OOo Base: Set

Databases are hardcore, complex and only for experienced hackers. Well, almost! **Andy Channelle** takes the bull by the horns and delves into *OpenOffice.org*'s *Base*.

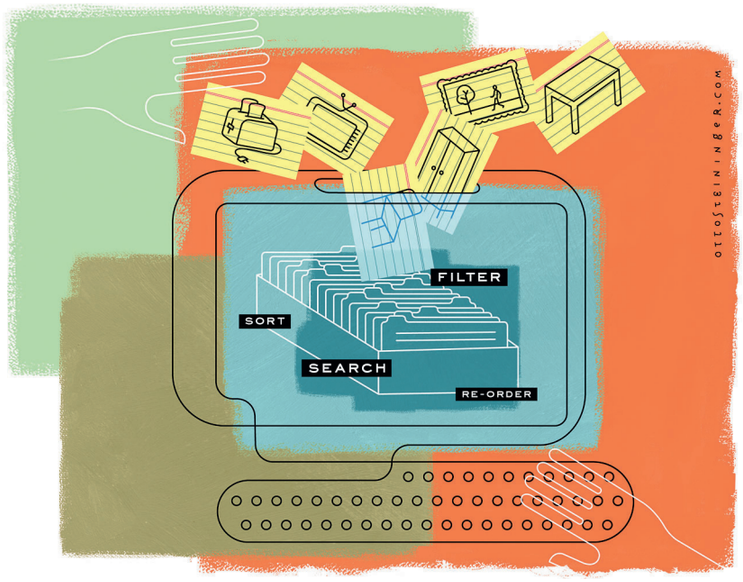

For many average computer users, the database is out of bounds; something too uncomfortably close to really hard work to be worthwhile. However, there are a couple of applications that are attempting to break this barrier and bring the joy (did I really say joy?) of databases to the masses. The most recent version of *KOffice* shipped with a fine DB application called *Kexi*, while the *OpenOffice.org* suite has been swelled with the addition of *Base*. We're concentrating on the latter here because it is less environment-centric than *Kexi* (though you can, of course use *Kexi* on a Gnome desktop) but also because *OOo* seems to be a little better at the handholding we'll require in the early stages of this project.

In fact, the database I'm creating in this tutorial is an inventory of the contents of my house just in case they're ever stolen or lost in some other way, and *OOo Base* has a wizard – that is, a pre-defined script – designed just for this purpose. So we'll be starting with that as an introduction to the concepts we are discussing and then make a few changes to the basic database to provide a little more depth.

On the coverdisc you'll find the latest stable version of *OOo*, version 2.0.3, which has *Base* ready and waiting for you along with the other elements of the office suite.

Our expert

Andy Channelle
Andy has been taking his first steps in Linux software for the past six years and has been interested in technology since the advent of the Dragon 32.

Part 1 Just what is a database?

Let's launch the application first. Select *OpenOffice.org Base* from (usually, depending on which distribution you're running) the Office menu. When the application window comes up, select Create A New Database. In the second screen, ensure that 'Yes, register the database for me' (which tells *OOo* where the data is being held) and 'Open the database for editing' are both selected. Hit the Finish button to get started.

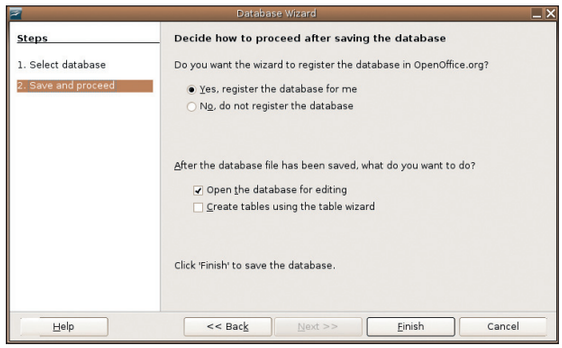

› Registering the database makes *OOo* aware of it. Useful if you want to create documents based on a big address book.

The job of building a database essentially involves two elements: structure and content. The content is really up to you, so we'll be concentrating here on the structure of the database, and this is initially defined using the Tables tool.

