore. If you are performing a [restore with external backup tools \[Page 40\]](#), the system displays a list of backups.
6. Choose *Database* → *Next Step*. Select the desired [external backup IDs](#).
7. Choose *Database* → *Next Step*.  
The system prompts you to insert the specified backup medium. If you are performing a

restore using external backup tools, the external backup IDs are displayed again at this point.

8. Choose *Database* → *Start* to start the recovery process.  
If you then want to restore other backup media without a backup history, proceed as outlined as of step 2 onward.
9. Choose *Database* → *Restart* to restart the database instance.



You can if you wish stop the restore process and resume it later on (see [Continuing an Interrupted Restore \[Page 35\]](#)).



## Displaying Information

To get a complete overview of all information and statistics for important database functions, select the desired [database instance](#), and choose *Instance* → *View Details*.

If you want to view detailed information about a particular area, choose *Instance* → *Information* and then the desired area. Information is available for the following topics:

[Displaying User Sessions \[Page 42\]](#)

[Displaying Cache Information \[Page 43\]](#)

[Displaying Data Area Information \[Page 43\]](#)

[Displaying Read and Write Operations \[Page 44\]](#)

[Displaying Log Area and Log Writer Information \[Page 44\]](#)

[Displaying the Backup History \[Page 45\]](#)

[Displaying Locks and Lock Requests \[Page 46\]](#)

[Displaying Versions \[Page 47\]](#)



## Displaying User Sessions

### Use

This menu function shows all database users who are currently connected to a database.

### Procedure

Choose *Instance* → *Information* → *Sessions*.

The following information is displayed:

Name	Name of the logged-on user
Terminal ID	A terminal identification specific to the particular operating system.
Task ID	Task number
Session ID	Number of the user session
Node	Name of the Database Manager server in the network
Catalog Cache Size	Size of <a href="#">catalog cache</a> available for user



## Displaying Cache Information

### Use

The hit rates for the [caches](#) show whether they were configured to be sufficiently large. The hit rates for the data cache and the converter cache are particularly important.

### Procedure

Choose *Instance* → *Information* → *Caches*.

Type	Cache type
Accesses	Total number of accesses
Successful	Number of successful accesses
Unsuccessful	Number of unsuccessful accesses
Hit Rate (%)	Hit rate (number of hits as a percentage of the total number of accesses)

The hit rate of the data cache should be at least 99% and the converter cache hit rate should be at least 98%.

Lower hit rates can occur for a short time, for example when tables are read for the first time, or when repeated table scans are being performed and the table does not fit into 10% of the data cache.

If the values are permanently lower than the recommended percentage, this is a sign of poor performance and the settings of the database parameters DATA\_CACHE and/or CONVERTER\_CACHE must be increased.



## Displaying Data Area Information

### Use

You can display information about the size and current utilization of the [data area](#).

### Procedure

Choose *Instance* → *Information* → *Data Area*.

Name	Name of <a href="#">data volume</a>
Value	Current value of the data volume

The following information is displayed:

Total space (KB)	Total size of the data volume in KB
Max. persistent space (KB)	Maximum permanent data area in KB
Used space (KB)	Used data area in KB
Used space (%)	Used data area as a percentage
Free space (KB)	Free data area in KB
Free space (%)	Free data area as a percentage

Updated perm pages

Updated permanent pages



## Displaying Read and Write Operations

### Use

The database server keeps statistics about the current number of the physical and logical read and write accesses since the [database instance was last started \[Page 22\]](#).

### Procedure

Choose *Instance* → *Information* → *I/O*.

Name	Name
Logical Reads	Number of logical read requests
Logical Writes	Number of logical write requests
Physical Reads	Number of physical read requests
Physical Writes	Number of physical write requests

The details are displayed for the following data areas:

Catalog	Data area occupied by the data dictionary of the database system. Frequent writing to this area is a sign of modifications made to the database design.
Permanent Data	The actual data area for permanent data.
Temporary Data	Data area on the hard disk that is used temporarily, that is required, for example, when building selected data sets.
Long Data	LONG columns
Leaf, Level 1, Level 2, Level 3	Data tree structure The database organizes the stored data as <a href="#">B* trees</a> .

