
MySQL++ User Manual

Kevin Atkinson
Sinisa Milivojevic
Monty Widenius
Warren Young

Copyright © 1998-2001, 2005 Kevin Atkinson (original author) MySQL AB Educational Technology Resources
\$Date: 2005-08-18 11:40:34 -0600 (Thu, 18 Aug 2005) \$

Table of Contents

1. Introduction	2
1.1. A Brief History of MySQL++	2
1.2. If You Have Questions.....	2
2. Overview	2
2.1. The Connection Object	3
2.2. The Query Object	3
2.3. Result Sets	3
2.4. Exceptions	4
3. Tutorial	4
3.1. Running the Examples	4
3.2. A Simple Example	5
3.3. Examples' Utility Module	6
3.4. A More Complicated Example	9
3.5. Exceptions	10
3.6. Quoting and Escaping	11
3.7. Specialized SQL Structures	12
3.8. Handling SQL Nulls	17
3.9. Which Query Type to Use?	18
3.10. Getting Field Meta-Information	19
3.11. Let's Do Something Useful	20
4. Template Queries	23
4.1. Setting up template queries	25
4.2. Setting the parameters at execution time	26
4.3. Using defaults	26
4.4. Combining the two	27
4.5. Error Handling	27
5. Specialized SQL Structures	27
5.1. sql_create	28
5.2. sql_create with Compare	28
5.3. sql_create with Additional Initializers	29
5.4. An Important Limitation of sql_create	30
5.5. Additional Features of Specialized SQL Structures	31
5.6. Harnessing SSQSL Internals	31
5.7. Alternate Creation Methods	34
5.8. Expanding SSQSL Macros	34
5.9. Extending the SSQSL Mechanism	35
6. Using Unicode with MySQL++	35
6.1. A Short History of Unicode	35
6.2. Unicode and Unix	35
6.3. Unicode and Win32	36
6.4. For More Information	36
7. Incompatible Library Changes	37
7.1. API Changes	37
7.2. ABI Changes	39

8. Licensing	40
8.1. GNU Lesser General Public License	41
8.2. MySQL++ User Manual License	48

1. Introduction

MySQL++ is a powerful C++ wrapper for MySQL's C API. Its purpose is to make working with queries as easy as working with STL containers.

The latest version of MySQL++ can be found at the official web site.

Support for MySQL++ can be had on the mailing list. That page hosts the mailing list archives, and tells you how you can subscribe.

1.1. A Brief History of MySQL++

MySQL++ was created in 1998 by Kevin Atkinson. It started out MySQL-specific, but there were early efforts to try and make it database-independent, and call it SQL++. This is where the old library name "sqlplus" came from. This is also why the old versions prefixed some class names with "Mysql" but not others: the others were supposed to be the database-independent parts.

Then in 1999, Sinisa Milivojevic unofficially took over maintenance of the library, releasing versions 1.0 and 1.1. (All of Kevin's releases were pre-1.0 point releases.) Kevin gave over maintenance to Sinisa officially with 1.2, and ceased to have any involvement with the library's maintenance. Sinisa went on to maintain the library through 1.7.9, released in mid-2001. Since Sinisa is an employee of MySQL AB, it seems to be during this time that the dream of multiple-database compatibility died.

With version 1.7.9, MySQL++ went into a period of stasis, lasting over three years. During this time, Sinisa ran the MySQL++ mailing list and supported its users, but made no new releases. There were many patches submitted during this period, some of which were ignored, others which were just put on the MySQL++ web site for people to try. A lot of these patches were mutually-incompatible, and not all of them gave a fully-functional copy of MySQL++.

In early August of 2004, the current maintainer (Warren Young) got fed up with this situation and took over. He released 1.7.10 later that month.

1.2. If You Have Questions...

If you want to email someone to ask questions about this library, we greatly prefer that you send mail to the MySQL++ mailing list. The mailing list is archived, so if you have questions, do a search to see if the question has been asked before.

You may find people's individual email addresses in various files within the MySQL++ distribution. Please do not send mail to them unless you are sending something that is inherently personal. Not all of the principal developers of MySQL++ are still active in its development; those who have dropped out have no wish to be bugged about MySQL++. Those of us still active in MySQL++ development monitor the mailing list, so you aren't getting any extra "coverage" by sending messages to additional email addresses.

2. Overview

MySQL++ has developed into a very complex and powerful library, with many different ways to accomplish the same task. Unfortunately, this means that figuring out how to perform a simple task can be frustrating for new users. In this section we will provide an overview of the most important user-facing components of the library.

The overall process for using MySQL++ is similar to that of most other database access APIs:

1. Open the connection

2. Form and execute the query
3. Iterate through the result set
4. Go to 2 :)

There is, however, a lot of extra functionality along each step of the way.

2.1. The Connection Object

A Connection object manages the connection to the MySQL server. You need at least one of these objects to do anything. A Connection object can either create Query objects for you, or you can execute queries directly through the Connection object. The separate Query object is the recommended way as it gives you far more power.

2.2. The Query Object

A Query object is the recommended way of building and executing queries. It is subclassed from `std::stringstream` which means you can write to it like any other C++ stream to form a query. The library includes stream manipulators that make it easy to generate syntactically-correct SQL.

You can also set up Template Queries with this class. Template queries work something like the C `printf()` function: you set up a fixed query string with tags inside that indicate where to insert the variable parts. If you have multiple queries that are structurally similar, you simply set up one template query, and use that in the various locations of your program.

A third method for building queries is to use Specialized SQL Structures (SSQLS). This feature presents your results as a C++ data structure, instead of making you access the data through MySQL++ intermediary classes. It also reduces the amount of embedded SQL code your program needs.

2.3. Result Sets

The field data in a result set are stored in a special `std::string`-like class called ColData. This class has conversion operators that let you automatically convert these objects to any of the basic C data types. Additionally, MySQL++ defines classes like DateTime, which you can initialize from a MySQL DATETIME string. These automatic conversions are protected against bad conversions, and can either set a warning flag or throw an exception, depending on how you set the library up.

As for the result sets as a whole, MySQL++ has a number of different ways of representing them:

Queries That Do Not Return Data

Not all SQL queries return data. An example is `CREATE TABLE`. For these types of queries, there is a special result type (`ResNSel`) that simply reports the state resulting from the query: whether the query was successful, how many rows it impacted (if any), etc.

Queries That Return Data: Dynamic Method

The easiest way to retrieve data from MySQL uses a Result object, which includes one or more Row objects. Because these classes are `std::vector`-like containers, you can treat the result set as a two-dimensional array. For example, you can get the 5th item on the 2nd row by simply saying `result[1][4]`. You can also access row elements by field name, like this: `result[2].lookup_by_name("price")`.

An alternate way of accessing your query results is through a ResUse object. This class acts more like an STL input iterator than a container: you walk through your result set one item at a time, always going forward. You can't seek around in the result set, and you can't know how many results are in the set until you find the end. This method is more efficient when there can be arbitrarily many results, which could pose a memory allocation problem with the

previous technique.

Queries That Return Data: Static Method

The Specialized SQL Structures (SSQLS) feature method above defines C++ structures that match the table structures in your database schema.

We call it the "static" method because the table structure is fixed at compile time. Indeed, some schema changes require that you update your SSQLS definitions and recompile, or else the program could crash or throw "bad conversion" exceptions when MySQL++ tries to stuff the new data into an outdated data structure. (Not all changes require a recompile. Adding a column to a table is safe, for instance, as the program will ignore the new column until you update the SSQLS definition.)

The advantage of this method is that your program will require very little embedded SQL code. You can simply execute a query, and receive your results as C++ data structures, which can be accessed just as you would any other structure. The results can be accessed through the Row object, or you can ask the library to dump the results into a sequential or set-associative STL container for you. Consider this:

```
vector<mystruct> v;
Query q = connection.query();
q << "SELECT * FROM mytable";
q.storein(v);
for (vector<mystruct>::iterator it = v.begin(); it != v.end(); ++it) {
    cout << "Price: " << it->price << endl;
}
```

Isn't that slick?

2.4. Exceptions

By default, the library throws exceptions whenever it encounters an error. You can ask the library to set an error flag instead, if you like, but the exceptions carry more information. Not only do they include a string member telling you why the exception was thrown, there are several exception types, so you can distinguish between different error types within a single `try` block.

3. Tutorial

This tutorial is meant to give you a jump start in using MySQL++. While it is a very complicated and powerful library, it's possible to make quite functional programs without tapping but a fraction of its power. This section will introduce you to the most useful fraction.

This tutorial assumes you know C++ fairly well, in particular the Standard Template Library (STL) and exceptions.

3.1. Running the Examples

All of the examples are complete running programs. They may or may not be built for you already, depending on how you installed the library.

If you installed MySQL++ from the source tarball on a Unixy system, the examples should have been built along with the library. If not, simply go into the examples directory and type `make`.

If you installed the library via RPM, the examples are in the `mysql++-devel` RPM. After installing that, the examples are in `/usr/src/mysql++/examples`. To build them, go into that directory and type `make -f Makefile.simple`. See the file `/usr/share/doc/mysql++-devel*/README.examples` for more details.

If you are on a Windows system, the build process for the library should have built the examples as well. Where the programs are depends on which compiler you're using. There should be a `README.*` file in the distribution specific to your compiler with further instructions.

Once you have the examples building, you need to initialize the sample database by running the `resetdb` example. The usage of `resetdb` is as follows:

```
resetdb [host [user [password [port]]]]
```

If you leave off host, localhost is assumed. If you leave off user, your current username is assumed. If you leave off the password, it is assumed that you don't need one. And if you leave off the port, it will use the standard MySQL port number.

The user you give `resetdb` needs to be an account with permission to create databases. Once the database is created you can use any account that has full permission to the sample database `mysql_cpp_data`.

You may also have to re-run `resetdb` after running some of the other examples, as they change the database.

3.2. A Simple Example

The following example demonstrates how to open a connection, execute a simple query, and display the results. This is `examples/simple1.cpp`:

```
#include "util.h"
#include <mysqlpp.h>
#include <iostream>
#include <iomanip>

using namespace std;

int
main(int argc, char *argv[])
{
    // Connect to the sample database.
    mysqlpp::Connection con(false);
    if (!connect_to_db(argc, argv, con)) {
        return 1;
    }

    // Retrieve a subset of the sample stock table set up by resetdb
    mysqlpp::Query query = con.query();
    query << "select item from stock";
    mysqlpp::Result res = query.store();

    // Display the result set
    cout << "We have:" << endl;
    if (res) {
        char buf[100];
        mysqlpp::Row row;
        mysqlpp::Row::size_type i;
        for (i = 0; row = res.at(i); ++i) {
            cout << '\t' << utf8trans(row.at(0), buf, sizeof(buf)) << endl;
        }
    } else {
        cerr << "Failed to get item list: " << query.error() << endl;
        return 1;
    }

    return 0;
}
```

This example simply gets the entire "item" column from the example table, and prints those values out.

Notice that MySQL++ lets you store result sets in STL containers, such as `std::vector`. We iterate through the

result set just as you would with any other STL container. The only tricky bit is the `it->at(0)` idiom. This dereferences the iterator, which yields a Row object, on which we call the `at()` method to retrieve the first field.