Able tables

Anyone who has worked with a spreadsheet will be familiar with the visual representation of data in a database table, based, as it is, on the notion of rows, columns and cells. In fact, there's a sound argument for choosing to use a spreadsheet rather than a database when dealing with small amounts of data; it's less complex, quicker to set up and easier to edit the structure of the whole thing visually in the future.

But the benefits of a database are that data can come from a number of sources and we can work – if necessary – with vast data sets that don't all need to be loaded into the computer's memory at once. Plus it's much easier to build up the data set using a well-designed 'form' containing text entry points, checkboxes and radio buttons than with a series of columns, and we can also define relationships between data sets to make managing large amounts of data less troublesome, if needed.

Convinced? Right, back to the tables and – the First Steps mantra – a little planning. The first step on the road to building a

» **Last month** We built a kid-friendly desktop with restricted access and web filters.

up a database

database is understanding exactly what it is supposed to contain and the relationships that you want to define, if any, between data. The best way to work this out is to grab a piece of paper and make a list of the information the database will have in it. In our case, we're building up a home inventory, so several elements – product name, price, purchase date and so on – would appear to be crucial. We're building a very simple data set on a single table, so these elements will become the database's 'fields', or headings that will require data input.

Defining the data

We also need to think about what type of data each entry will need. The Product Name field, for instance, will need to be a text field, so the user can enter plain old text into it, whereas the Price field will need to be a number and, more explicitly, a decimal or currency field so that prices entered in here will automatically be formatted in the correct way.

We will only be using a few data types in our database at the moment, including text, decimal and date, but in future it would be possible to add particular categories (kitchen, computer, music etc) in a second table, and these might include a 'yes/no' option to make it easy to define the categories either using checkboxes or a drop-down list in the form. A more extensive database can use two or more tables. For example, a school database could have one table containing data about teachers and a second housing records for children; these could then be linked on the basis of the relationship between an individual teacher and the pupils in their class. A third table might have exam results, which could then be linked in with both teacher and pupil tables.

I now have a list of seven field headings: item, manufacturer, model, purchase price, purchase place, purchase date and serial number. These need to be joined by a Record ID field, so it's time to launch the first wizard in *OOo Base*: the Table Wizard.

Burglar alarm

It's no good having all this important information on a computer if the computer is one of the things that goes up in smoke or is taken by your local stripey-jumpered burglar. Backups, therefore, are vital; and if you're even a little paranoid, off-site backups (where you store CDs, DVDs or files elsewhere) are a very good idea. Fortunately, in the case of our simple database, the whole shebang is wrapped up into a nice, compact ODB file that can be attached to an email and sent to an off-site webmail account or burned to a CD and stashed away in your granddad's attic.

Part 2 A little light wizardry

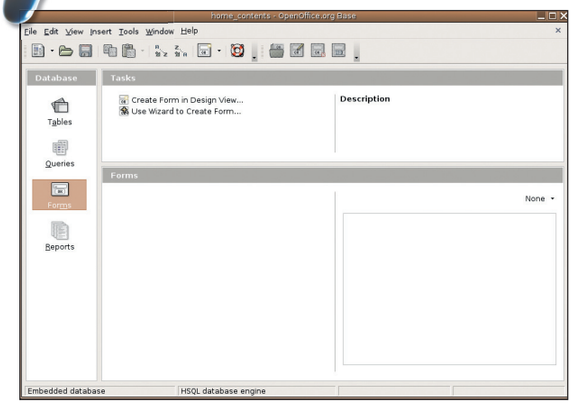

1 Summon the wizard

In the main interface page, select the Tables icon in the left pane and click on the Use Wizard to Create Table option. This will launch the wizard. If, by the way, you're working on something other than a home inventory and *OOo* doesn't have a wizard for it, select Create Table In Design View instead. The process of building the table is the same; you'll just need to do a little more typing.