SUM shows the total number of accesses.



## Displaying Log Area and Log Writer Information

### Use

You are provided with information about the state and size of the [log area](#) and the activities of the Log Writer.

### Procedure

Choose *Instance* → *Information* → *Log*.

Log mode	<a href="#">Mode</a> in which the log is operated
Max. size	Maximum available storage space in the log area
Backup Segment size	Size of a segment

Used size (KB)	Used space in KB
Used size in %	Used space as percentage
Not saved (KB)	Size of log area not yet saved, in KB
Not saved (%)	Size of log area not yet saved as a percentage
Log since last data backup (KB)	Number of log pages written since the last <a href="#">data backup</a>
Savepoints written	Number of <a href="#">savepoints</a> written since last <a href="#">log backup</a>
Checkpoints written	Number of checkpoints written (Only for SAP DB versions < 7.4)
Physical reads	Read activities in the log area since the last <a href="#">restart [Page 22]</a> of the database system
Physical writes	Write activities in the log area since the last restart of the database system
Queue size	Current size of the log queue
Queues allocated	Current log queue allocation
Queue entries	Maximum size of the log queue
Queue overflows	Number of log queue overflows
Group commits	Number of group commits
Waits for logwriter	Number of wait states for log write operations
Max. waits	Maximum number of wait states per log page
Avg. Waits	Average number of wait states per log page
OMS Log used	Log area used by the Object Management System (OMS)
OMS min. free pages	Minimum number of free pages for the Object Management System (OMS)



## Displaying the Backup History

### Use

The system displays a chronological list of all backup and restore operations that have taken place ([Backup History](#)).

### Procedure

Choose *Instance* → *Information* → *Backup History*

You will see the following items of information for each action:

Label	<a href="#">Backup ID</a>
Action	The action (backup or recover) and the <a href="#">operational state</a> in which it took place.
Beginning	Time when the action began
End	Time when the action ended

Result	Result
Media Name	Name of the <a href="#">backup medium</a>
Log Required	<a href="#">Log backups</a> required
Size (KB)	Size (KB)
Size (Pages)	Size (pages)
Volumes	Volumes
Next Log Page	Log page
From Page	From page number
To Page	To page number

For additional information, choose *BackupHistory* and one of the following functions:

- → *Show details*  
The system displays media information. If you are using [external backup tools](#), the system displays the [external backup IDs](#) and the availability of the backup media.
- → *Columns*  
You can create your own personal display variant of the backup history here.
- → *Refresh*  
You update the display of the backup history.
- → *Refresh with External Backup Info*  
You update the backup history and the information about the backups that were performed with external backup tools.



## Displaying Locks and Lock Requests

### Use

You are given information about the current locks and lock requests in the database server.

### Procedure

Choose *Instance* → *Information* → *Locks*.

The following information is displayed:

Max Entries	Maximum number of locks held and requested
Average Used Entries	Average number of locks held and requested
Collisions	Number of collisions that occurred
Escalations	Number of escalations that occurred
Row Locks	Row locks currently held
Table Locks	Table locks currently held
Request Timeout	This indicates the time a task may wait for the setting of a requested lock before it is canceled.



## Displaying Versions

### Use

The system displays the current versions of the database software and the runtime environment.

### Procedure

Select *Instance* → *Information* → *Versions*.



To display the current version of the Database Manager, choose *Help* → *About Database Manager*.



## Options for Diagnosing Problems

In the event of errors or performance problems, the Database Manager GUI offers you several different options for diagnosing the problem.