The only thing that isn't explicitly handled in the code block above is that we delegate connection establishment to `connect_to_db()` in the `util` module. We do this only because that function also handles the command line parsing for the examples, so they have a consistent interface.

3.3. Examples' Utility Module

I referred to the `util` module above. Following is the source for that module, which also contains other functions used by other examples. It isn't important to understand this module in detail, but understanding its outlines will make the following examples more clear.

```
#include "util.h"

#include <iostream>
#include <iomanip>
#include <stdlib.h>

using namespace std;

const char* kpcSampleDatabase = "mysql_cpp_data";

// Converts a Unicode string encoded in UTF-8 form (which the MySQL
// database uses) to a form suitable for outputting to the standard C++
// cout stream. Functionality is platform-specific.

char*
utf8trans(const char* utf8_str, char* out_buf, int buf_len)
{
#ifdef MYSQLPP_PLATFORM_WINDOWS
    // It's Win32, so assume console output, where output needs to be in
    // local ANSI code page by default.
    wchar_t ucs2_buf[100];
    static const int ub_chars = sizeof(ucs2_buf) / sizeof(ucs2_buf[0]);

    // First, convert UTF-8 string to UCS-2
    if (MultiByteToWideChar(CP_UTF8, 0, utf8_str, -1,
                            ucs2_buf, ub_chars) > 0) {
        // Next, convert UCS-2 to local code page.
        CPINFOEX cpi;
        GetCPInfoEx(CP_OEMCP, 0, &cpi);
        WideCharToMultiByte(cpi.CodePage, 0, ucs2_buf, -1,
                            out_buf, buf_len, 0, 0);
        return out_buf;
    }
    else {
        int err = GetLastError();
        if (err == ERROR_NO_UNICODE_TRANSLATION) {
            cerr << "Bad data in UTF-8 string" << endl;
        }
        else {
            cerr << "Unknown error in Unicode translation: " <<
                GetLastError() << endl;
        }
        return 0;
    }
#else
    // Assume a modern Unixy platform, where the system's terminal I/O
    // code handles UTF-8 directly. (e.g. common Linux distributions
    // since 2001 or so, recent versions of Mac OS X, etc.)
    strncpy(out_buf, utf8_str, buf_len);
    return out_buf;
#endif
}
```

```
//// print_stock_header /////////////////////////////////
// Display a header suitable for use with print_stock_rows().

void
print_stock_header(int rows)
{
    cout << "Records found: " << rows << endl << endl;
    cout.setf(ios::left);
    cout << setw(21) << "Item" <<
        setw(10) << "Num" <<
        setw(10) << "Weight" <<
        setw(10) << "Price" <<
        "Date" << endl << endl;
}

//// print_stock_row /////////////////////////////////
// Print out a row of data from the stock table, in a format compatible
// with the header printed out in the previous function.

void
print_stock_row(const std::string& item, mysqlpp::longlong num,
                 double weight, double price, const mysqlpp::Date& date)
{
    // We do UTF-8 translation on item field because there is Unicode
    // data in this field in the sample database, and some systems
    // cannot handle UTF-8 sent directly to cout.
    char buf[100];
    cout << setw(20) << utf8trans(item.c_str(), buf, sizeof(buf)) << ' ' <<
        setw(9) << num << ' ' <<
        setw(9) << weight << ' ' <<
        setw(9) << price << ' ' <<
        date << endl;
}

//// print_stock_row /////////////////////////////////
// Take a Row from the example 'stock' table, break it up into fields,
// and call the above version of this function.

void
print_stock_row(const mysqlpp::Row& row)
{
    // The brief code below illustrates several aspects of the library
    // worth noting:
    //
    // 1. You can subscript a row by integer (position of the field in
    // the row) or by string (name of field in the row). The former is
    // more efficient, while the latter trades some efficiency for
    // robustness in the face of schema changes. (Consider using SSQLS
    // if you need a tradeoff in between these two positions.)
    //
    // 2. You can also get at a row's field's with Row::at(), which is
    // much like Row::operator[](int). Besides the syntax difference,
    // the only practical difference is that only at() can access field
    // 0: this is because '0' can be converted to both int and to const
    // char*, so the compiler rightly complains that it can't decide
    // which overload to call.
    //
    // 3. Notice that we make an explicit temporary copy of the first
    // field, which is the only string field. We must tolerate the
    // inefficiency of this copy, because Row::operator[] returns a
    // ColData object, which goes away after it is converted to some
    // other form. So, while we could have made print_stock_row()
    // take a const char* argument (as past versions mistakenly did!)
    // this would result in a dangling pointer, since it points into the
    // ColData object, which is dead by the time the pointer is
    // evaluated in print_stock_row(). It will probably even work this
    // way, but like any memory bug, it can wreak subtle havoc.
    std::string item(row.at(0));
```

```
    print_stock_row(item, row["num"], row[2], row[3], row[4]);  
}  
  
//// print_stock_rows ///////////////////////////////  
// Print out a number of rows from the example 'stock' table.  
  
void  
print_stock_rows(mysqlpp::Result& res)  
{  
    print_stock_header(res.size());  
  
    // Use the Result class's read-only random access iterator to walk  
    // through the query results.  
    mysqlpp::Result::iterator i;  
    for (i = res.begin(); i != res.end(); ++i) {  
        // Notice that a dereferenced result iterator can be converted  
        // to a Row object, which makes for easier element access.  
        print_stock_row(*i);  
    }  
}  
  
//// get_stock_table ///////////////////////////////  
// Retreive the entire contents of the example 'stock' table.  
  
void  
get_stock_table(mysqlpp::Query& query, mysqlpp::Result& res)  
{  
    // Reset the query object, in case we're re-using it.  
    query.reset();  
  
    // You can write to the query object like you would any ostream.  
    query << "select * from stock";  
  
    // Show the query string. If you call preview(), it must be before  
    // you call execute() or store() or use().  
    cout << "Query: " << query.preview() << endl;  
  
    // Execute the query, storing the results in memory.  
    res = query.store();  
}  
  
//// connect_to_db ///////////////////////////////  
// Establishes a connection to a MySQL database server, optionally  
// attaching to database kdb. This is basically a command-line parser  
// for the examples, since the example programs' arguments give us the  
// information we need to establish the server connection.  
  
bool  
connect_to_db(int argc, char *argv[], mysqlpp::Connection& con,  
             const char *kdb)  
{  
    if (argc < 1) {  
        cerr << "Bad argument count: " << argc << '!' << endl;  
        return false;  
    }  
  
    if (!kdb) {  
        kdb = kpcSampleDatabase;  
    }  
  
    if ((argc > 1) && (argv[1][0] == '-')) {  
        cout << "usage: " << argv[0] <<  
            " [host] [user] [password] [port]" << endl;  
        cout << endl << "\tConnects to database ";  
        if (strlen(kdb) > 0) {  
            cout << ' "' << kdb << "'";  
        }  
    }  
    else {
```

```
        cout << "server";
    }
    cout << " on localhost using your user" << endl;
    cout << "\tname and no password by default." << endl << endl;
    return false;
}

if (argc == 1) {
    con.connect(kdb);
}
else if (argc == 2) {
    con.connect(kdb, argv[1]);
}
else if (argc == 3) {
    con.connect(kdb, argv[1], argv[2]);
}
else if (argc == 4) {
    con.connect(kdb, argv[1], argv[2], argv[3]);
}
else if (argc >= 5) {
    con.connect(kdb, argv[1], argv[2], argv[3], atoi(argv[4]));
}

if (con) {
    return true;
}
else {
    cerr << "Database connection failed: " << con.error() << endl;
    return false;
}
}
```

This is actually an abridged version of util.cpp, with the Unicode stuff removed. The interaction between MySQL, MySQL++ and Unicode is covered in a later chapter, Using Unicode with MySQL++.

3.4. A More Complicated Example

The simple1 example above was pretty trivial. Let's get a little deeper. Here is examples/simple2.cpp:

```
#include "util.h"
#include <mysql++.h>

#include <iostream>
#include <iomanip>

using namespace std;

int
main(int argc, char *argv[])
{
    // Connect to the sample database.
    mysqlpp::Connection con(false);
    if (!connect_to_db(argc, argv, con)) {
        return 1;
    }

    // Retrieve the sample stock table set up by resetdb
    mysqlpp::Query query = con.query();
    query << "select * from stock";
    mysqlpp::Result res = query.store();

    // Display results
    if (res) {
        // Display header
        cout.setf(ios::left);
```

```
cout << setw(21) << "Item" <<
    setw(10) << "Num" <<
    setw(10) << "Weight" <<
    setw(10) << "Price" <<
    "Date" << endl << endl;

// Get each row in result set, and print its contents
char buf[100];
mysqlpp::Row row;
mysqlpp::Row::size_type i;
for (i = 0; row = res.at(i); ++i) {
    cout << setw(20) <<
        utf8trans(row["item"], buf, sizeof(buf)) << ' ' <<
        setw(9) << row["num"] << ' ' <<
        setw(9) << row["weight"] << ' ' <<
        setw(9) << row["price"] << ' ' <<
        setw(9) << row["sdate"] <<
        endl;
}
else {
    cerr << "Failed to get stock table: " << query.error() << endl;
    return 1;
}

return 0;
}
```

This example illustrates several new concepts.

First, notice that we store the result set in a `Result` object. Like the a `std::vector` we used in the `simple1` example, `Result` is a container type, so iterating through it is straightforward. The main difference is that `Result` is less distanced from the way the underlying MySQL C API works, so it is somewhat more efficient.

Second, we access each row's data indirectly through a `Row` object. This affords several nice features, such as the ability to access a field by name. You can also access fields by position, as we did in the `simple1` example, which is more efficient, but makes your code less flexible.

3.5. Exceptions

By default, MySQL++ uses exceptions to signal errors. Most of the examples have a full set of exception handlers. This is worthy of emulation.

All of MySQL++'s custom exceptions derive from a common base class, `Exception`. That in turn derives from the Standard C++ exception base class, `std::exception`. Since the library can indirectly cause exceptions to come from the Standard C++ Library, it's possible to catch all exceptions from MySQL++ by just catching `std::exception` by reference. However, it's usually better to catch the all of the concrete exception types that you expect, and add a handler for `Exception` or `std::exception` to act as a "catch-all" for unexpected exceptions.

Some of these exceptions are optional. When disabled, the object signals errors in some other way, typically by returning an error code or setting an error flag. Classes that support this feature derive from `OptionalExceptions`. Moreover, when such an object creates another object that also derives from this interface, it passes on its exception flag. Since everything flows from the `Connection` object, disabling exceptions on it at the start of the program disables all optional exceptions. You can see this technique at work in the "simple" examples, which keeps them, well, simple.

Real-world code typically can't afford to lose out on the additional information and control offered by exceptions. But at the same time, it is still sometimes useful to disable exceptions temporarily. To do this, put the section of code that you want to not throw exceptions inside a block, and create a `NoExceptions` object at the top of that block. When created, it saves the exception flag of the `OptionalExceptions` derivative you pass to it, and then disables exceptions on it. When the `NoExceptions` object goes out of scope at the end of the block, it restores the excep-

tions flag to its previous state. See `examples/resetdb.cpp` to see this technique at work.