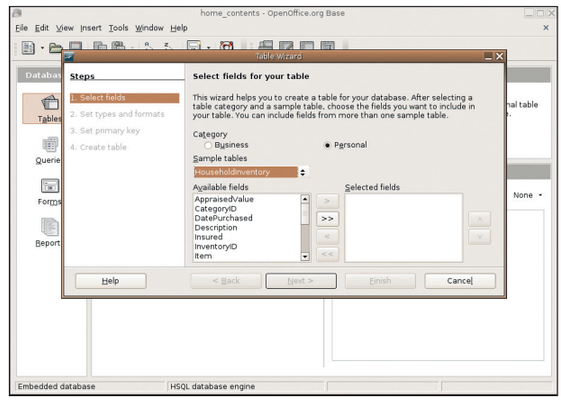

2 Pick your fields

Choose from the available options using the Business and Personal categories then use the drop-down list in the middle of the window to select the project – in this case we're using Personal and Household Inventory – which will then provide a list of available fields in the left of the two boxes in the centre of the window. You can add fields to the database by selecting one with a mouse and using the top arrow key in the centre to shift it to the right (the double arrows act as a select-all option). Once added, the position of fields within the table can be adjusted with the up and down arrows on the right.

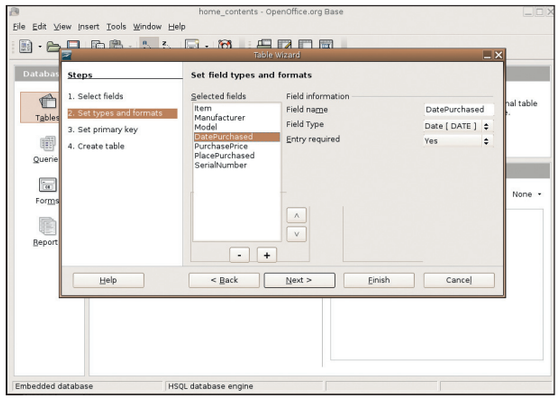

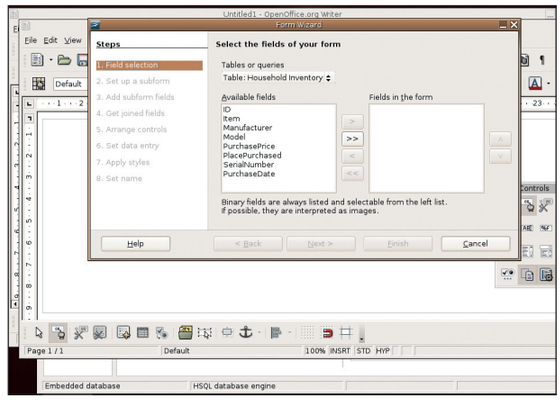

3 What kind of data?

The next section allows you to set the data type for each field. It will mostly be Text, but we also want one Decimal (for the price) and one Date (for the purchase date). We can set the Required status here too; if we set this to Yes, the user won't be able to save a record until this field has been completed. There's also an option in here to rearrange the order again and add or remove fields using the + and - keys at the bottom. If you don't use the wizard, these last two stages are done in a single window which looks like a spreadsheet. You just enter the field names on the left and then use the drop-down list in the following column to select the data type.

4 Create a form

The next stage is to run the Form wizard. A form does exactly what its name implies: presents on-screen spaces for the user to fill in, and then links these spaces to the table. Choose the Forms icon on the left then select Use Wizard To Create Form. The resulting window is quite similar to the Table wizard, with available fields on the left that can be sent to the right pane with the arrow button. Again, just select each field in turn. We'll skip straight through the second section as we don't have any sub-forms in this database (these are forms within the form that, for example, might be linked to a second table).

