CHAPTER 3

Field Types and Import/Export Formats

Knowing Your Data

Besides just knowing the raw statistics and capacities of your software tools ("speeds and feeds," as the machinists like to say), it's wise to understand the basic data formats with which your database software can work. In this chapter we give a brief overview of each of the underlying FileMaker field data types, and we also discuss the various data formats available for import and export.

First, a word about data types in FileMaker: Unlike certain other database engines, FileMaker is somewhat forgiving about data types. You can add text characters to a number field, for example, without being stopped by FileMaker. The price you pay for this flexibility is the possibility that you'll make a mistake or permit something undesirable to happen. So it's necessary to have a firm grasp of FileMaker's data types, and how each of them works and relates to the others.

FileMaker Field Types

In the previous chapter, on FileMaker specifications, we called out a number of the raw capacities of FileMaker's field types, in the briefest possible format. Here we dwell on each data type and its characteristics in a bit more detail.

➔ *For more information on FileMaker fields and field types, **see*** Special Edition Using FileMaker 8*, Chapter 3, "Defining and Working with Fields.*

Text

FileMaker's text field type is able to hold up to 2 gigabytes of character data per text field. FileMaker 8 stores text internally as Unicode, which requires about 2 bytes per character (the exact number varies by the specific character and encoding). So a single FileMaker text field can hold about a billion characters of data (nearly half a million regular pages of text).

There's no need to specify the "size" of the text field in advance; FileMaker will automatically accommodate any text you enter, up to the field size maximum.

FileMaker's text field can be indexed on a word-by-word basis, so it's possible to search quickly for one or more words anywhere within a text field.

In other database systems, such as those based on SQL, it is often necessary to specify the maximum size of a text field in advance. Many SQL databases do go beyond the SQL standard in offering a text data type of flexible size.

Number

The number field type is FileMaker's only numeric data type. Unlike some other tools, FileMaker does not have separate data types for integers and floating point numbers. All numbers are capable of being treated as floating point.

FileMaker 8's numeric data type can store numbers in the range 10^{-400} through 10^{400}, and -10^{400} through -10^{-400}, as well as the value 0. The data type can account for 400 digits of precision. By default, though, numeric values will only account for 16 digits of precision. To achieve higher precision, you'll need to use the SetPrecision() function.

Numeric fields can accept character data entry, though the character data will be ignored. If the field's validation is also set to Strict Data Type: Numeric Only, even the entry of character data will be disallowed.

In other database systems, such as those based on SQL, integers and floating point numbers tend to be considered as two different data types, and it's often necessary to choose in advance between lower-precision and higher-precision floating point numbers, though SQL does also offer arbitrary-precision numbers as well.

Date

FileMaker's date type can store dates from 1/1/0001 to 12/31/4000. Internally, these dates are stored as integer values between 1 and 1,460,970, indicating the number of days elapsed between 1/1/0001 and the date in question. This is significant in that it means that integer math can be performed on dates: 12/31/2001 + 1 = 1/1/2002, and 12/21/2001 – 365 = 12/21/2000.

FileMaker will correctly interpret oddly formed dates, as long as they're within its overall date range. For example, 4/31/2005 will be interpreted as 5/1/2005, whereas 12/32/2005 will be interpreted as 1/1/2006. But 12/31/5000 will be rejected outright.

FileMaker's date type is stricter than the text and number types, and will reject any date outside the range just mentioned, including those containing textual data. Field validation options can also be configured to force entry of a full four-digit year.

Though the date type is stored as an integer, it can be entered and displayed in a wide variety of well-known or local date display formats. By default it will use settings it inherits from the current user's operating system, but the date display format can be overridden on a field-by-field basis on each individual layout.

By comparison, SQL-based database systems also offer a date data type, measured in days.