When one `OptionalExceptions` derivative creates another such object and passes on its exception flag, it passes a copy of the flag. Therefore, the two objects' flags operate independently after the new one is created. There's no way to globally enable or disable this flag on existing objects in a single call.

There are a few classes of exceptions MySQL++ can throw that are not optional:

- The largest set of non-optional exceptions are those from the Standard C++ Library. For instance, if your code said "row[21]" on a row containing only 5 fields, the `std::vector` underlying the row object will throw an exception. (It will, that is, if it conforms to the standard.) You might consider wrapping your program's main loop in a try block catching `std::exceptions`, just in case you trigger one of these exceptions.
- `ColData` will always throw `BadConversion` when you ask it to do an improper type conversion. For example, you'll get an exception if you try to convert "1.25" to `int`, but not when you convert "1.00" to `int`. In the latter case, MySQL++ knows that it can safely throw away the fractional part.
- If you use template queries and don't pass enough parameters when instantiating the template, `Query` will throw a `BadParamCount` exception.

It's educational to modify the examples to force exceptions. For instance, misspell a field name, use an out-of-range index, or change a type to force a `ColData` conversion error.

3.6. Quoting and Escaping

SQL syntax often requires certain data to be quoted. Consider this query:

```
SELECT * FROM stock WHERE item = 'Hotdog Buns'
```

Because the string "Hotdog Buns" contains a space, it must be quoted. With MySQL++, you don't have to add these quote marks manually:

```
string s = "Hotdog Buns";
Query q = conn.query();
q << "SELECT * FROM stock WHERE item = " << quote_only << s;
```

That code produces the same query string as in the previous example. We used the MySQL++ `quote_only` manipulator, which causes single quotes to be added around the next item inserted into the stream. This works for various string types, for any type of data that can be converted to MySQL++'s `ColData` type, and for Specialized SQL Structures. (The next section introduces the SSQLS feature.)

Quoting is pretty simple, but SQL syntax also often requires that certain characters be "escaped". Imagine if the string in the previous example was "Frank's Brand Hotdog Buns" instead. The resulting query would be:

```
SELECT * FROM stock WHERE item = 'Frank's Brand Hotdog Buns'
```

That's not valid SQL syntax. The correct syntax is:

```
SELECT * FROM stock WHERE item = 'Frank''s Brand Hotdog Buns'
```

As you might expect, MySQL++ provides that feature, too, through its `escape` manipulator. But here, we want both quoting and escaping. That brings us to the most widely useful manipulator:

```
string s = "Frank's Brand Hotdog Buns";
Query q = conn.query();
q << "SELECT * FROM stock WHERE item = " << quote << s;
```

The quote manipulator both quotes strings, and escapes any characters that are special in SQL.

3.7. Specialized SQL Structures

Retrieving data

The next example introduces one of the most powerful features of MySQL++: Specialized SQL Structures (SSQLS). This is `examples/custom1.cpp`:

```
#include "util.h"

#include <mysql++.h>
#include <custom.h>

#include <iostream>
#include <iomanip>
#include <vector>

using namespace std;
using namespace mysqlpp;

// The following is calling a very complex macro which will create
// "struct stock", which has the member variables:
//
//   string item
//   ...
//   Date sdate
//
// plus methods to help populate the class from a MySQL row
// among other things that I'll get to in a later example.
sql_create_5(stock,
              1, 5,                      // explained in the user manual
              string, item,
              longlong, num,
              double, weight,
              double, price,
              Date, sdate)

int
main(int argc, char *argv[])
{
    // Wrap all MySQL++ interactions in one big try block, so any
    // errors are handled gracefully.
    try {
        // Establish the connection to the database server.
        Connection con(use_exceptions);
        if (!connect_to_db(argc, argv, con)) {
            return 1;
        }

        // Retrieve the entire contents of the stock table, and store
        // the data in a vector of 'stock' SSQLS structures.
        Query query = con.query();
        query << "select * from stock";
        vector<stock> res;
        query.storein(res);

        // Display the result set
        print_stock_header(res.size());
        vector<stock>::iterator it;
        for (it = res.begin(); it != res.end(); ++it) {
            print_stock_row(it->item, it->num, it->weight, it->price,
                            it->sdate);
        }
    }
    catch (const BadQuery& er) {
        // Handle any query errors
        cerr << "Query error: " << er.what() << endl;
    }
}
```

```
        return -1;
    }
    catch (const BadConversion& er) {
        // Handle bad conversions; e.g. type mismatch populating 'stock'
        cerr << "Conversion error: " << er.what() << endl <<
            "\tretrieved data size: " << er.retrieved <<
            ", actual size: " << er.actual_size << endl;
        return -1;
    }
    catch (const Exception& er) {
        // Catch-all for any other MySQL++ exceptions
        cerr << "Error: " << er.what() << endl;
        return -1;
    }
}
return 0;
}
```

As you can see, SSQLS is very powerful.

Adding data

SSQLS can also be used to add data to a table. This is examples/custom2.cpp:

```
#include "util.h"

#include <mysql++.h>
#include <custom.h>

#include <iostream>
#include <string>
#include <vector>

using namespace std;
using namespace mysqlpp;

sql_create_5(stock,
    1, 5,
    string, item,
    longlong, num,
    double, weight,
    double, price,
    Date, sdate)

int
main(int argc, char *argv[])
{
    try {
        // Establish the connection to the database server.
        Connection con(use_exceptions);
        if (!connect_to_db(argc, argv, con)) {
            return 1;
        }

        // Create and populate a stock object.  We could also have used
        // the set() member, which takes the same parameters as this
        // constructor.
        stock row("Hot Dogs", 100, 1.5, 1.75, "1998-09-25");

        // Form the query to insert the row into the stock table.
        Query query = con.query();
        query.insert(row);

        // Show the query about to be executed.
        cout << "Query: " << query.preview() << endl;

        // Execute the query.  We use execute() because INSERT doesn't
```

```
// return a result set.
query.execute();

// Print the new table.
Result res;
get_stock_table(query, res);
print_stock_rows(res);
}
catch (const BadQuery& er) {
    // Handle any query errors
    cerr << "Query error: " << er.what() << endl;
    return -1;
}
catch (const BadConversion& er) {
    // Handle bad conversions
    cerr << "Conversion error: " << er.what() << endl <<
        "\tretrieved data size: " << er.retrieved <<
        ", actual size: " << er.actual_size << endl;
    return -1;
}
catch (const Exception& er) {
    // Catch-all for any other MySQL++ exceptions
    cerr << "Error: " << er.what() << endl;
    return -1;
}
return 0;
}
```

That's all there is to it!

There is one subtlety: MySQL++ automatically quotes and escapes the data when building SQL queries using SSQLS structures. It's efficient, too: MySQL++ is smart enough to apply quoting and escaping only for those data types that actually require it.

Because this example modifies the sample database, you may want to run resetdb after running this program.

Modifying data

It almost as easy to modify data with SSQLS. This is examples/custom3.cpp:

```
#include "util.h"

#include <mysql++.h>
#include <custom.h>

#include <iostream>
#include <string>
#include <vector>

using namespace std;
using namespace mysqlpp;

sql_create_5(stock,
    1, 5,
    string, item,
    longlong, num,
    double, weight,
    double, price,
    Date, sdate)

int
main(int argc, char *argv[])
{
    try {
        // Establish the connection to the database server.
```

```
Connection con(use_exceptions);
if (!connect_to_db(argc, argv, con)) {
    return 1;
}

// Build a query to retrieve the stock item that has Unicode
// characters encoded in UTF-8 form.
Query query = con.query();
query << "select * from stock where item = \"Nürnberger Brats\";

// Retrieve the row, throwing an exception if it fails.
Result res = query.store();
if (res.empty()) {
    throw BadQuery("UTF-8 bratwurst item not found in "
                  "table, run resetdb");
}

// Because there should only be one row in the result set,
// there's no point in storing the result in an STL container.
// We can store the first row directly into a stock structure
// because one of an SSQSL's constructors takes a Row object.
stock row = res.at(0);

// Create a copy so that the replace query knows what the
// original values are.
stock orig_row = row;

// Change the stock object's item to use only 7-bit ASCII, and
// to deliberately be wider than normal column widths printed
// by print_stock_table().
row.item = "Nuerenberger Bratwurst";

// Form the query to replace the row in the stock table.
query.update(orig_row, row);

// Show the query about to be executed.
cout << "Query: " << query.preview() << endl;

// Run the query with execute(), since UPDATE doesn't return a
// result set.
query.execute();

// Print the new table contents.
get_stock_table(query, res);
print_stock_rows(res);
}
catch (const BadQuery& er) {
    // Handle any query errors
    cerr << "Query error: " << er.what() << endl;
    return -1;
}
catch (const BadConversion& er) {
    // Handle bad conversions
    cerr << "Conversion error: " << er.what() << endl <<
        "\tretrieved data size: " << er.retrieved <<
        ", actual size: " << er.actual_size << endl;
    return -1;
}
catch (const Exception& er) {
    // Catch-all for any other MySQL++ exceptions
    cerr << "Error: " << er.what() << endl;
    return -1;
}

return 0;
}
```

When you run the example you will notice that in the WHERE clause only the 'item' field is checked for. This is because SSQSL also also less-than-comparable.

Don't forget to run resetdb after running the example.

Less-than-comparable

SSQLS structures can be sorted and stored in STL associative containers as demonstrated in the next example. This is examples/custom4.cpp:

```
#include "util.h"

#include <mysql++.h>
#include <custom.h>

#include <iostream>
#include <iomanip>
#include <string>
#include <vector>

using namespace std;
using namespace mysqlpp;

sql_create_5(stock,
 1, // This number is used to make a SSQLS less-than-comparable.
 // When comparing two SSQLS structures, the first N elements are
 // compared. In this instance, we are saying that we only want
 // the first element ('item') to be used when comparing two
 // stock structures.

 5, // Each SSQLS structure includes a number of constructors. Some
 // of these are fixed in nature, but one of these will have this
 // number of arguments, one for each of the first N elements in
 // the structure; it is an initialization ctor. Since N is the
 // same as the number of structure elements in this instance,
 // that ctor will be able to fully initialize the structure. This
 // behavior is not always wanted, however, so the macro allows
 // you make the constructor take fewer parameters, leaving the
 // remaining elements uninitialized. An example of when this is
 // necessary is when you have a structure containing only two
 // integer elements: one of the other ctors defined for SSQLS
 // structures takes two ints, so the compiler barfs if you pass
 // 2 for this argument. You would need to pass 0 here to get
 // that SSQLS structure to compile.
 string, item,
 longlong, num,
 double, weight,
 double, price,
 Date, sdate)

int
main(int argc, char *argv[])
{
    try {
        // Establish the connection to the database server.
        Connection con(use_exceptions);
        if (!connect_to_db(argc, argv, con)) {
            return 1;
        }

        // Retrieve all rows from the stock table and put them in an
        // STL set. Notice that this works just as well as storing them
        // in a vector, which we did in custom1.cpp. It works because
        // SSQLS objects are less-than comparable.
        Query query = con.query();
        query << "select * from stock";
        set<stock> res;
        query.storein(res);

        // Display the result set. Since it is an STL set and we set up
        // the SSQLS to compare based on the item column, the rows will
        // be sorted by item.
    }
}
```

```
print_stock_header(res.size());
set<stock>::iterator it;
cout.precision(3);
for (it = res.begin(); it != res.end(); ++it) {
    print_stock_row(it->item.c_str(), it->num, it->weight,
                    it->price, it->sdate);
}

// Use set's find method to look up a stock item by item name.
// This also uses the SSQLS comparison setup.
it = res.find(stock("Hotdog Buns"));
if (it != res.end()) {
    cout << endl << "Currently " << it->num <<
        " hotdog buns in stock." << endl;
}
else {
    cout << endl << "Sorry, no hotdog buns in stock." << endl;
}
} catch (const BadQuery& er) {
    // Handle any query errors
    cerr << "Query error: " << er.what() << endl;
    return -1;
}
catch (const BadConversion& er) {
    // Handle bad conversions
    cerr << "Conversion error: " << er.what() << endl <<
        "\tretrieved data size: " << er.retrieved <<
        ", actual size: " << er.actual_size << endl;
    return -1;
}
catch (const Exception& er) {
    // Catch-all for any other MySQL++ exceptions
    cerr << "Error: " << er.what() << endl;
    return -1;
}
return 0;
}
```

For more details on the SSQLS feature, see the Specialized SQL Structures chapter.