The following functions are available:

[Reading Database Files \[Page 52\]](#)

[Activating the Database Trace Function \[Page 47\]](#)

[Activating the Event Monitor \[Page 49\]](#)

[Using the Command Line Version of the Program \[Page 49\]](#)

[Checking the Logical Consistency of the Volumes \[Page 50\]](#)

[Displaying OMS Caches \[Page 50\]](#)

[Displaying OMS Class Containers \[Page 50\]](#)

[Activating the OMS Monitor \[Page 51\]](#)

[Checking the Physical Consistency of a Backup \[Page 51\]](#)

[Checking the Server of a Database Instance \[Page 52\]](#)



## Activating the Database Trace Function

### Use

The Database Manager gives you the option of enabling the database trace function if errors occur.



Use this function when asked to do so by your [SAP DB Support](#) organization. We would not recommend that you activate the database trace function unless it is really necessary, because it seriously affects the performance of the [database instance](#).

Disable the database trace function as soon as the actions needed for the analysis have been logged.

The database trace documents all reactions of the database kernel to database statements. This also means that it can be used not only to trace errors that occur when statements are processed but also to provide a more exact classification of inconsistencies caused, for example, by hardware errors. If the database trace is enabled, you should be running the database instance with the smallest load possible, and only those actions needed to reproduce the error should be performed on the database instance.

## Procedure

You can execute the log functions below either individually or in a combination that you can define yourself.

### Logging Database Statements in General

Choose *Instance* → *Check* → *Kernel Trace*.

On the *Options* tab page, select the required database trace options in the list. If the options support the selection of a trace level, the system default value is displayed in the column *Level*.

If you then choose *KernelTrace* → *Selected Items ON*, you activate the database trace for the selected database trace options with the system default value for the trace level.

If you want to change the trace level for the selected database trace options, choose *KernelTrace* → *Selected Items ON with Level...* You are now able to choose a different trace level. The higher the trace level, the greater amount of detail is provided in the information collected by the database trace.

Choose *KernelTrace* → *Selected Items OFF* to deactivate the database trace.

### Logging Database Statements in a Specific User Session

Choose *Instance* → *Check* → *Kernel Trace*.

The database trace is written for all user sessions by default. The default setting can be changed:

To do this, choose the *Advanced* tab page and select *Trace session*. Then choose whether the database trace is to be written for the current session or for a specific user session. In this case, enter the session ID of the desired user session.



To locate the session ID of a user session, choose *Instance* → *Information* → *Sessions*.

Choose *KernelTrace* → *Set Advanced Options*.

### Logging the Occurrence of Specific Errors

Choose *Instance* → *Check* → *Kernel Trace*.

On the *Advanced* tab page, select *Stop on error*. Enter the code of the error whose occurrence is to stop the creation of the database trace.

Then choose *KernelTrace* → *Flush* to write the information still contained in the buffer to the trace and therefore update the database trace before the log file is created.

### Creating and Displaying the Database Trace Log

Choose *Instance* → *Check* → *Kernel Trace*.

On the *Protocol* tab page, select the required options.

Then choose *Kernel Trace* → *Flush* to write the information still contained in the buffer to the trace and thus update the database trace before the log file is created.

Choose *KernelTrace* → *Make Protocol*. The database trace is then converted to a legible format. To open the log, choose *Instance* → *Check* → *Diagnosis Files*. Select the `Kernel Trace Protocol` file and choose *DiagnosisFiles* → *View*.



## Activating the Event Monitor

### Use

The database manager is equipped with a mechanism to log events. A series of preconfigured events enables you to monitor the most important operating data, such as the data and log fill level, and to display it continuously.

### Procedure

1. First choose *File* → *Properties* and select the *Event* field on the *General* tab page.
2. Then choose *Instance* → *Configuration* → *Parameters*
3. Choose *Parameters* → *Show Extended and Support Parameters*.  
On the *Extended* tab page, change the `EVENT_ALIVE_CYCLE` parameter from 0 to 10 (see [Displaying and Changing Current Database Parameters \[Page 18\]](#)).  
On the *Support* tab page, change the `_EVENTSIZE` parameter from 0 to 100 and `_MAXEVENTTASKS` from 0 to 2.
4. Save your changes.