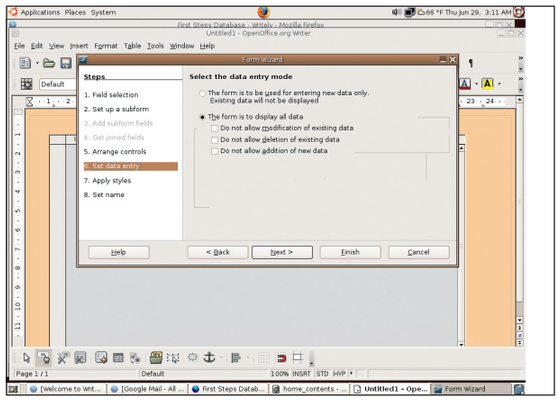

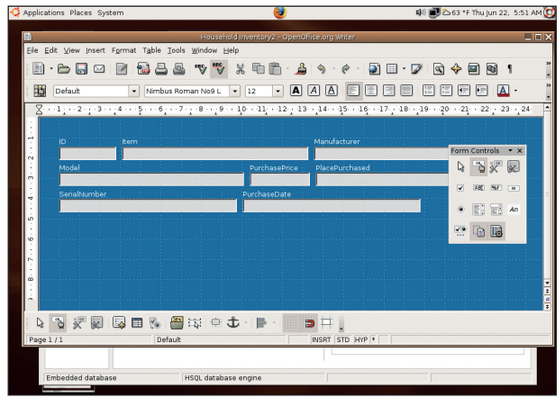

5 Entry restrictions

Now we'll choose the style of the form – that is, where each entry field and label will appear. Choose whichever suits and move on using the Next button to the Set Data Entry section. This is a space where we can set restrictions on how data can be entered and edited later on. There are two options to start with, and we're choosing the second (This Form Is To Display All Data), which simply means that users will be able to navigate through database entries with the *OOo* forms user interface. Leave all the options below this unchecked.

6 Final step: Style it up

Finally there is a chance to choose from *OOo Base*'s form styles, give the whole thing a name and view the result: a form ready for the inputting of data! It's worth noting that you can build tons of different forms for the same database just by re-running the wizard. So, in this case, you might want to have one form that's complete, and a second one that doesn't, for example, include the Purchase Price element, making it possible to hide from your partner how much money you've spent on computers, guitars or rare LPs. Or is that just me?

Brigitte Bardot uses Mandriva

Well, how else was I going to get you to read this bit about database models? But while you're here... Wikipedia lists six common types of database, but the most common are flat file and relational. The flat file has a two-dimensional array of data: columns and rows. The rows are related – they are all aspects of the same entity – and the columns will be of a similar value – they all represent a monetary value, date or whatever. The database we're creating in this tutorial is a good example of a flat file. In contrast, a relational database might contain any number of flat files (or tables) and is good for defining and exploring relationships between them. Splitting the data across a number of tables and then using virtual wires to highlight the relationships between them reduces the need for data replication, and allows you to examine patterns with far more depth than a single flat file might. My exam results table could be associated with a teacher, a cohort of pupils and an education authority or country.

Part 3 Adding and playing with data

The two wizards have given us a space to add the content of our database, and we can begin adding our records straight away. There are two ways to do this. The first is to go into the Table section of the database, by selecting the Table icon in the left pane, then right-click on the table that we created earlier and select Open. This will take you into the spreadsheet-like table entry window with the database fields ranged across the top row of cells and a line of blank cells beneath them. Click in a cell, type in some data and hit Tab to advance to the next cell. Enter data in each cell until all the data has been entered. You may notice that as you began to type the first bit of text into the first cell, a new series of cells appears below, ready for the next record. When you get to the last cell of the first line, hitting Tab will take the cursor into the first cell of the next line ready for a new record.

Oh, and the second way to populate a database? You'll just have to tune in next time...

Walk with me

There are some important tools in this table entry window and it is possible to do all your work in here, so it's worth exploring. Along the top of the window is the menu bar with all the usual options. Below this is the tool bar, where you'll find all the stuff to search, filter and reorder the entries in your database. In the image above right, the first six icons are the usual cut, copy and paste bunch; however, take note of the very first icon (a small floppy disc with a black triangle). This is the button to save the current record, and if you make a significant change to a record, it's a good idea to hit this and commit the change to the database. Note that once a record has been entered and you tab down to a new line, it is saved automatically, but this is not the case when you edit one.

The next really interesting icon is the magnifying glass, which opens a search dialog. It's not enormously useful for our paltry data set at the moment, but as things get more expansive it can be really useful. It has options for searching through individual fields for a text string, or through every single field, plus the ability to examine a whole field, just the beginning or just the end. As with most search tools, we define our search criteria, hit the Search button and sit back to await the results.