3.8. Handling SQL Nulls

There is no equivalent of SQL's null in the standard C++ type system.

The primary distinction is one of type: in SQL, null is a column attribute, which affects whether that column can hold a SQL null. This effectively doubles the number of types in SQL. MySQL++ handles this the same way SQL does: it provides the `Null` template to allow the creation of distinct "nullable" versions of existing C++ types. For each column type MySQL understands, the library instantiates this template for the closest C++ type. (See the top of `lib/type_info.cpp` for the list.)

Template instantiations are first-class types in the C++ language, on par with any other type. You can use nullable MySQL++ types anywhere you'd use the plain version of that type, you can assign plain values to a nullable object and vice versa, etc.

There's a secondary distinction between SQL null and anything available in the standard C++ type system: SQL null is a distinct value, equal to nothing else. C++'s `NULL` is ambiguous, being equal to 0 in integer context, so MySQL++ has a global `null` object which you can assign to any nullable object to get a SQL null.

By default, if you try to convert a SQL null to a plain C++ data type, MySQL++ will throw a `BadNullConversion` exception to enforce this distinction. If you insert a SQL null into a C++ stream, you get "(NULL)". The `NullisNull` "behavior" type encapsulates these two rules; it is the default for one of template `Null`'s parameters. To relax this distinction, you can instantiate the `Null` template with a different behavior type: `NullisZero` or `NullisBlank`. As you

might guess from their names, SQL nulls using these behaviors get converted to 0 or a blank C string, respectively.

3.9. Which Query Type to Use?

There are three major ways to execute a query in MySQL++: `Query::execute()`, `Query::store()`, and `Query::use()`. Which should you use, and why?

`execute()` is for queries that do not return data *per se*. For instance, `CREATE INDEX`. You do get back some information from the MySQL server, which `execute()` returns to its caller in a `ResNSel` object. In addition to the obvious — a flag stating whether the query succeeded or not — this object also contains things like the number of rows that the query affected. If you only need the success status, there's `Query::exec()`, which just returns `bool`.

If your query does pull data from the database, the simplest option is `store()`. This returns a `Result` object, which contains an in-memory copy of the result set. The nice thing about this is that `Result` is a sequential container, like `std::vector`, so you can iterate through it forwards and backwards, access elements with subscript notation, etc. There are also the `storein()` methods, which actually put the result set into an STL container of your choice. The downside of these methods is that a sufficiently large result set will give your program memory problems.

For these large result sets, the superior option is a `use()` query. This returns a `ResUse` object, which is similar to `Result`, but without all of the random-access features. This is because a "use" query tells the database server to send the results back one row at a time, to be processed linearly. It's analogous to a C++ stream's input iterator, as opposed to a random-access iterator that a container like `vector` offers. By accepting this limitation, you can process arbitrarily large result sets. This technique is demonstrated in `examples/simple3.cpp`:

```
#include "util.h"
#include <mysql++.h>
#include <iostream>
#include <iomanip>

using namespace std;

int
main(int argc, char *argv[])
{
    // Connect to the sample database.
    mysqlpp::Connection con(false);
    if (!connect_to_db(argc, argv, con)) {
        return 1;
    }

    // Ask for all rows from the sample stock table set up by resetdb.
    // Unlike simple2 example, we don't store result set in memory.
    mysqlpp::Query query = con.query();
    query << "select * from stock";
    mysqlpp::ResUse res = query.use();

    // Retreive result rows one by one, and display them.
    if (res) {
        // Display header
        cout.setf(ios::left);
        cout << setw(21) << "Item" <<
            setw(10) << "Num" <<
            setw(10) << "Weight" <<
            setw(10) << "Price" <<
            "Date" << endl << endl;

        // Get each row in result set, and print its contents
        mysqlpp::Row row;
        while (row = res.fetch_row()) {
            cout << setw(20) << row["item"] << ' ' <<
                setw(9) << row["num"] << ' ' <<
                setw(9) << row["weight"] << ' ' <<
                setw(9) << row["price"] << ' ' <<
                setw(9) << row["sdate"] <<
```

```
        endl;
    }

    return 0;
}
else {
    cerr << "Failed to get stock item: " << query.error() << endl;
    return 1;
}
}
```

This example does the same thing as `simple2`, only with a "use" query instead of a "store" query. If your program uses exceptions, you should instead look at `examples/usequery.cpp`, which does the same thing as `simple`, but with exception-awareness.

3.10. Getting Field Meta-Information

The following example demonstrates how to get information about the fields in a result set, such as the name of the field and the SQL type. This is `examples/fieldinfo1.cpp`:

```
#include "util.h"

#include <mysql++.h>

#include <iostream>
#include <iomanip>

using namespace std;
using namespace mysqlpp;

int
main(int argc, char *argv[])
{
    try {
        Connection con(use_exceptions);
        if (!connect_to_db(argc, argv, con)) {
            return 1;
        }

        Query query = con.query();
        query << "select * from stock";
        cout << "Query: " << query.preview() << endl;

        Result res = query.store();
        cout << "Records Found: " << res.size() << endl << endl;

        cout << "Query Info:\n";
        cout.setf(ios::left);

        for (unsigned int i = 0; i < res.names().size(); i++) {
            cout << setw(2) << i
                // this is the name of the field
                << setw(15) << res.names(i).c_str()
                // this is the SQL identifier name
                // Result::types(unsigned int) returns a mysql_type_info which in many
                // ways is like type_info except that it has additional sql type
                // information in it. (with one of the methods being sql_name())
                << setw(15) << res.types(i).sql_name()
                // this is the C++ identifier name which most closely resembles
                // the sql name (its is implementation defined and often not very readable)
                << setw(20) << res.types(i).name()
                << endl;
        }

        cout << endl;
    }
}
```

```
if (res.types(0) == typeid(string)) {
    // this is demonstrating how a mysql_type_info can be
    // compared with a C++ type_info.
    cout << "Field 'item' is of an SQL type which most "
        "closely resembles\nthe C++ string type\n";
}

if (res.types(1) == typeid(longlong)) {
    cout << "Field 'num' is of an SQL type which most "
        "closely resembles\nC++ long long int type\n";
}
else if (res.types(1).base_type() == typeid(longlong)) {
    // you have to be careful as if it can be null the actual
    // type is Null<TYPE> not TYPE. So you should always use
    // the base_type method to get at the underlying type.
    // If the type is not null than this base type would be
    // the same as its type.
    cout << "Field 'num' base type is of an SQL type which "
        "most closely\nresembles the C++ long long int type\n";
}

catch (const BadQuery& er) {
    // Handle any query errors
    cerr << "Query error: " << er.what() << endl;
    return -1;
}
catch (const BadConversion& er) {
    // Handle bad conversions
    cerr << "Conversion error: " << er.what() << endl <<
        "\tretrieved data size: " << er.retrieved <<
        ", actual size: " << er.actual_size << endl;
    return -1;
}
catch (const Exception& er) {
    // Catch-all for any other MySQL++ exceptions
    cerr << "Error: " << er.what() << endl;
    return -1;
}

return 0;
}
```

3.11. Let's Do Something Useful

These next few examples demonstrate just how powerful C++ can be, allowing you to do a lot of work in few lines of code without losing efficiency.

Since the code is meant to be re-used as-is, constants that can differ from one case to another have been grouped in order to simplify editing. Also, all of these examples have full error checking code, showing off the power of MySQL++'s exception handling features.

Loading binary file in a BLOB column

Since MySQL 3.23, BLOB columns have been available, but their use is sometimes not straightforward. Besides showing how easy it can be with MySQL++, this example demonstrates several features of MySQL++. The program requires one command line parameter, which is a full path to the binary file. This is `examples/load_file.cpp`:

```
#include <mysql++.h>
#include <sys/stat.h>
#include <fstream>
#include <stdlib.h>
```

```
using namespace std;
using namespace mysqlpp;

const char MY_DATABASE[] = "telcent";
const char MY_TABLE[] = "fax";
const char MY_HOST[] = "localhost";
const char MY_USER[] = "root";
const char MY_PASSWORD[] = "";
const char MY_FIELD[] = "fax"; // BLOB field

int
main(int argc, char *argv[])
{
    if (argc < 2) {
        cerr << "Usage : load_file full_file_path" << endl << endl;
        return -1;
    }

    Connection con(use_exceptions);
    try {
        con.connect(MY_DATABASE, MY_HOST, MY_USER, MY_PASSWORD);
        Query query = con.query();
        ostringstream strbuf;
        ifstream In(argv[1], ios::in | ios::binary);
        struct stat for_len;
        if ((In.rdbuf())->is_open()) {
            if (stat(argv[1], &for_len) == -1)
                return -1;
            unsigned int blen = for_len.st_size;
            if (!blen)
                return -1;
            char *read_buffer = new char[blen];
            In.read(read_buffer, blen);
            string fill(read_buffer, blen);
            strbuf << "INSERT INTO " << MY_TABLE << " (" << MY_FIELD <<
                ") VALUES(\"" << mysqlpp::escape << fill << "\")" << endl;
            query.exec(strbuf.str());
            delete[] read_buffer;
        }
        else
            cerr << "Failed to open " << argv[1] <<
                '.' << endl;
    }
    catch (const BadQuery& er) {
        // Handle any query errors
        cerr << "Query error: " << er.what() << endl;
        return -1;
    }
    catch (const BadConversion& er) {
        // Handle bad conversions
        cerr << "Conversion error: " << er.what() << endl <<
            "\tretrieved data size: " << er.retrieved <<
            ", actual size: " << er.actual_size << endl;
        return -1;
    }
    catch (const Exception& er) {
        // Catch-all for any other MySQL++ exceptions
        cerr << "Error: " << er.what() << endl;
        return -1;
    }
    return 0;
}
```

Notice that we used the `escape` manipulator when building the `INSERT` query above. This is because we're not using one of the MySQL++ types that does automatic escaping and quoting.

