



**LibreOffice**  
The Document Foundation

Base Handbook

*Chapter 10*  
*Database Maintenance*

## Copyright

---

This document is Copyright © 2013–2015 by the LibreOffice Documentation Team. Contributors are listed below. You may distribute or modify it under the terms of either the GNU General Public License (<http://www.gnu.org/licenses/gpl.html>), version 3 or later, or the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), version 4.0 or later.

All trademarks within this guide belong to their legitimate owners.

## Contributors

Jochen Schiffers  
Hazel Russman

Robert Großkopf  
Jean Hollis Weber

Jost Lange

## Feedback

Please direct any comments or suggestions about this document to the Documentation Team's mailing list: [documentation@global.libreoffice.org](mailto:documentation@global.libreoffice.org)

Note: Everything you send to a mailing list, including your email address and any other personal information that is written in the message, is publicly archived and cannot be deleted.

## Acknowledgments

This chapter is based on an original German document and was translated by Hazel Russman.

## Publication date and software version

Published 23 December 2015. Based on LibreOffice 5.0.

## Note for Mac users

---

Some keystrokes and menu items are different on a Mac from those used in Windows and Linux. The table below gives some common substitutions for the instructions in this chapter. For a more detailed list, see the application Help.

<b>Windows or Linux</b>	<b>Mac equivalent</b>	<b>Effect</b>
<b>Tools &gt; Options</b> menu selection	<b>LibreOffice &gt; Preferences</b>	Access setup options
<i>Right-click</i>	<i>Control+click</i>	Open a context menu
<i>Ctrl (Control)</i>	<i>⌘ (Command)</i>	Used with other keys
<i>F5</i>	<i>Shift+⌘+F5</i>	Open the Navigator
<i>F11</i>	<i>⌘+T</i>	Open the Styles and Formatting window

# Contents

---

<b>Copyright</b> .....	<b>2</b>
Contributors.....	2
Feedback.....	2
Acknowledgments.....	2
Publication date and software version.....	2
<b>Note for Mac users</b> .....	<b>2</b>
<b>General remarks on maintaining databases</b> .....	<b>4</b>
<b>Compacting a database</b> .....	<b>4</b>
<b>Resetting autovalues</b> .....	<b>4</b>
<b>Querying database properties</b> .....	<b>4</b>
<b>Exporting data</b> .....	<b>5</b>
<b>Testing tables for unnecessary entries</b> .....	<b>5</b>
Testing entries using relationship definition.....	5
Editing entries using forms and subforms.....	6
Queries for finding orphan entries.....	7
<b>Database search speed</b> .....	<b>8</b>
Effect of queries.....	8
Effect of listboxes and comboboxes.....	8
Influence of the database system used.....	8

## General remarks on maintaining databases

---

Frequent alteration of the data in a database, in particular many deletions, has two effects. First, the database grows steadily even though it may not actually contain more data. Second, the automatically created primary key continues to increment regardless of what value for the next key is actually necessary. Important maintenance is described in this chapter.

## Compacting a database

---

The behavior of HSQLDB is to preserve storage space for deleted records. Databases that are filled with test data, especially if this includes images, retain the same size even if all these records are subsequently deleted.

To free this storage space, the database records must be rewritten (tables, table descriptions, etc).

On the main Base interface, using **Tools > SQL**, you can directly enter a simple command (which on server databases is only available to the system administrator):

```
SHUTDOWN COMPACT
```

The database is taken down and freed from all accumulated waste space. After this is done, Base must be restarted in order to access the database.

Since LibreOffice version 3.6, compression is carried out by default just before the database is closed.

## Resetting autovalues

---

A database is created, all possible functions tested with examples, and corrections made until everything works. By this time, on average, many primary key values will have risen to over 100. Bad luck if the primary key has been set to auto-increment, as is commonplace! If the tables are emptied in preparation for normal usage or prior to handing the database on to another person, the primary key continues to increment from its current position instead of resetting itself to zero.

The following SQL command, again entered using **Tools > SQL**, lets you reset the initial value:

```
ALTER TABLE "Table_name" ALTER COLUMN "ID" RESTART WITH 'New value'
```

This assumes that the primary key field has the name "ID" and has been defined as an autovalue field. The new value should be the one that you want to be automatically created for the next new record. So, for example, if current records go up to 4, the new value should be 5.

## Querying database properties

---

All information on the tables of the database is stored in table form in a separate part of HSQLDB. This separate area can be reached using the name "INFORMATION\_SCHEMA".

The following query can be used to find out field names, field types, column sizes, and default values. Here an example is given for a table named Searchtable.

```
SELECT "COLUMN_NAME", "TYPE_NAME", "COLUMN_SIZE", "COLUMN_DEF" AS "Default Value" FROM "INFORMATION_SCHEMA"."SYSTEM_COLUMNS" WHERE "TABLE_NAME" = 'Searchtable' ORDER BY "ORDINAL_POSITION"
```

All special tables in HSQLDB are described in the appendix to this handbook. Information on the content of these tables is most easily obtained by direct queries:

```
SELECT * FROM "INFORMATION_SCHEMA"."SYSTEM_PRIMARYKEYS"
```

The asterisk ensures that all available columns of the table are shown. The table searched for above gives essential information on the primary keys of the various tables.