### Displaying the List of Events for a Database Instance

1. Select the required [database instance](#) in the list of registered database instances.
2. Choose *Instance* → *Check* → *Event Monitor*.  
A list is then displayed with the following information:

Name	Priority
Name of the Event	low (Information)   medium (Warning)   high (Problem)



## Using the Command Line Version of the Program

In the Database Manager GUI program, you can also use the command line version of the program called the Database Manager CLI.

Choose *Instance* → *Command Line* .

Enter the desired [DBM commands](#) in the `Command` field and select the menu entry *Command* → *Execute*.

**See also:**

[Database Manager CLI: SAP DB 7.4](#)



## Checking the Logical Consistency of the Volumes

### Use

You check the consistency of the internal data structures of the current [database instance](#). If there are serious inconsistencies, the database instance must be recovered in the same way as after a disk failure ([Recovery \[Page 34\]](#)).

If you execute this command in the ADMIN [operational state](#), free storage pages wrongly recorded as used following an abnormal termination of database operation are released to the free space management.



We recommend that you execute this function before each [full data backup](#).

### Prerequisites

The database instance is the `ONLINE` or the `ADMIN` operational state.

### Procedure

Choose *Instance* → *Check* → *Database Structures* and then *DatabaseStructures* → *Execute* to start the logical consistency check.

### Result

The information supplied by the database kernel is written to the `Database Errors` file. To view the contents of the file, choose *Instance* → *Check* → *Diagnosis Files*. Select the `Database Errors` file and choose *DiagnosisFiles* → *View*.



## Displaying OMS Caches

### Use

This function supports the [liveCache](#) performance analysis.

### Prerequisites

The current [database instance](#) has the [instance type](#) `liveCache`. You are in the [operational state](#) `ONLINE`.

### Procedure

Choose *Instance* → *Information* → *OMS Caches* to display the OMS caches.



## Displaying OMS Class Containers

### Use

This function supports the [liveCache](#) performance analysis.

## Prerequisites

The current [database instance](#) has the [instance type](#) liveCache. You are in the [operational state](#) ONLINE.

## Procedure

Choose *Instance* → *Information* → *OMS Class Containers* to display the information about the class containers.



## Activating the OMS Monitor

### Use

This function supports the [liveCache](#) performance analysis.

### Prerequisites

The current [database instance](#) has the [instance type](#) type liveCache. You are in the [operational state](#) ONLINE.

### Procedure

Choose *Instance* → *Information* → *OMS Monitor* to activate the OMS Monitor.



## Checking the Physical Consistency of a Backup

### Use

You check the backup on the selected [backup medium](#).



This procedure, however, only checks the physical consistency of the imported backup. The media definition is checked to determine whether it can be imported, whether the backup is imported entirely (correct number of pages), and whether the [pages](#) are consistent.

The [logical consistency \[Page 50\]](#) is **not** checked.

### Prerequisites

The [database instance](#) can be in any [operational state](#).

### Procedure

1. Choose *Instance* → *Check* → *Backup*.
2. Select the medium on which the backup is stored.
3. Choose *CheckMedium* → *Next Step* and then *CheckMedium* → *Start* to start the physical consistency check.

### Result

The information supplied by the database kernel is displayed in the message output area.



## Reading Database Files

### Use

The SAP DB database system has a number of [database files](#). These include the [files of the Database Manager](#), which are written by the Database Manager.

You can use this command to access the existing database files.

### Procedure

Choose *Instance* → *Check* → *Diagnosis Files*.



## Checking the Server of a Database Instance

### Use

The Database Manager GUI provides you with an overview of the operating system resources that the database system is using, how the database sessions have been distributed among the operating system processes/threads, and the status of the active [database sessions](#). There are additional functions available, but these should primarily be reserved for support employees and developers.