Time

Like the date type, FileMaker's time data type stores its information in an underlying integer representation. In the case of the time field type, what's being stored is the number of seconds since midnight of the previous day, yielding a range of possible values from 1 to 86,400. As with the date type, it's possible to perform integer math with time values: So `Get(CurrentTime) + 3600` will return a time an hour ahead of the current time. And `15:30:00 – 11:00:00` returns 4:30:00, so "time math" can be used to compute time intervals correctly as well.

Also like the date type, the time data type can accept entry and display of time data in a variety of formats. "Overlapping" times work in the same way as overlapping dates: FileMaker will interpret the value 25:15:00 as representing 1:15 a.m. on the following day. (This is exactly what happens when you store a value greater than 86,400 in a time field, which FileMaker will certainly allow.)

By comparison, SQL-based database systems generally offer a time data type, sometimes with the capability to track a time zone as well.

Timestamp

FileMaker's timestamp data type is like a combination of the date and time types. A timestamp will be displayed as something like "11/20/2005 12:30:00." Internally, like both dates and times, it is stored as an integer. In this case the internally stored number represents the number of seconds since 1/1/0001. The timestamp can represent date/time combinations ranging from 1/1/0001 00:00:00 to 12/31/4000 11:59:59.999999 (a range of over a hundred billion seconds).

As with the related date and time field types, timestamps can be displayed and entered in a variety of formats. It's also possible to perform math with timestamps as with dates and times, but since timestamps can measure time in days but are stored in seconds, the math may get a bit unwieldy.

The timestamp data type is commonplace in the database world. Some flavors of timestamp offer the capability to track a time zone as well.

Container

The container data type is FileMaker's field type for binary data (meaning data that, unlike text and numbers, does not have an accepted plain-text representation). Binary data generally represents an electronic file of some sort, such as a picture, a movie, a sound file, or a file produced by some other software application such as a word processor or page layout program.

A single container field can store a single binary object up to 4 gigabytes in size. Though any type of electronic file can be stored in a container field (one file per field), FileMaker has specialized knowledge of a few types of binary data, and can play or display such objects directly by calling on operating system services. Specifically, FileMaker can play or display pictures, sounds, and movies (and in fact any "movie-like" file supported by QuickTime, if QuickTime is installed). Additionally, on Windows, a container field can store and display Object Linking and Embedding (OLE) objects.

When working with container fields, it's important to understand the distinction between storing the data directly (embedding the entire file in the FileMaker database and increasing its file size by tens or hundreds of K or indeed by megabytes) and storing it by reference (storing just the path to the specified file). The former leads to larger database files; the latter relies on the binary files being stored in a fixed location and not moving relative to the database.

Many other database systems are able to handle binary data, but there's little in the way of an agreed-upon standard for doing so. In general, each database tool has its own means for dealing with binary data.

Calculation

From the point of view of data types, calculations are not a data type at all. They are instead a *field* type, and can be constructed to return their results as any of the six fundamental data types: text, number, date, time, timestamp, or container.

Summary

Summary fields are another example of a field type rather than a data type. Summary fields perform summary operations on sets of grouped records, and always return a numeric result.

Importing Data

FileMaker can import data from one or more static, external files on the hard drive of the FileMaker client machine or on a shared network volume. FileMaker can also import data from remote data services. The following sections outline the various sources for importing data and the specific requirements and limitations of each.

➔ *For more information on importing data in FileMaker,* ***see*** Special Edition Using FileMaker 8*, Chapter 19 "Importing Data into FileMaker."*

File-Based Data Sources

FileMaker can import data from individual files, available on a local hard drive or networked volume, in any of the following formats:

- Tab-separated text
- Comma-separated text
- SYLK
- DIF
- WKS
- BASIC
- Merge files
- DBase files

In addition to importing from these common text file formats, FileMaker can also perform more specialized imports from files created in Excel, or in FileMaker itself.

Importing from Excel

When importing data from an Excel file, FileMaker can detect multiple worksheets within the source Excel file. FileMaker can also detect the existence of any *named ranges* in the source document. When importing, if named ranges or multiple worksheets are detected, FileMaker will give you a choice as to whether to import from a worksheet or named range, and allow you to select the specific worksheet or range from which to import.