Displaying images in HTML from BLOB column

This example is also a very short one, considering the function that it performs. Although all modern versions of MySQL have a command that dumps data from a BLOB column to a binary file, this example shows how to do it in your code instead, without requiring a temporary file on disk. This is `examples/cgi_image.cpp`:

```
#include <mysql++.h>

using namespace std;
using namespace mysqlpp;

#define MY_DATABASE "telcent"
#define MY_TABLE     "fax"
#define MY_HOST      "localhost"
#define MY_USER      "root"
#define MY_PASSWORD  ""
#define MY_FIELD     "fax"          // BLOB field
#define MY_KEY       "datet"        // PRIMARY KEY

int
main(int argc, char *argv[])
{
    if (argc < 2) {
        cerr << "Usage : cgi_image primary_key_value" << endl << endl;
        return -1;
    }

    cout << "Content-type: image/jpeg" << endl;
    Connection con(use_exceptions);
    try {
        con.connect(MY_DATABASE, MY_HOST, MY_USER, MY_PASSWORD);
        Query query = con.query();
        query << "SELECT " << MY_FIELD << " FROM " << MY_TABLE << " WHERE "
              << MY_KEY << " = " << argv[1];
        ResUse res = query.use();
        Row row = res.fetch_row();
        long unsigned int *jj = res.fetch_lengths();
        cout << "Content-length: " << *jj << endl << endl;
        fwrite(row.raw_data(0), 1, *jj, stdout);
    }
    catch (const BadQuery& er) {
        // Handle any query errors
        cerr << "Query error: " << er.what() << endl;
        return -1;
    }
    catch (const Exception& er) {
        // Catch-all for any other MySQL++ exceptions
        cerr << "Error: " << er.what() << endl;
        return -1;
    }
    return 0;
}
```

DELETE or UPDATE from SELECT

MySQL's SELECT statement has more power to winnow out just the items of interest from the database than do DELETE or UPDATE queries. Therefore, many people have wanted the ability to execute a SELECT statement that in fact deletes or updates the rows matched, rather than returning them. This example implements that feature in just a few lines of code. It is `examples/updel.cpp`:

```
#include <mysql++.h>
#include <string>
```

```
using namespace std;
using namespace mysqlpp;

#define MY_DATABASE "telcent"
#define MY_TABLE     "nazivi"
#define MY_HOST      "localhost"
#define MY_USER      "root"
#define MY_PASSWORD   ""
#define MY_FIELD     "naziv"
#define MY_QUERY     "SELECT URL from my_table as t1, my_table as t2 where t1.field = t2.field"

int
main()
{
    Connection con(use_exceptions);
    try {
        ostringstream strbuf;
        unsigned int i = 0;
        con.connect(MY_DATABASE, MY_HOST, MY_USER, MY_PASSWORD);
        Query query = con.query();
        query << MY_QUERY;
        ResUse res = query.use();
        Row row;
        strbuf << "delete from " << MY_TABLE << " where " << MY_FIELD <<
            " in (";
        // for UPDATE just replace the above DELETE FROM with UPDATE statement
        for (; row = res.fetch_row(); i++)
            strbuf << row.at(0) << ",";
        if (!i)
            return 0;
        string output(strbuf.str());
        output.erase(output.size() - 1, 1);
        output += ")";
        query.exec(output);
        //cout << output << endl;
    }
    catch (const BadQuery& er) {
        // Handle any query errors
        cerr << "Query error: " << er.what() << endl;
        return -1;
    }
    catch (const BadConversion& er) {
        // Handle bad conversions
        cerr << "Conversion error: " << er.what() << endl <<
            "\tretrieved data size: " << er.retrieved <<
            ", actual size: " << er.actual_size << endl;
        return -1;
    }
    catch (const Exception& er) {
        // Catch-all for any other MySQL++ exceptions
        cerr << "Error: " << er.what() << endl;
        return -1;
    }
    return 0;
}
```

Notice that the row values used in the IN clause aren't escaped or quoted. This is because row elements are ColData types, so they have automatic escaping and quoting, as appropriate to the type being inserted. If you want to disable this feature, it's easily done: click the ColData link for the details.

Users of this example should beware that one more check is required in order to run this query safely: in some extreme cases, the size of the query might grow larger than MySQL's maximum allowed packet size. This check should be added.

4. Template Queries

Another powerful feature of MySQL++ is being able to set up template queries. These are kind of like C's `printf()` facility: you give MySQL++ a string containing the fixed parts of the query and placeholders for the variable parts, and you can later substitute in values into those placeholders.

The following program demonstrates how to use this feature. This is `examples/resetdb.cpp`, the program you've run a few times now if you've worked through all the examples:

```
#include "util.h"
#include <mysql++.h>
#include <iostream>
using namespace std;

int
main(int argc, char *argv[])
{
    // Connect to database server
    mysqlpp::Connection con;
    try {
        cout << "Connecting to database server..." << endl;
        connect_to_db(argc, argv, con, "");
    }
    catch (exception& er) {
        cerr << "Connection failed: " << er.what() << endl;
        return 1;
    }

    // Create new sample database, or re-create it.  We suppress
    // exceptions, because it's not an error if DB doesn't yet exist.
    bool new_db = false;
    {
        mysqlpp::NoExceptions ne(con);
        mysqlpp::Query query = con.query();
        if (con.select_db(kpcSampleDatabase)) {
            // Toss old table, if it exists.  If it doesn't, we don't
            // really care, as it'll get created next.
            cout << "Dropping existing stock table..." << endl;
            query.execute("drop table stock");
        }
        else {
            // Database doesn't exist yet, so create and select it.
            if (con.create_db(kpcSampleDatabase) &&
                con.select_db(kpcSampleDatabase)) {
                new_db = true;
            }
            else {
                cerr << "Error creating DB: " << con.error() << endl;
                return 1;
            }
        }
    }

    // Create sample data table within sample database.
    cout << "Creating new stock table..." << endl;
    try {
        // Send the query to create the table and execute it.
        mysqlpp::Query query = con.query();
        query << "create table stock (item char(20) not null, "
                 "num bigint, weight double, price double, sdate date)";
        query.execute();

        // Set up the template query to insert the data.  The parse()
        // call tells the query object that this is a template and
        // not a literal query string.
        query << "insert into %5:table values (%0q, %1q, %2, %3, %4q)";
        query.parse();
    }
}
```

```
// Set the template query parameter "table" to "stock".
query.def["table"] = "stock";

// Notice that we don't give a sixth parameter in these calls,
// so the default value of "stock" is used. Also notice that
// the first row is a UTF-8 encoded Unicode string! All you
// have to do to store Unicode data in recent versions of MySQL
// is use UTF-8 encoding.
cout << "Populating stock table..." << endl;
query.execute("N  rnberger Brats", 92, 1.5, 8.79, "2005-03-10");
query.execute("Pickle Relish", 87, 1.5, 1.75, "1998-09-04");
query.execute("Hot Mustard", 75, .95, .97, "1998-05-25");
query.execute("Hotdog Buns", 65, 1.1, 1.1, "1998-04-23");

cout << (new_db ? "Created" : "Reinitialized") <<
      " sample database successfully." << endl;
}

catch (const mysqlpp::BadQuery& er) {
    // Handle any query errors
    cerr << "Query error: " << er.what() << endl;
    return 1;
}
catch (const mysqlpp::BadConversion& er) {
    // Handle bad conversions
    cerr << "Conversion error: " << er.what() << endl <<
        "\tretrieved data size: " << er.retrieved <<
        ", actual size: " << er.actual_size << endl;
    return 1;
}
catch (const mysqlpp::Exception& er) {
    // Catch-all for any other MySQL++ exceptions
    cerr << "Error: " << er.what() << endl;
    return 1;
}

return 0;
}
```

The line just before the call to `query.parse()` sets the template, and the `parse` call puts it into effect. From that point on, you can re-use this query by calling any of several `Query` member functions that accept query template parameters. In this example, we're using `Query::execute()`.

Let's dig into this feature a little deeper.

4.1. Setting up template queries

To set up a template query, you simply insert it into the `Query` object, using numbered placeholders wherever you want to be able to change the query. Then, you call the `parse()` function to tell the `Query` object that the query string is a template query, and it needs to parse it:

```
query << "select (%2:field1, %3:field2) from stock where %1:wheref = %0q:what";
query.parse();
```

The format of the placeholder is:

```
%##(modifier)(:name)(:)
```

Where '##' is a number up to three digits. It is the order of parameters given to a `SQLQueryParms` object, starting from 0.

'modifier' can be any one of the following:

%	Print an actual "%"
---	---------------------

""	Don't quote or escape no matter what.
q	This will quote and escape the item using the MySQL C API function mysql-escape-string if it is a string or char *, or another MySQL-specific type that needs to be quoted.
Q	Quote but don't escape based on the same rules as for 'q'. This can save a bit of processing time if you know the strings will never need quoting
r	Always quote and escape even if it is a number.
R	Always quote but don't escape even if it is a number.

"`:name`" is for an optional name which aids in filling SQLQueryParams. Name can contain any alpha-numeric characters or the underscore. You can have a trailing colon, which will be ignored. If you need to represent an actual colon after the name, follow the name with two colons. The first one will end the name and the second one won't be processed.

4.2. Setting the parameters at execution time

To specify the parameters when you want to execute a query simply use `Query::store(const SQLString &parm0, [..., const SQLString &parm11])`. This type of multiple overload also exists for `Query::use()` and `Query::execute()`. '`parm0`' corresponds to the first parameter, etc. You may specify up to 12 parameters. For example:

```
Result res = query.store("Dinner Rolls", "item", "item", "price")
```

with the template query provided above would produce:

```
select (item, price) from stock where item = "Dinner Rolls"
```

The reason we didn't put the template parameters in numeric order...

```
select (%0:field1, %1:field2) from stock where %2:wheref = %3q:what
```

...will become apparent shortly.

4.3. Using defaults

You can also set the parameters one at a time by means of class Query's public data member `def`. To change the values of the `def`, simply use the subscript operator. You can refer to the parameters either by number or by name. The following two examples have the same effect:

```
query.def[0] = "Dinner Rolls";
query.def[1] = "item";
query.def[2] = "item";
query.def[3] = "price";
```

and

```
query.def["what"] = "Dinner Rolls";
```

```
query.def["whereref"] = "item";
query.def["field1"] = "item";
query.def["field2"] = "price";
```

Once all the parameters are set simply execute as you would have executed the query before you knew about template queries:

```
Result res = query.store()
```

4.4. Combining the two

You can also combine the use of setting the parameters at execution time and setting them via the def object by calling `Query::store()` (or `use()` or `execute()`) without passing the full number of parameters that the template supports:

```
query.def["field1"] = "item";
query.def["field2"] = "price";
Result res1 = query.store("Hamburger Buns", "item");
Result res2 = query.store(1.25, "price");
```

Would store the query:

```
select (item, price) from stock where item = "Hamburger Buns"
```

for `res1` and

```
select (item, price) from stock where price = 1.25
```

for `res2`.

Now you see why we ordered the placeholders in the template above as we did: we used positions 0 and 1 for the ones we want to change frequently, and used 2 and 3 for the parameters that seldom change.