### Prerequisites

The [database instance](#) is in `ONLINE` or `ADMIN` [operational state](#).

### Procedure

1. Choose *Instance* → *Check* → *Database Server*.
2. Select the required information type and then choose *DatabaseServer* → *View* to display detailed information.



## Installing an Update of the Database Software

### Prerequisites

You have imported the relevant patch for the database software update.

### Procedure

1. Choose *Instance* → *Online* to start the [database instance](#).  
The first time the Database Manager is started after there has been a change in the version of the database software, it automatically adjusts the database instance kernel parameters to the new version of the database software.
2. To update the system tables, choose *Instance* → *Configuration* → *Upgrade System Tables* (See also: [Updating the System Tables \[Page 53\]](#)).

### Result

The Database Manager creates a copy of the parameter file each time a parameter is changed. The serial number extension indicates the version, whereby 01 is the latest version:

UNIX: /usr/spool/sql/config/<database\_name>.<running\_number>

Microsoft Windows NT/Windows 95:

%DBROOT%\CONFIG\<database\_name>.<running\_number>



Should errors occur while loading the parameters, the operation can be undone by restoring the original parameter file. To do so, use the [Database Manager CLI](#) program to copy the file <database\_name>.<running\_number> (copy of parameter file) back to <database\_name> (active parameter file).

If an error occurs when the [database instance is restarted \[Page 22\]](#) again after the update of the software version, contact den [SAP DB Support](#).



## Updating the System Tables

### Use

You can use the Database Manager GUI to update the system tables when you upgrade the database software release.

### Prerequisites

The database instance is in the ONLINE [operational state](#).

### Procedure

1. Choose *Instance* → *Configuration* → *Upgrade System Tables*.
2. Specify the [DBA](#) user and the [DOMAIN](#) user if they are not logged on to the Database Manager, and then choose *Start* to load the system tables.

### Result

The system tables are loaded and updated. At the end of the process, the log is displayed.



If an error occurs, the `ERROR` status appears. You can then look at the installation log with the file name `dbm.ins` for more information on the error ([Reading Database Files \[Page 52\]](#)).



## Accessing the SQL Studio Program

### Use

You can call the SQL Studio program from the Database Manager GUI and access data of a SAP DB [database instance](#) with it.

### Prerequisites

You have the authorization of a [database user](#) of this database instance.

## Procedure

Choose *Tools* → *SQL Studio*.

**See also:**

[SQL Studio: SAP DB 7.4](#)



## Screen Areas of the Database Manager GUI

The screen of the Database Manager GUI is structure into nine areas. The size of these screen areas can be changed.

1		
2		
3	4	
5	6	
7		
	8	9

### Contents of the Screen Areas

Number in the Overview	Name of the Window	Description
1.	Menu Bar	Selection of functions and displays for the current <a href="#">database instance</a> . The menu bar dynamically adjusts to the selected functions.
2.	Icon Bar	The icon bar contains icons you can use to call up functions shown on the menu bar for the current database instance.
3.	Directory structure of the registered database instances	For organizing the registered database instances as in Microsoft Windows Explorer  You can display or hide this structure. To do this, choose <i>View</i> → <i>Tree</i> and select or deselect <i>Tree</i> .
4.	List of Registered Database Instances	Database instances registered in the Database Manager and their <a href="#">operational state</a> . You can obtain an overview of the statistics and important operational values by double clicking a database

		instance
5.	Menu List of the Current Database Instance	You can display or hide this list. To do this, choose <i>View</i> → <i>Menu</i> and select or deselect <i>Menu</i> .
6.	Output Area	Central screen area The name of the current database instance, the input screens for functions, and the output values for displays are displayed here.
7.	Message display	Display of general and error messages
8.	Displaying the Registration Name	Name for the registration of the current database instance in the Database Manager GUI and current operational state
9.	Display of the database name	Name of the current database instance and the database server