

# Visual Studio 2005 and MySQL

[Connector/NET 5.0](#) [Project specification](#) [Tools needed](#) [Class layout](#)  
[Set up project](#) [Encapsulate connection](#) [Login](#) [Properties](#)  
[Select database](#) [Select tables for grids](#) [Grid paging](#) [Action buttons](#) [Find](#)  
[Configure and populate grids](#)  
[Build queries](#) [Grid manager](#) [Detail grid WHERE clause](#) [Event handlers](#)  
[Just-in-time grid lookups](#) [Grid popup menu](#) [ListView lookups](#) [Event declarations](#)  
[SQL Server version](#) [Installing and running the code](#)

Chapter 15 introduced the basics of using MySQL with Microsoft .NET. But .NET is huge, .NET 2.0 and Visual Studio 2005 are bigger than earlier versions, and MySQL Connector/NET 5.0 is a different beast from 1.0.7, which does not work well with VS 2005. In this chapter we investigate whether you can stay with MySQL 5 for substantial Visual Studio database application development, or whether you need to migrate to SQL Server 2005, which is expensive, or SQL Server Express, which is free but limited in capability.

There's lots of code here. Feel free to lift it whole or in [chunks](#), which include:

- how to get the info you need from `information_schema`,
- login forms, including reading from and writing to the Windows registry,
- populating ComboBoxes from a database,
- how to implement a clientside cache for database data,
- configuring DataGridViews to use such caches for tables of any size,
- databinding and update management with BindingSource objects,
- context menus,
- custom events,
- on-the-fly lookup data lookup browse windows using a DataGridView,
- on-the fly lookup data lookup browse windows using a ListView,
- on-the-fly column data find dialogues,
- how to write for multiple DBMS backends.

We assume installations of .NET 2.0, Visual Studio 2005, and MySQL 5.0.22 or later. To compare MySQL and SQL Server versions, you also need SQL Server or SQL Express.

## Connector/NET 5

The alpha release of MySQL Connector/NET 5.0 appeared in August 2006. Download the most recent release from <http://dev.mysql.com/downloads/connector/net/5.0.html>. Installation is point-and-click. Once you have installed it, navigate to its *docs* folder and drag a Shortcut for its help file *MySql.Data.chm* to a convenient desktop location.

Connector/NET 5.0 is stable for an early release. Like 1.0, it implements two classes, both referenced ([MySql.Data.MySqlClient](#), [MySql.Data.Types](#)) on the MySQL website and in the help file.

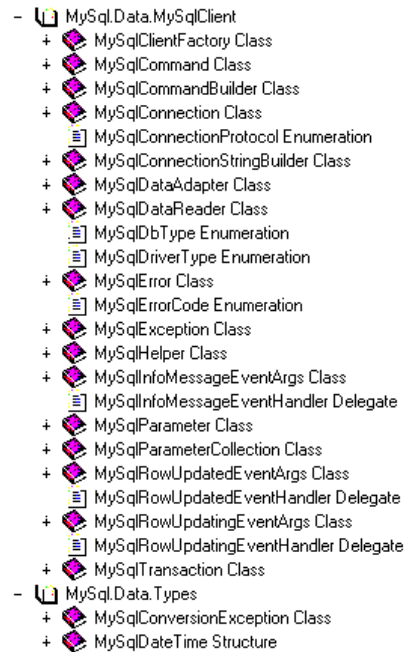
Connector/NET 5 is not graphical. For that you need the MySQL Visual Studio plugin [MySQL Tools for Microsoft Visual Studio](#), first released in September 2006. It does not yet support graphical DataSet design, but MySQL says it will eventually. Sadly then, point-and-click MySQL development in Visual Studio is not here yet. For some projects, that makes a good case for using SQL Express, or perhaps Connector/ODBC with MySQL, or SQL Server. If you are willing to write a little code, though, can MySQL cut it in the the .NET 2.0 environment?

We wrote *TheUsual* in PHP ([Chapter 12](#)) as a paging browser for any data table, or any master-detail pair of tables, in any MySQL database for which the user has appropriate privileges. It pages through any huge table; it supports row edits, inserts and deletes; and it supports finding a row on the master grid primary key. As a challenge can we write a more powerful version of *TheUsual* for MySQL Visual Studio 2005?

- page the detail grid too, so it can also handle huge tables;
- give the app automatic popup lookup browse windows for any column that is a foreign key referencing another table;
- provide a find dialog in each grid for finding any column value;
- tooltips;
- make the app DBMS-agnostic—encapsulate MySQL-specific code so the app can be taught to run against another DBMS simply by replacing MySQL modules with modules for another DBMS ( for example Sql Server 2005).

## The specification

*Login* should connect to a database server via on-the-fly authentication of username and password, or optionally via user-selectable persistent logins, and it must leave behind a connection object for use by the application. SQL Server 2005 can use Windows authentication but MySQL cannot, so the login module must be DBMS-specific. We are writing for developers and DBAs rather than end-users, so we delegate database, table and column privileges to login control. The connection object needs a DBMS-agnostic wrapper.



**Fig 22-1: Connector/NET classes**

*Databases:* Like *TheUsual* for PHP, the application should offer the user a list of available databases as reported by `information_schema`. Validation consists simply of setting the connection's default database to the user's choice.

*Master table:* Once a database is selected from a pick list, the application should offer a pick list of available physical tables or Views, again as reported by a query to `information_schema`.

*Detail table:* Once a master table is chosen from a pick list, the application should offer a list of detail tables having a foreign key that refers to the selected master table (another `information_schema` query). If a detail table is selected, the application should open coordinated master and detail grids for the two tables. Absent choice of a detail table, the application should simply browse the master table.