One thing to watch out for, however, is that `Query::store(const char* q)` is also defined for executing the query `q`. Therefore, when you call `Query::store()` (or `use()`, or `execute()`) with only one item and that item is a `const char*`, you need to explicitly convert it into a `SQLString` to get the right overload:

```
Result res = query.store(SQLString("Hamburger Buns"));
```

4.5. Error Handling

If for some reason you did not specify all the parameters when executing the query and the remaining parameters do not have their values set via `def`, the query object will throw a `BadParamCount` object. If this happens, you can get an explanation of what happened by checking the value of `SQLQueryNEParams::string`, like so:

```
query.def["field1"] = "item";
query.def["field2"] = "price";
Result res = query.store(1.25);
```

This would throw `SQLQueryNEParams` because the `whereref` is not specified.

In theory, this exception should never be thrown. If the exception is thrown it probably a logic error in your program.

5. Specialized SQL Structures

The Specialized SQL Structure (SSQLS) feature lets you easily define C++ structures that match the form of your SQL tables. Because of the extra functionality that this feature builds into these structures, MySQL++ can populate them automatically when retrieving data from the database; with queries returning many records, you can ask MySQL++ to populate an STL container of your SSQLS records with the results. When updating the database, MySQL++ can use SSQLS structures to match existing data, and it can insert SSQLS structures directly into the database.

You define an SSQLS using one of several macros. (These are in the file custom.h, and in the file that it includes, custom-macros.h.) There are a bunch of different macros, for different purposes. The following sections will discuss each macro type separately, beginning with the easiest and most generally useful.

5.1. sql_create

This is the most basic sort of SSQLS declaration:

```
sql_create_5(stock, 0, 0,
             string, item,
             int, num,
             double, weight,
             double, price,
             mysqlpp::Date, date)
```

This creates a C++ structure called 'stock' containing five member variables, along with some constructors and other member functions useful with MySQL++.

One of the generated constructors takes a reference to a mysqlpp::Row object, allowing you to easily populate a vector of stocks like so:

```
vector<stock> result;
query.storein(result);
```

That's all there is to it. The only requirements are that the table structure be compatible with the SSQLS's member variables, and that the fields are in the same order.

The general format of this set of macros is:

```
sql_create_#(NAME, KEYS, INITPARMS, TYPE1, ITEM1, ... TYPE#, ITEM#)
```

Where # is the number of member variables, NAME is the name of the structure you wish to create, TYPEx is the type name for a member variable, and ITEMx is that variable's name.

The KEYS and INITPARMS arguments can always be zero, to keep things simple. We will discuss what happens if you use different values in the next few sections.

5.2. sql_create with Compare

SSQLS structures can also have member functions that allow you to compare one structure to another. You simply change the first 0 in the previous example (KEYS) to a higher value. If this number is N, then two structures are considered equal if the first N members of each are equal.

For example:

```
sql_create_5(stock, 1, 0,
             string, item,
             int, num,
             double, weight,
             double, price,
             mysqlpp::Date, date)
```

Here we are saying that the 'item' field is a kind of key field: it is always unique between any two 'stock' items, so if two stock records have equal item values, they are the same stock item.

That change adds the following members to the SSQLS:

```
struct stock {
    ...
    stock(const std::string &p1);
    set(const std::string &p1);
    bool operator ==(const stock &other) const;
    bool operator !=(const stock &other) const;
    bool operator >(const stock &other) const;
    bool operator <(const stock &other) const;
    bool operator >=(const stock &other) const;
    bool operator <=(const stock &other) const;
    int cmp(const stock &other) const;
    int compare(const stock &other) const;
}

int compare(const stock &x, const stock &y);
```

The global `compare()` function compares `x` to `y` and returns `<0` if `x < y`, `0` if `x = y`, and `>0` if `x > y`. `stock::cmp()` and `stock::compare()` are the same thing as `compare(*this, other)`.

The additional constructor initializes the key fields of the structure and leaves the other member variables undefined. This is useful for creating temporary objects to use for comparisons like `x <= stock("Hotdog")`.

Because `stock` is now less-than-comparable you can store the query results in an STL associative container:

```
std::set<stock> result;
query.storein(result);
```

And you can now use it like any other set:

```
cout << result.lower_bound(stock("Hamburger"))->item << endl;
```

This will return the first item that begins with "Hamburger".

You can also use it with any STL algorithm that require the values to be less-than-comparable.

5.3. sql_create with Additional Initializers

If third parameter for this macro (INITPARMS) is nonzero, the SSQLS will have two additional members functions that make it easier to initialize the structure's data members. For example:

```
sql_create_5(stock, 1, 5,
    string, item,
    int, num,
    double, weight,
    double, price,
    mysqlpp::Date, date)
```

will add these functions to the structure relative to that in the previous example:

```
struct stock {
    ...
    stock(const string&, const int&, const double&,
          const double&, const mysqlpp::Date&);
    void set(const string&, const int&, const double&,
             const double&, const mysqlpp::Date&);
}
```

There is one trick with this: because each SSQSL has at least one other constructor besides the one defined by this feature, not every logical value for INITPARMS results in working code. A simple example is setting KEYS and INITPARMS to the same value: you get two identical constructor definitions, so the compiler refuses to compile the code. If you are getting compiler errors having to do with duplicate definitions, try changing this value to zero.

5.4. An Important Limitation of `sql_create`

The features described in the two previous sections work together nicely most of the time. However, if you try to use the same value for the second and third parameters to `sql_create_N`, your program will fail to compile.

Why is this?

The second parameter sets up SSQSL comparisons, one aspect of which defines a constructor taking just the table's key fields. The third parameter sets up an initialization constructor, taking as many fields as you request. When these two values are equal, you get two identical constructor definitions, which is illegal in C++.

The solution is to use 0 for the third parameter, indicating that you do not need a separate full-initialization constructor. `examples/custom6.cpp` illustrates this:

```
#include "util.h"
#include <mysql++.h>
#include <custom.h>

#include <iostream>
#include <iomanip>
#include <vector>

using namespace std;
using namespace mysqlpp;

// To store a subset of a row, we define an SSQSL containing just the
// fields that we will store. There are complications here that are
// covered in the user manual.
sql_create_1(stock_subset,
             1, 0,
             string, item)

int
main(int argc, char *argv[])
{
    try {
        // Establish the connection to the database server.
        Connection con(use_exceptions);
        if (!connect_to_db(argc, argv, con)) {
            return 1;
        }

        // Retrieve a subset of the stock table, and store the data in
        // a vector of 'stock_subset' SSQSLS structures.
        Query query = con.query();
        query << "select item from stock";
        vector<stock_subset> res;
        query.storein(res);

        // Display the result set
        cout << "We have:" << endl;
        vector<stock_subset>::iterator it;
        for (it = res.begin(); it != res.end(); ++it) {
            cout << '\t' << it->item << endl;
        }
    }
    catch (const BadQuery& er) {
        // Handle any query errors
        cerr << "Query error: " << er.what() << endl;
        return -1;
    }
}
```

```
    }
    catch (const BadConversion& er) {
        // Handle bad conversions; e.g. type mismatch populating 'stock'
        cerr << "Conversion error: " << er.what() << endl <<
            "\tretrieved data size: " << er.retrieved <<
            ", actual size: " << er.actual_size << endl;
        return -1;
    }
    catch (const Exception& er) {
        // Catch-all for any other MySQL++ exceptions
        cerr << "Error: " << er.what() << endl;
        return -1;
    }
}
return 0;
}
```

This example shows one other thing, which is how to retrieve a subset of a table using SSQLS. Because we wanted only one column from the table, we had to pass 0 for the third parameter to `sql_create_N` to get the code to compile.

(See the `simple1` example in the Tutorial for another way to accomplish the same thing.)

5.5. Additional Features of Specialized SQL Structures

Up to this point, we haven't been using all of the features in the SSQLS structures we've been generating. We could have used the `sql_create_basic_*` macros instead, which would have worked just as well for what we've seen so far, and the generated code would have been smaller.

Why is it worth ignoring the "basic" variants of these macros, then? Consider this:

```
query.insert(s);
```

This does exactly what you think it does: it inserts 's' into the database. This is possible because a standard SSQLS has functions that the query object can call to get the list of fields and such, which it uses to build an insert query. `query::update()` and `query::replace()` also rely on this SSQLS feature. A basic SSQLS lacks these functions.

Another feature of standard SSQLSes you might find a use for is changing the table name used in queries. By default, the table in the MySQL database is assumed to have the same name as the SSQLS structure type. But if this is inconvenient, you can globally change the table name used in queries like this:

```
stock::table() = "MyStockData";
```

5.6. Harnessing SSQLS Internals

Continuing the discussion in the previous section, there is a further set of methods that the non-"basic" versions of the `sql_create` macros define for each SSQLS. These methods are mostly for use within the library, but some of them are useful enough that you might want to harness them for your own ends. Here is some pseudocode showing how the most useful of these methods would be defined for the `stock` structure used in all the `custom*.cpp` examples:

```
// Basic form
template <class Manip>
stock_value_list<Manip> value_list(cchar *d = "", 
    Manip m = mysqlpp::quote) const;

template <class Manip>
stock_field_list<Manip> field_list(cchar *d = "",
```

```
Manip m = mysqlpp::do_nothing) const;

template <class Manip>
stock_equal_list<Manip> equal_list(cchar *d = ",",
    cchar *e = " = ", Manip m = mysqlpp::quote) const;

// Boolean argument form
template <class Manip>
stock_cus_value_list<Manip> value_list([cchar *d, [Manip m,] ]
    bool i1, bool i2 = false, ..., bool i5 = false) const;

// List form
template <class Manip>
stock_cus_value_list<Manip> value_list([cchar *d, [Manip m,] ]
    stock_enum i1, stock_enum i2 = stock_NULL, ...,
    stock_enum i5 = stock_NULL) const;

// Vector form
template <class Manip>
stock_cus_value_list<Manip> value_list([cchar *d, [Manip m,] ]
    vector<bool> *i) const;

...Plus the obvious equivalents for field_list() and equal_list()
```

Rather than try to learn what all of these methods do at once, let's ease into the subject. Consider this code:

```
stock s("Dinner Rolls", 75, 0.95, 0.97, "1998-05-25");
cout << "Value list: " << s.value_list() << endl;
cout << "Field list: " << s.field_list() << endl;
cout << "Equal list: " << s.equal_list() << endl;
```

That would produce something like:

```
Value list: 'Dinner Rolls',75,0.95,0.97,'1998-05-25'
Field list: item,num,weight,price,date
Equal list: item = 'Dinner Rolls',num = 75,weight = 0.95, price = 0.97,date = '1998-05-25'
```

That is, a "value list" is a list of data member values within a particular SSQLS instance, a "field list" is a list of the fields (columns) within that SSQLS, and an "equal list" is a list in the form of an SQL equals clause.

Just knowing that much, it shouldn't surprise you to learn that `Query::insert()` is implemented more or less like this:

```
*this << "INSERT INTO " << v.table() << " (" << v.field_list() <<
") VALUES (" << v.value_list() << ")";
```

where 'v' is the SSQLS you're asking the Query object to insert into the database.