When importing from Excel, as with all imports, FileMaker brings in only the raw data it finds in the source file. Formulas, macros, and other programming logic are not imported.

FileMaker Pro assigns an appropriate field type (text, number, date, or time) if all rows in the column hold the same Excel data type. Otherwise, a column becomes a text field when imported into FileMaker.

Importing from FileMaker

FileMaker can also import data from other FileMaker files. These may be files that are present on the local client machine, or files that are hosted on another machine.

Much as an Excel file can contain multiple worksheets, a FileMaker database can contain multiple tables. It's necessary to choose a single table as your data source when importing from a FileMaker file.

Importing from FileMaker can be particularly convenient if the source file has a structure that matches that of the target file. In this case, rather than manually configuring the import mapping on a field-by-field basis, it's possible to choose the Arrange by Matching Field Names option. When you do so, fields of the same name in the source and target tables will be paired in the import mapping.

Importing Multiple Files at Once

It's also possible to import data from multiple files at a time. You can import data from either text or image files. You can import both the raw data in the file, and also extra data about each source file's name and location.

In each case, the files being imported must all be grouped in or underneath a single folder. You can specify whether to look simply inside the one folder, or whether to search all the way down through any subfolders.

Importing from Multiple Text Files

When importing from a batch of text files, you may import up to three pieces of data from each text file:

- Filename
- Full path to file
- Text contents

You may choose to import any or all of these.

Unlike with a regular import from a text file, the internal structure of the text file is disregarded. The entire contents, whether containing tabs, carriage returns, commas, or other potential delimiters, is imported into a single target field.

Importing from Multiple Image Files

Importing from multiple image files is quite similar to importing from multiple text files. When importing a batch of images, you may import any or all of the following data fields:

- Filename
- Full path to file
- Image
- Image thumbnail

In addition to filename, file path, and file contents (an image, in this case), FileMaker allows you to import an image thumbnail, either in addition to or instead of the full image. You may wish to do this to save file space or screen space. FileMaker creates these thumbnails via its own algorithms, so you have no control over the exact details of thumbnail size or quality.

When importing images, you have the choice (as you always do when working with data in container fields) of importing the full image into the database, or merely storing a reference. Importing full images will take up more space in the database (probably much more), whereas importing only the references will mean that you'll need to make the original files continuously available from a hard drive or network volume that all users of the database can access.

Importing Digital Photos (Mac OS)

On the Mac OS, FileMaker can also import images directly from a digital camera, or from any device capable of storing digital photos. This process is quite similar to a batch import of images from a single folder with only a few differences.

FileMaker will allow you to specify which images to import. You may choose them individually, or via a range such as "last 12 images." FileMaker will also handle transferring the files from the storage device to a download location of your choice. And, as with other imports, you can choose whether to import the full image into FileMaker, or simply store a reference.

Whereas the regular batch import of images brings in only four pieces of data about each image, a digital image import may have access to much more data about each image. If the selected images contain EXIF data (a standard for embedding extra data into an image file), FileMaker can also detect and import many additional pieces of data about the image such as shutter speed, ISO setting, and the like.

Importing from an ODBC Data Source

FileMaker can import data from a data source accessed via ODBC. Many types of data can be accessed via ODBC, but it's most commonly used to retrieve data from a remote database, often one running some flavor of the SQL language.

Working with OBDC data sources requires three things:

- **A data source able to provide data via ODBC**—Again, this is most often a remote database server of some kind. The administrators of the data source may need to perform specific configuration of the data source before it can accept ODBC connections.
- **An ODBC driver, installed on the local computer that's running FileMaker, that's able to talk to the specific ODBC data source in question**—ODBC drivers need to be installed on each computer that will access a data source. So, much like a FileMaker plug-in, ODBC drivers generally need to be installed on the computer of each FileMaker user who will be using ODBC access. ODBC drivers are specific to a particular data source (the PostgreSQL or Sybase databases, for example), and also specific to a particular platform (Mac or PC). In order to connect to an ODBC data source, you must have a driver specific to both your data source and platform (Sybase 12 driver for Mac OS, for example).
- **A DSN (Data Source Name) that specifies the details of how to connect to a specific data source**—DSNs are configured differently on each platform, and generally contain information about a specific data source (server name, user name, password, database name, and the like).