Following the search is the Refresh/Rebuild button, which I'll introduce you to next time. We then have three sort tools. The first

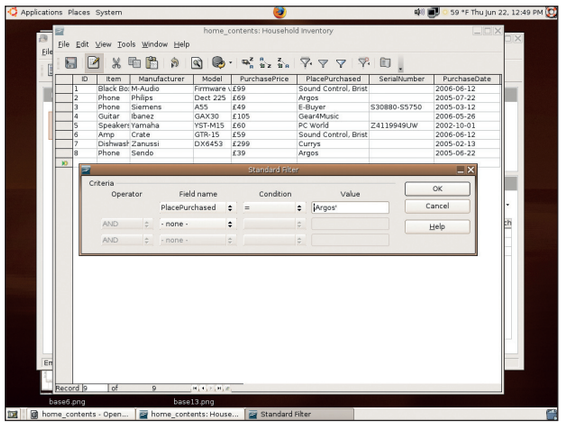

› Filtering is great for isolating parts of the table. For what it's worth, I'm a very thrifty shopper.

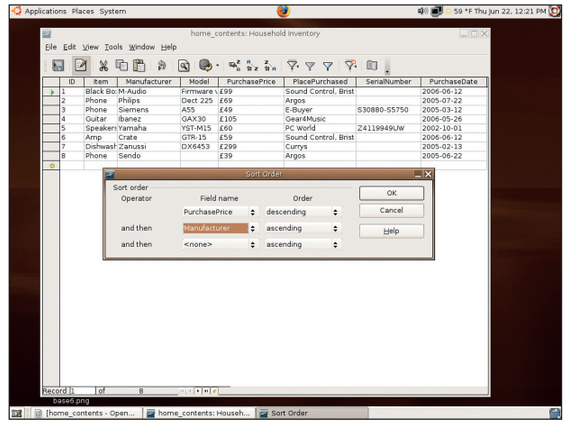

› Sorting changes the order of the displayed data.

is a catch-all Sort option. Clicking on it will bring up a dialog box allowing us to, for example, sort all of our records into a list based on how much everything costs (choose Ascending to start with the cheapest, or the letter A and Descending to go with the most expensive or reverse alphabetically) or the date of purchase. The Sort option works in the same manner as the search dialog. Define a pattern and hit OK to have the data reordered to it. If you like, you can perform less complicated reordering – sorting by purchase date, for example – by simply selecting a field, clicking on its name, and then hitting either of the other two sort icons. These are, respectively, Sort Ascending and Sort Descending. You can always get back to the original state of the database by selecting the ID field, which we've used to number each record consecutively, and choosing the Sort Ascending button.

Caught in the filter

Following the Sort options come the filters. These allow us to isolate particular records or batches of records based on definable criteria. The first icon is Auto Filter. This works by taking the currently selected cell as its filter. If I selected a Purchase Price cell with the value £65 in it then hit the Auto Filter button, every entry that didn't contain £65 in that field would disappear from the screen, and those that did would be shown at the top of the screen. To get rid of the Auto Filter and return to the main data set, use the Remove Filter/Sort button, which is the fourth of this quartet of icons. The third filter icon is the standard filter, and this provides more flexible options for filtering on a number of different fields (as with the Sort option above).

At the bottom of the table entry window is a notification and navigation area. This comes into its own in larger databases, telling the user what record they are currently browsing or how many there are in total, and providing a convenient method of shifting from record to record – use the icons or type a number in the Record X Of X box – or going to either the first or last entry in the database.

Well, our database is complete, and how smug do we feel? Shovelling data into a database can be a tedious business, but once the whole thing is set up and running, it's easy to add new items as and when they are purchased. It's a case of short-term pain for long-term gain. Bring on the burglars! LXF

Quick tip

When you use Sort, it is possible to put a second criterion in, so you could order on price first, say, and then alphabetically by manufacturer.

» Next month We'll bring in data from other sources and try queries and reports.