Now let's look at a complete example, which uses one of the more complicated forms of `equal_list()`. This example builds a query with fewer hard-coded strings than the most obvious technique requires, which makes it more robust in the face of change. Here is `examples/custom5.cpp`:

```
#include "util.h"

#include <mysql++.h>
#include <custom.h>

#include <iostream>
#include <iomanip>
#include <vector>

using namespace std;
using namespace mysqlpp;
```

```
sql_create_5(stock,
    1, 5,
    string, item,
    longlong, num,
    double, weight,
    double, price,
    Date, sdate)

int
main(int argc, char *argv[])
{
    try {
        Connection con(use_exceptions);
        if (!connect_to_db(argc, argv, con)) {
            return 1;
        }

        // Get all the rows in the stock table.
        Query query = con.query();
        query << "select * from stock";
        vector<stock> res;
        query.storein(res);

        if (res.size() > 0) {
            // Build a select query using the data from the first row
            // returned by our previous query.
            query.reset();
            query << "select * from stock where " <<
                res[0].equal_list(" and ", stock_weight, stock_price);

            // Display the finished query.
            cout << "Custom query:\n" << query.preview() << endl;
        }
    }
    catch (const BadQuery& er) {
        // Handle any query errors
        cerr << "Query error: " << er.what() << endl;
        return -1;
    }
    catch (const BadConversion& er) {
        // Handle bad conversions
        cerr << "Conversion error: " << er.what() << endl <<
            "\tretrieved data size: " << er.retrieved <<
            ", actual size: " << er.actual_size << endl;
        return -1;
    }
    catch (const Exception& er) {
        // Catch-all for any other MySQL++ exceptions
        cerr << "Error: " << er.what() << endl;
        return -1;
    }
    return 0;
}
```

This example uses the list form of `equal_list()`. The arguments `stock_weight` and `stock_price` are enum values equal to the position of these columns within the `stock` table. `sql_create_x` generates this enum for you automatically.

The boolean argument form of that `equal_list()` call would look like this:

```
query << "select * from stock where " <<
    res[0].equal_list(" and ", false, false, true, true, false);
```

It's a little more verbose, as you can see. And if you want to get really complicated, use the vector form:

```
vector<bool> v(5, false);
v[stock_weight] = true;
v[stock_price] = true;
query << "select * from stock where " <<
    res[0].equal_list(" and ", v);
```

This form makes the most sense if you are building many other queries, and so can re-use that vector object.

Many of these methods accept manipulators and custom delimiters. The defaults are suitable for building SQL queries, but if you're using these methods in a different context, you may need to override these defaults. For instance, you could use these methods to dump data to a text file using different delimiters and quoting rules than SQL.

At this point, we've seen all the major aspects of the SSQLS feature. The final sections of this chapter look at some of the peripheral aspects.

5.7. Alternate Creation Methods

If for some reason you want your SSQLS data members to have different names than used in the MySQL database, you can do so like this:

```
sql_create_c_names_5(stock, 1, 5,
    string, item, "item",
    int, num, "quantity",
    double, weight, "weight",
    double, price, "price"
    mysqlpp::Date, date, "shipment")
```

If you want your SSQLS to have its data members in a different order from those in the MySQL table, you can do it like this:

```
sql_create_c_order_5(stock, 2, 5,
    mysqlpp::Date, date, 5,
    double, price, 4,
    string, item, 1,
    int, num, 2,
    double, weight, 3)
```

You can combine the custom names and custom ordering like this:

```
sql_create_complete_5(stock, 2, 5,
    mysqlpp::date, date, "shipment", 5,
    double, price, "price", 4,
    string, item, "item", 1,
    int, num, "quantity", 2,
    double, weight, "weight", 3)
```

All three of these macro types have "basic" variants that work the same way. Again, basic SSQLSes lack the features necessary for automatic insert, update and replace query creation.

5.8. Expanding SSQLS Macros

If you ever need to see the code that a given SSQLS declaration expands out to, use the utility `doc/ssqls-pretty`, like so:

```
ssqls-pretty < myprog.cpp | less
```

This locates the first SSQLS declaration in that file and uses the C++ preprocessor to expand that macro. You may have to change the script to tell it where your MySQL++ header files are.

5.9. Extending the SSQSL Mechanism

The SSQSL headers — `custom.h` and `custom-macros.h` — are automatically generated by the Perl script `custom.pl`. Although it is possible to change this script to get additional functionality, it's usually better to do that through inheritance.

A regular user may find it helpful to change the limit on the maximum number of SSQSL data members allowed. It's 25 out of the box. A smaller value may speed up compile time, or you may require a higher value because you have more complex tables than that. Simply change the `max_data_members` variable at the top of `custom.pl` and say '`make`'. The limit for Visual C++ is 31, according to one report. There doesn't seem to be a practical limit with GCC 3.3 at least: I set the limit to 100 and the only thing that happened is that `custom-macros.h` went from 1.3 MB to 18 MB and the build time for `examples/custom.*` got a lot longer.

6. Using Unicode with MySQL++

6.1. A Short History of Unicode

...with a focus on relevance to MySQL++

In the old days, computer operating systems only dealt with 8-bit character sets. This only gives you 256 possible characters, but the modern Western languages have more characters combined than that by themselves. Add in all the other languages of the world, plus the various symbols people use, and you have a real mess! Since no standards body held sway over things like international character encoding in the early days of computing, many different character sets were invented. These character sets weren't even standardized between operating systems, so heaven help you if you needed to move localized Greek text on a Windows machine to a Russian Macintosh! The only way we got any international communication done at all was to build standards on the common 7-bit ASCII subset. Either people used approximations like a plain "c" instead of the French "ç", or they invented things like HTML entities ("ç" in this case) to encode these additional characters using only 7-bit ASCII.

Unicode solves this problem. It encodes every character in the world, using up to 4 bytes per character. The subset covering the most economically valuable cases takes two bytes per character, so most Unicode-aware programs deal in 2-byte characters, for efficiency.

Unfortunately, Unicode came about two decades too late for Unix and C. Converting the Unix system call interface to use multi-byte Unicode characters would break all existing programs. The ISO lashed a wide character sidecar onto C in 1995, but in common practice C is still tied to 8-bit characters.

As Unicode began to take off in the early 1990s, it became clear that some sort of accommodation with Unicode was needed in legacy systems like Unix and C. During the development of the Plan 9 operating system (a kind of successor to Unix) Ken Thompson invented the UTF-8 encoding. UTF-8 is a superset of 7-bit ASCII and is compatible with C strings, since it doesn't use 0 bytes anywhere as multi-byte Unicode encodings do. As a result, many programs that deal in text will cope with UTF-8 data even though they have no explicit support for UTF-8. (Follow the last link above to see how the design of UTF-8 allows this.)

The MySQL database server comes out of the Unix/C tradition, so it only supports 8-bit characters natively. All versions of MySQL could store UTF-8 data, but sometimes the server actually needs to understand the data; when sorting, for instance. To support this, explicit UTF-8 support was added to MySQL in version 4.1.

Because MySQL++ does not need to understand the text flowing through it, it neither has nor needs explicit UTF-8 support. C++'s `std::string` stores UTF-8 data just fine. But, your program probably *does* care about the text it gets from the database via MySQL++. The remainder of this chapter covers the choices you have for dealing with UTF-8 encoded Unicode data in your program.

6.2. Unicode and Unix

Modern Unices support UTF-8 natively. Red Hat Linux, for instance, has had system-wide UTF-8 support since version 8. This continues in the Enterprise and Fedora forks of Red Hat Linux, of course.

On such a Unix, the terminal I/O code understands UTF-8 encoded data, so your program doesn't require any special

code to correctly display a UTF-8 string. If you aren't sure whether your system supports UTF-8 natively, just run the simple1 example: if the first item has two high-ASCII characters in place of the "ü" in "Nürnberg Brats", you know it's not handling UTF-8.

If your Unix doesn't support UTF-8 natively, it likely doesn't support any form of Unicode at all, for the historical reasons I gave above. Therefore, you will have to convert the UTF-8 data to the local 8-bit character set. The standard Unix function `iconv()` can help here. If your system doesn't have the `iconv()` facility, there is a free implementation available from the GNU Project. Another library you might check out is IBM's ICU. This is rather heavy-weight, so if you just need basic conversions, `iconv()` should suffice.

6.3. Unicode and Win32

Each Win32 API function that takes a string actually has two versions. One version supports only 1-byte "ANSI" characters (a superset of ASCII), so they end in 'A'. Win32 also supports the 2-byte subset of Unicode called UCS-2. Some call these "wide" characters, so the other set of functions end in 'W'. The `MessageBox()` API, for instance, is actually a macro, not a real function. If you define the `UNICODE` macro when building your program, the `MessageBox()` macro evaluates to `MessageBoxW()`; otherwise, to `MessageBoxA()`.

Since MySQL uses UTF-8 and Win32 uses UCS-2, you must convert data going between the Win32 API and MySQL++. Since there's no point in trying for portability — no other OS I'm aware of uses UCS-2 — you might as well use native Win32 functions for doing this translation. The following code is distilled from `utf8_to_win32_ansi()` in `examples/util.cpp`:

```
void utf8_to_win32_ansi(const char* utf8_str, char* ansi_str, int ansi_len)
{
    wchar_t ucs2_buf[100];
    static const int ub_chars = sizeof(ucs2_buf) / sizeof(ucs2_buf[0]);

    MultiByteToWideChar(CP_UTF8, 0, utf8_str, -1, ucs2_buf, ub_chars);
    CPINFOEX cpi;
    GetCPInfoEx(CP_OEMCP, 0, &cpi);
    WideCharToMultiByte(cpi.CodePage, 0, ucs2_buf, -1,
                        ansi_str, ansi_len, 0, 0);
}
```

The examples use this function automatically on Windows systems. To see it in action, run simple1 in a console window (a.k.a. "DOS box"). The first item should be "Nürnberg Brats". If not, see the last paragraph in this section.

`utf8_to_win32_ansi()` converts `utf8_str` from UTF-8 to UCS-2, and from there to the local code page. "Wait-aminnit," you shout! "I thought we were trying to get away from the problem of local code pages!" The console is one of the few Win32 facilities that doesn't support UCS-2 by default. It can be put into UCS-2 mode, but that seems like more work than we'd like to go to in a portable example program. Since the default code page in most versions of Windows includes the "ü" character used in the sample database, this conversion works out fine for our purposes.

If your program is using the GUI to display text, you don't need the second conversion. Prove this to yourself by adding the following to `utf8_to_win32_ansi()` after the `MultiByteToWideChar()` call:

```
MessageBox(0, ucs2_buf, "UCS-2 version of Item", MB_OK);
```

All of this assumes you're using Windows NT or one of its direct descendants: Windows 2000, Windows XP, Windows 2003 Server, and someday "Longhorn". Windows 95/98/ME and Windows CE do not support UCS-2. They still have the 'W' APIs for compatibility, but they just smash the data down to 8-bit and call the 'A' version for you.

6.4. For More Information

The Unicode FAQs page has copious information on this complex topic.

When it comes to Unix and UTF-8 specific items, the UTF-8 and Unicode FAQ for Unix/Linux is a quicker way to

find basic information.