➔ *For more information on configuring ODBC access in FileMaker,* ***see*** Special Edition Using FileMaker 8*, Chapter 19, "Importing Data into FileMaker," and Chapter 20, "Exporting Data from FileMaker."*

Once you have successfully configured and connected to an ODBC data source, the process for selecting data to import is a bit different than for regular imports. Before proceeding to the field mapping dialog, you'll need to build a SQL query that selects the fields and records you want. (For example, your SQL query might read SELECT name_last, name_first, city, state, zip FROM customer). Once you've done this, you'll be able to map the resulting fields to those in your FileMaker database.

Importing from an XML Data Source

FileMaker can import data from XML-based data sources. The source can either be a physical file, stored on a locally accessible volume, or a remote XML data stream accessed over HTTP.

In order to import data from an XML source, the XML data must be presented in the correct format, which FileMaker calls a *grammar*. FileMaker can only import data from XML sources that use the FMPXMLRESULT grammar. If the XML data is not in the FMPXMLRESULT grammar, you will need to specify an XSLT stylesheet at the time of import; that stylesheet needs to be written in such a way as to transform the inbound XML into the FMPXMLRESULT grammar.

➔ *For more information on FileMaker's XML import capabilities,* ***see*** Special Edition Using FileMaker 8, *Chapter 22, "FileMaker and Web Services."*

Creating New Tables on Import

NEW FileMaker 8 has a nice new capability. When importing data, you can now specify that the inbound data should be placed in a new table rather than adding to or updating an existing table. The new table will take its field names from those present in the data source.

Exporting Data

FileMaker can export its data to a variety of data types, as follows:

- Tab-separated text
- Comma-separated text
- SYLK
- DIF
- WKS
- BASIC
- Merge
- HTML table
- FileMaker Pro
- Excel

Most of these are straightforward: The records in the found set will be exported to a file of the specified format.

A few of these formats deserve special mention and are detailed in the sections that follow.

➔ *For more information on exporting data in FileMaker,* ***see*** Special Edition Using FileMaker 8*, Chapter 20, "Exporting Data from FileMaker."*

Exporting to HTML

When exporting to "HTML Table," the resulting file will be a complete HTML document consisting of an HTML table that displays the chosen data, somewhat like FileMaker's Table View.

Exporting to FileMaker Pro

The result of this export choice will be a new FileMaker Pro 7/8 file, with fields and field types that generally match those of exported fields. But note that the logic underlying calculation and summary fields will not be preserved. Those fields will be re-created based on their underlying data type: A calculation field returning a text result will be inserted into a Text field, whereas summary field data, if exported, will be inserted into a Number field. Only the raw data from the original file will be exported; none of the calculation or scripting logic will carry over.

Exporting to Excel

When exporting records to Excel, you have the option to create a header row where the database field names will appear as column names. You can also specify a worksheet name, document title, document subject, and document author.

NEW Exporting records to Excel is similar to, but a bit more flexible than, FileMaker 8's new Save Records As Excel feature. Save Records As Excel will only export fields from the current layout, and will not export fields on tab panels other than the active one.

Automatically Opening or Emailing Exported Files

NEW New in FileMaker 8 is the capability to automatically open and/or email an exported file. These choices are selected via two checkboxes in the Export dialog.

Automatically opening a file is a convenience for the user: A newly created Excel file would open right away in Excel. Automatic email is an even more powerful tool. When this choice is selected, a new email will be created in the user's default email client with the exported file as an attachment. (The user will still need to specify the email recipients manually.)