*Grids* should be paged if the rowcount justifies paging, should be quick regardless of table size, should offer clickable column sorting and searching.

*Add, edit, delete, lookup:* If a table has a primary key, browsing should permit adding new rows, adding rows seeded as copies of existing rows except for the primary key, deleting rows, and looking up foreign key values in popup browse windows. Popup lookup browses should be brisk even with large tables. They should support sorting on any column, and one-click selection. All data updates should be two-stage: (i) make the changes in the browser, (ii) commit them to the database. It must be possible to either undo stage (i) changes, or refresh the browser without saving them.

*Flexibility:* The application should permit the user to iteratively browse any available table, or table pair, simply by selecting databases and tables from pick lists.

*Scalability:* All this should work for any available table, large or small. Even for very large tables, paging or caching should provide crisp performance.

*Language:* C# is OO-friendlier than Visual Basic, and much easier to write and maintain than C++, so that's the choice.

## **Tools needed**

The login dialog can be a standard Windows Form with TextBoxes for server name, username and password, and a CheckBox to tell *theUsual* whether to remember connection parameters. *TheUsual* itself can also live in a standard Form. It doesn't need a menu (yet). Standard ComboBoxes are ideal for database and table selections. Ordinary Buttons will do fine for navigation, Go, Update and Exit interfaces. Visual Studio has a nice array of ToolStrip tools we can use for paging parameters. All these widgets can be dropped onto a form from the Visual Studio Toolbox, as is. It takes just a few minutes to name and arrange them as desired.

In previous versions of Visual Studio, the data browsing tool of choice was the DataGrid. VS 2005 introduced the DataGridView, with two enhancements that are particularly useful in a table browser:

- data binding can be generalised through use of a BindingSource, hiding the details and simplifying code for queries and for updates;
- in VirtualMode, the DataGridView can cache very large datasources such that accessing any part of the table is nearly instantaneous; implementing VirtualMode is not trivial, but not forbidding either;

Unfortunately a DataGridView can turn just one of these tricks at a time. A BindingSource streamlines update code very nicely, but caching has to be done serverside, for example via a LIMIT clause in the query. With VirtualMode, on the other hand, you lose update streamlining, but automatic clientside caching lets you forget all about paging.

Since the main browsing grids must support updates, they should use a BindingSource with serverside caching via a LIMIT clause. Lookup grids will be readonly so they can run in VirtualMode with a clientside cache. For performance comparison, we will implement an alternate lookup grid as a ListView.

The project thus needs only standard elements of VS 2005, MySQL and Connector/NET. The SQL Server version of course needs an installation of that DBMS.

## Class layout

Let the application namespace be `theUsual`, and let the principal class be a Windows Form subclassed as `TheUsual` with dropdown user input controls for selecting the database, master table and detail table; two data browsing grids; command buttons for filling the grids, for updating after edits, and for exiting from the program; a statusbar for general program messages; navigation buttons; and a toolstrip for paging parameters.

How to partition DBMS-agnostic and DBMS-specific code? DotNET has a set of DBMS-agnostic data classes (for example `DataTable`, `DataSet`, `BindingSource`) that mediate between DBMS-specific data classes (for example `MySqlConnection`, `MySqlCommand`, `MySqlDataAdapter`) and Windows controls (for example `ComboBox`, `DataGridView`). We will give *TheUsual* a thick DBMS-agnostic layer that talks to data-mediating classes, and a thin DBMS-specific layer.

How to implement that partition? One way would be to subclass all DBMS-specific functionalities. That will be useful here and there, but as a global approach it risks a blizzard of vexing cross-class references. The .NET concept of *partial class* offers a simpler overall solution. `Partial class foo` can exist in multiple files if each file declares it `Partial` in the same namespace. DotNET uses partial classes to separate code which it generates from code which you write. *TheUsual* will partition DBMS-specific and DBMS-agnostic code in the same way.

The specification also implies

- a DBMS-specific login class;
- a manager class to encapsulate some DBMS-specific grid management;
- an info class for passing database and table info to a lookup browser;
- for lookups, a class that implements a data-driven DataGridView running in VirtualMode; we can write that class generically but it will need a custom database-specific subclass, too;
- a class implementing a lookup as a ListView, so we can compare its capabilities and performance to the VirtualMode DataGridView, and
- a class to implement column find.

Some files will be common to each DBMS implementation. Does Visual Studio support sharing common files across multiple projects? Sure, if you purchase Microsoft Visual Source Safe (VSS). But that's another \$500. We won't assume you have a copy. Without VSS, files are shared across projects by copying them. Ugh. Are there Visual Studio tools for comparing source files? Sure, if you purchase the VS Team System for another \$500. We will jump off that bridge only if we have to; meanwhile we can use familiar text editors (e.g., *TextPad*) for that.

The overall class structure is:

```
namespace theUsual {
  class TheUsual {
    class DgvManager           • TheUsual
    class ListViewTheUsual    • manage DBMS specifics in grids
    class rowInfo             • popup lookup ListView browse window
    class TheUsualConnection • row info passed to called classes
  }                           • connection object hiding DBMS specifics
  class Cache                • clientside cache for class DgvFormJIT
  class DataRetriever        • DBMS-specific data retriever for Cache
  class DgvFormJIT           • Just-in-time cached lookup browse window
  class dgvInfo              • grid info for called classes
  class DgvRowFinder         • popup find dialogue
  class FindDlg              • dialog for searching for a column value
  class LoginForm            • DBMS-specific login dialog
  class ServerDlg            • dialog for choosing saved login
}
```

To read the rest of this and other chapters, [buy a copy of the book](#)

---

[TOC](#) [Previous](#) [Next](#)

---