This information is useful above all for macros. Instead of having to provide detailed information on each freshly created table or database, procedures are written to fetch this information directly out of the database and are therefore universally applicable. The example database shows this, among other things, in one of the maintenance modules, where foreign keys are determined.

## Exporting data

---

Along with the possibility of exporting data by opening the \*.odb file, there is a much simpler method. Directly at the Base interface, you can use **Tools > SQL** to enter a simple command that, in server databases, is reserved for the system administrator.

```
SCRIPT 'Database'
```

This creates a complete SQL extraction of the database with all table definitions, relationships between tables, and records. Queries are not accessed since they were created in the user interface and are not stored in the internal database. However all views are included.

By default, the exported file is a normal text file. It can also be provided in binary or zipped form, especially for large databases. However this makes re-importing it into LO somewhat more complicated.

The format of the exported file can be changed using:

```
SET SCRIPTFORMAT {TEXT | BINARY | COMPRESSED};
```

The file can be read in using **Tools > SQL**, creating a new database with the same data. In the case of an internal database, the following lines must be removed before import:

```
CREATE SCHEMA PUBLIC AUTHORIZATION DBA
```

The schema already exists. It will therefore be treated as an invalid schema name.

```
CREATE USER SA PASSWORD ""
GRANT DBA TO SA
SET WRITE_DELAY 60
SET SCHEMA PUBLIC
```

These entries deal with the user profile and other default settings, which are already set for LO internal databases. They are found directly before the contents that will be inserted into the tables using the "INSERT" command.

## Testing tables for unnecessary entries

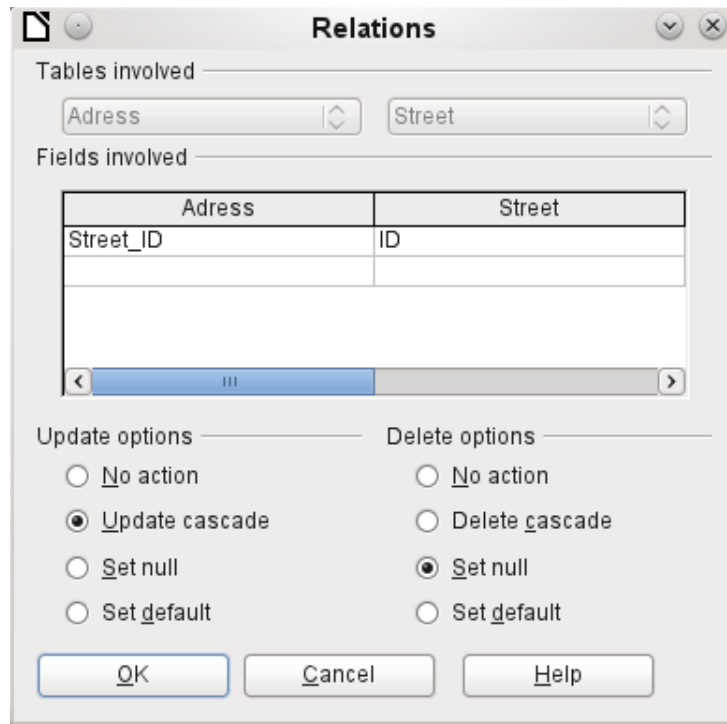
---

A database consists of one or more main tables, which contain foreign keys from other tables. In the example database, these are the Media and Address tables. In the Address table the primary key of the postcode occurs as a foreign key. If a person moves home, the address gets changed. The result may be that no foreign key Postcode\_ID corresponding to this postcode exists any longer. In principle therefore, the postcode itself could be deleted. However, it is not apparent during normal usage that the record is no longer needed. There are various ways to prevent this sort of problem arising.

## Testing entries using relationship definition

The integrity of the data can be ensured while defining relationships. In other words, you can prevent the deletion or alteration of keys from leading to errors in the database.

The following dialog is accessible through **Tools > Relationships**, followed by a right-click on the connector between two tables.



Here the tables Address and Street are considered. **All specified actions apply to the Address table**, which contains the foreign key Street\_ID. Update options refer to an update of the ID field in the Street table. If the numeric key in the "Street"."ID" field is altered, "No action" means that the database resists this change if a "Street"."ID" with that key number occurs as a foreign key in the Address table.

"Update cascade" means that the key number is simply carried over. If the street 'Burgring' in the Street table has the ID '3' and is also represented in "Address"."Street\_ID", the ID can be safely altered, for example, to '67' – the corresponding "Address"."Street\_ID" values will be automatically be changed to '67'.

If *Set null* is chosen, altering the ID makes "Address"."Street\_ID" an empty field.

The delete options are handled similarly.

For both options, the GUI currently does not allow the possibility *Set default*, as the GUI default settings are different from those of the database. See Chapter 3, Tables.

Defining relationships helps keep the relationships themselves clean, but it does not remove unnecessary records that provide their primary key as a foreign key in the relationship. There may be any number of streets without corresponding addresses.