7. Incompatible Library Changes

This chapter documents those library changes since the epochal 1.7.9 release that break end-user programs. You can dig this stuff out of the ChangeLog, but the ChangeLog focuses more on explaining and justifying the facets of each change, while this section focuses on how to migrate your code between these library versions.

Since pure additions do not break programs, those changes are still documented only in the ChangeLog.

7.1. API Changes

This section documents files, functions, methods and classes that were removed or changed in an incompatible way. If your program uses the changed item, you will have to change something in your program to get it to compile after upgrading to each of these versions.

v1.7.10

Removed `Row::operator[]()` overloads except the one for `size_type`, and added `Row::lookup_by_name()` to provide the "subscript by string" functionality. In practical terms, this change means that the `row["field"]` syntax no longer works; you must use the new `lookup_by_name` method instead.

Renamed the generated library on POSIX systems from `libsqlplus` to `libmysqlpp`.

v1.7.19

Removed `SQLQuery::operator=()`, and the same for its `Query` subclass. Use the copy constructor instead, if you need to copy one query to another query object.

v1.7.20

The library used to have two names for many core classes: a short one, such as `Row` and a longer one, `MysqlRow`. The library now uses the shorter names exclusively.

All symbols within MySQL++ are in the `mysqlpp` namespace now if you use the new `mysql++ .h` header. If you use the older `sqlplus .hh` or `mysql++ .hh` headers, these symbols are hoisted up into the global namespace. The older headers cause the compiler to emit warnings if you use them, and they will go away someday.

v2.0.0

Library-wide changes

Where possible, all distributed Makefiles only build dynamic libraries. (Shared objects on most Unices, DLLs on Windows, etc.) Unless your program is licensed under the GPL or LGPL, you shouldn't have been using the static libraries from previous versions anyway.

Connection class changes

- `Connection::create_db()` and `drop_db()` return `true` on success. They returned `false` in v1.7.x! This change will only affect your code if you have exceptions disabled.
- Renamed `Connection::real_connect()` to `connect()`, made several more of its parameters default, and removed the old `connect()` method, as it's now a strict subset of the new one. The only practical consequence is that if your program was using `real_connect()`, you will have to change it to `connect()`.
- Replaced `Connection::read_option()` with new `set_option()` mechanism. In addition to changing the

name, programs using this function will have to use the new `Connection::Option` enumerated values, accept a `true` return value as meaning success instead of 0, and use the proper argument type. Regarding the latter, `read_option()` took a `const char*` argument, but because it was just a thin wrapper over the MySQL C API function `mysql_options`, the actual value being pointed to could be any of several types. This new mechanism is properly type-safe.

Exception-related changes

- Classes `Connection`, `Query`, `Result`, `ResUse`, and `Row` now derive from `OptionalExceptions` which gives these classes a common interface for disabling exceptions. In addition, almost all of the per-method exception-disabling flags were removed. The preferred method for disabling exceptions on these objects is to create an instance of the new `NoExceptions` class on the stack, which disables exceptions on an `OptionalExceptions` subclass as long as the `NoExceptions` instance is in scope. You can instead call `disable_exceptions()` on any of these objects, but if you only want them disabled temporarily, it's easy to forget to re-enable them later.
- In the previous version of MySQL++, those classes that supported optional exceptions that could create instances of other such classes were supposed to pass this flag on to their children. That is, if you created a `Connection` object with exceptions enabled, and then asked it to create a `Query` object, the `Query` object also had exceptions disabled. The problem is, this didn't happen in all cases where it should have in v1.7. This bug is fixed in v2.0. If your program begins crashing due to uncaught exceptions after upgrading to v2.0, this is the most likely cause. The most expeditious fix in this situation is to use the new `NoExceptions` feature to return these code paths to the v1.7 behavior. A better fix is to rework your program to avoid or deal with the new exceptions.
- All custom MySQL++ exceptions now derive from the new `Exception` interface. The practical upshot of this is that the variability between the various exception types has been eliminated. For instance, to get the error string, the `BadQuery` exception had a string member called `error` plus a method called `what()`. Both did the same thing, and the `what()` method is more common, so the error string was dropped from the interface. None of the example programs had to be changed to work with the new exceptions, so if your program handles MySQL++ exceptions the same way they do, your program won't need to change, either.
- Renamed `SQLQueryNEParams` exception to `BadParamCount` to match style of other exception names.
- Added `BadOption`, `ConnectionFailed`, `DBSelectionFailed`, `EndOfResults`, `EndOfResultSets`, `LockFailed`, and `ObjectNotInitialized` exception types, to fix overuse of `BadQuery`. Now the latter is used only for errors on query execution. If your program has a "catch-all" block taking a `std::exception` for each try block containing MySQL++ statements, you probably won't need to change your program. Otherwise, the new exceptions will likely show up as program crashes due to unhandled exceptions.

Query class changes

- In previous versions, `Connection` had a querying interface similar to class `Query`'s. These methods were intended only for `Query`'s use; no example ever used this interface directly, so no end-user code is likely to be affected by this change.
- A more likely problem arising from the above change is code that tests for query success by calling the `Connection` object's `success()` method or by casting it to `bool`. This will now give misleading results, because queries no longer go through the `Connection` object. Class `Query` has the same success-testing interface, so use it instead.
- `Query` now derives from `std::ostream` instead of `std::stringstream`.

Result/ResUse class changes

- Renamed `ResUse::mysql_result()` to `raw_result()` so it's database server neutral.
- Removed `ResUse::eof()`, as it wrapped the deprecated and unnecessary MySQL C API function `mysql_eof`. See the `simple2` and `usequery` examples to see the proper way to test for the end of a result set.

Row class changes

- Removed "field name" form of `Row::field_list()`. It was pointless.
- Row subscripting works more like v1.7.9: one can subscript a `Row` with a string (e.g. `row["myfield"]`), or with an integer (e.g. `row[5]`). `lookup_by_name()` was removed. Because `row[0]` is ambiguous (0 could mean the first field, or be a null pointer to `const char*`), there is now `Row::at()`, which can look up any field by index.

Miscellaneous changes

- Removed the backwards-compatibility headers `sqlplus.hh` and `mysql++.hh`. If you were still using these, you will have to change to `mysql++.h`, which will put all symbols in namespace `mysqlpp`.
- Can no longer use arrow operator (`->`) on the iterators into the `Fields`, `Result` and `Row` containers.

7.2. ABI Changes

This section documents those library changes that require you to rebuild your program so that it will link with the new library. Most of the items in the previous section are also ABI changes, but this section is only for those items that shouldn't require any code changes in your program.

If you were going to rebuild your program after installing the new library anyway, you can probably ignore this section.

v1.7.18

The `Query` classes now subclass from `stringstream` instead of the deprecated `strstream`.

v1.7.19

Fixed several `const`-incorrectnesses in the `Query` classes.

v1.7.22

Removed "reset query" parameters from several `Query` class members. This is not an API change, because the parameters were given default values, and the library would ignore any value other than the default. So, any program that tried to make them take another value wouldn't have worked anyway.

v1.7.24

Some freestanding functions didn't get moved into namespace `mysqlpp` when that namespace was created. This release fixed that. It doesn't affect the API if your program's C++ source files say using namespace `mysqlpp` within them.

v2.0.0

Removed `Connection::infoo()`. (I'd call this an API change if I thought there were any programs out there actu-

ally using this...)

Collapsed the `Connection` constructor taking a `bool` (for setting the `throw_exceptions` flag) and the default constructor into a single constructor using a default for the parameter.

Classes `Connection` and `Query` are now derived from the `Lockable` interface, instead of implementing their own lock/unlock functions.

In several instances, functions that took objects by value now take them by `const` reference, for efficiency.

Merged `SQLQuery` class's members into class `Query`.

Merged `RowTemplate` class's members into class `Row`.

Reordered member variable declarations in some classes. The most common instance is when the private section was declared before the public section; it is now the opposite way. This can change the object's layout in memory, so a program linking to the library must be rebuilt.

Simplified the date and time class hierarchy. `Date` used to derive from `mysql_date`, `Time` used to derive from `mysql_time`, and `DateTime` used to derive from both of those. All three of these classes used to derive from `mysql_dt_base`. All of the `mysql_*` classes' functionality and data has been folded into the leaf classes, and now the only thing shared between them is their dependence on the `DTbase` template. Since the leaf classes' interface has not changed and end-user code shouldn't have been using the other classes, this shouldn't affect the API in any practical way.

`mysql_type_info` now always initializes its private `num` member. Previously, this would go uninitialized if you used the default constructor. Now there is no default ctor, but the ctor taking one argument (which sets `num`) has a default.

8. Licensing

The primary copyright holders on the MySQL++ library and its documentation are Kevin Atkinson (1998), MySQL AB (1999 through 2001) and Educational Technology Resources, Inc. (2004 through the date of this writing). There are other contributors, who also retain copyrights on their additions; see the `ChangeLog` file in the MySQL++ distribution tarball for details.

The MySQL++ library and its Reference Manual are released under the GNU Lesser General Public License (LGPL), reproduced below.

The MySQL++ User Manual — excepting some example code from the library reproduced within it — is offered under a license closely based on the Linux Documentation Project License (LDPL) v2.0, included below. (The MySQL++ documentation isn't actually part of the Linux Documentation Project, so the main changes are to LDP-related language. Also, generic language such as "author's (or authors)" has been replaced with specific language, because the license applies to only this one document.)

These licenses basically state that you are free to use, distribute and modify these works, whether for personal or commercial purposes, as long as you grant the same rights to those you distribute the works to, whether you changed them or not. See the licenses below for full details.

8.1. GNU Lesser General Public License

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification").

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the

Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not nor-

mally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to ad-

dress new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the library's name and a brief idea of what it does.>

Copyright © <year> <name of author>

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this lib-

rary; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library `Frob' (a library for tweaking knobs) written by James Random Hacker.

<signature of Ty Coon>, 1 April 1990

Ty Coon, President of Vice

That's all there is to it!

8.2. MySQL++ User Manual License

I. COPYRIGHT

The copyright to the MySQL++ User Manual is owned by its authors.

II. LICENSE

The MySQL++ User Manual may be reproduced and distributed in whole or in part, in any medium physical or electronic, provided that this license notice is displayed in the reproduction. Commercial redistribution is permitted and encouraged. Thirty days advance notice via email to the authors of redistribution is appreciated, to give the authors time to provide updated documents.

A. REQUIREMENTS OF MODIFIED WORKS

All modified documents, including translations, anthologies, and partial documents, must meet the following requirements:

1. The modified version must be labeled as such.
2. The person making the modifications must be identified.
3. Acknowledgement of the original author must be retained.
4. The location of the original unmodified document be identified.
5. The original authors' names may not be used to assert or imply endorsement of the resulting document without the original authors' permission.

In addition it is requested that:

1. The modifications (including deletions) be noted.
2. The authors be notified by email of the modification in advance of redistribution, if an email address is provided in the document.

Mere aggregation of the MySQL++ User Manual with other documents or programs on the same media shall not cause this license to apply to those other works.

All translations, derivative documents, or modified documents that incorporate the MySQL++ User Manual may not have more restrictive license terms than these, except that you may require distributors to make the resulting document available in source format.