## Editing entries using forms and subforms

In principle the whole interrelationship between tables can be displayed within forms. This is easiest of course when a table is related to only one other table. Thus in the following example, the author's primary key becomes the foreign key in the table rel\_Media\_Author; rel\_Media\_Author also contains a foreign key from Media, so that the following arrangement shows a n:m-relationship with three forms. Each is presented through a table.

The first figure shows that the title *I hear you knocking* belongs to the author *Dave Edmunds*. Therefore *Dave Edmunds* must not be deleted – otherwise information required for the media

*I hear you knocking* will be missing. However the listbox allows you to choose a different record instead of *Dave Edmunds*.

The screenshot shows a data form with three tables. The first table, 'Last Name' and 'First Name', lists authors: Edmunds (Dave), Götze (Dr. Lutz), Hawking (Steven W.), Heller (Dr. Klaus), and Hermann (Ursula). The second table, 'Author' and 'Author\_Sort', shows 'Edmunds, I' selected. The third table, 'Title', shows 'I hear you knockin'. The status bar indicates 'Record 1 of 10'.

In the form there is a built-in filter whose activation can tell you which categories are not needed in the Media table. In the case just described, almost all the example authors are in use. Only the *Erich Kästner* record can be deleted without any consequences for any other record in *Media*.

The screenshot shows the same data form but with only one record: 'Kästner' (Erich). A yellow gear icon (filter) is visible in the first table. The status bar indicates 'Record 1 of 1'. A yellow button labeled 'Apply Filter' is shown below the form.

The filter is hard-coded in this case. It is found in the form properties. Such a filter is activated automatically when the form is launched. It can be switched off and on. If it is deleted, it can be accessed again by a complete reload of the form. This means more than just updating the data; the whole form document must be closed and then reopened.

The 'Form Properties' dialog box is shown with the 'Data' tab selected. The 'Filter' property is set to the SQL query: `"ID" NOT IN (SELECT "Author_ID" FROM "rel_Media_Author")`. Other properties include Content type (Table), Content (Author), Analyze SQL command (Yes), Sort ("LastName" ASC, "FirstName" ASC), Allow additions (Yes), Allow modifications (Yes), Allow deletions (Yes), Add data only (No), Navigation bar (No), and Cycle (Default).

## Queries for finding orphan entries

The above filter is part of a query which can be used to find orphaned entries.

```
SELECT "Surname", "Firstname" FROM "Author" WHERE "ID" NOT IN (SELECT "Author_ID"
FROM "rel_Media_Author")
```

If a table contains foreign keys from several other tables, the query needs to be extended accordingly. This affects, for example, the Town table, which has foreign keys in both the Media table and the Postcode table. Therefore, records in the Town table which are to be deleted should not be referenced in either of these tables. This is determined by the following query:

```
SELECT "Town" FROM "Town" WHERE "ID" NOT IN (SELECT "Town_ID" FROM "Media")
AND "ID" NOT IN (SELECT "Town_ID" FROM "Postcode")
```

Orphaned entries can then be deleted by selecting all the entries that pass the set filter, and using the Delete option in the context menu of the record pointer, called up by right-clicking.

## Database search speed

---

### Effect of queries

It is just these queries, used in the previous section to filter data, that prove unsatisfactory in regard to the maximum speed of searching a database. The problem is that in large databases, the subquery retrieves a correspondingly large amount of data with which each single displayable record must be compared. Only comparisons with the relationship IN make it possible to compare a single value with a set of values. The query

```
... WHERE "ID" NOT IN (SELECT "Author_ID" FROM "rel_Media_Author")
```

can contain a large number of possible foreign keys from the rel\_Media\_Author table, which must first be compared with the primary keys in the table Authors for each record in that table. Such a query is therefore not suitable for daily use but may be required for database maintenance. For daily use, search functions need to be constructed differently so that the search for data is not excessively long and does not damage day-to-day work with the database.

### Effect of listboxes and comboboxes

The more listboxes that are built into a form, and the more they contain, the longer the form takes to load, since these listboxes must be created.

The better the Base program sets up the graphical interface and initially reads the listbox contents only partially, the less delay there will be.

Listboxes are created using queries, and these queries must be run when the form is loaded for each listbox.

The same query structure for more listboxes is better done using a common View, instead of repeatedly creating fields with the same syntax using the stored SQL commands in the listboxes. Views are above all preferable for external database systems, as here the server runs significantly faster than a query which has to be put together by the GUI and freshly put to the server. The server treats Views as complete local queries.

### Influence of the database system used

The internal HSQLDB database is set up so as to ensure that Base and Java work well together. Unfortunately since Version LO 3.5, Base has had problems with the speed of processing in precisely this area. This becomes particularly noticeable when dealing with large tables with several thousand records. These problems have various causes and only disappeared with Version 4.1.

External databases run significantly faster. Direct connections to MySQL or PostgreSQL, and connections using ODBC run at practically the same speed. JDBC also depends on cooperation with Java, but still works faster than an internal connection using HSQLDB.