

MySQL GUI Tools Manual

MySQL GUI Tools Manual

Abstract

This manual describes all the MySQL GUI Tools.

Document generated on: 2006-09-30 (revision: 3513)

Copyright 2004-2006 MySQL AB

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms: You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how MySQL disseminates it (that is, electronically for download on a website with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of MySQL AB. MySQL AB reserves any and all rights to this documentation not expressly granted above.

Please email [<docs@mysql.com>](mailto:docs@mysql.com) for more information or if you are interested in doing a translation.

Table of Contents

Preface	xi
1. Installation	1
1.1. Introduction	1
1.2. Installing Under Windows	1
1.2.1. Using the Windows Installer	1
1.2.2. Installing From the ZIP File	1
1.3. Installing Under Linux	2
1.3.1. Installing the Generic Tarball	2
1.3.2. RPM Installation	2
1.4. Installing Under Mac OS X	3
2. Running MySQL GUI Tools	4
2.1. Running Under Windows	4
2.2. Running Under Linux	4
2.3. Running On Mac OS X	5
I. MySQL Administrator	6
3. Introduction to MySQL Administrator	9
4. Connection Dialog	10
4.1. Connection Dialog	10
4.2. Configure-service Mode	11
5. Main Window	13
5.1. Introduction	13
5.2. Sidebar	13
5.3. Working Area	14
5.4. Menus	14
5.4.1. File	15
5.4.2. Edit	15
5.4.3. View	15
5.4.4. Tools	15
5.4.5. Help	16
6. Server Information	17
7. Service Control	18
7.1. Introduction	18
7.2. Start/Stop Service	18
7.3. Configure Service	18
8. Startup Variables	21
8.1. Introduction	21
9. Server Connections	22
9.1. Introduction	22
9.2. Threads	22
9.3. User Connections	22
10. User Administration	23
10.1. Introduction	23
10.2. User Accounts	23
10.3. User Information	25
10.4. Global Privileges	26
10.5. Schema Privileges	27
10.6. Table Column Privileges	27
11. Health	28
11.1. Introduction	28
11.2. Connection Health	28
11.3. Memory Health	28
11.4. Creating Custom Health Graphs	29
11.5. Status Variables	31

11.6. System Variables	31
12. Server Logs	33
12.1. Introduction	33
12.2. Error Log	34
12.3. Slow Log	34
12.4. General Log	34
13. Backup	35
13.1. Introduction	35
13.2. Backup Project	35
13.3. Advanced Options	35
13.3.1. Backup Execution Method	35
13.3.2. Output File Options	36
13.4. Schedule	36
14. Restore	38
14.1. Introduction	38
14.2. The General Tab	38
14.3. Restore Content	38
15. Replication Status	40
15.1. Introduction	40
15.2. Configuring Replication Servers	40
15.3. Using the Replication Status Section	40
16. Catalogs	42
16.1. Introduction	42
16.2. Schema Tables	42
16.2.1. Table Status	43
16.2.2. Row Status	43
16.3. Schema Indexes	43
II. MySQL Query Browser	44
17. Introduction to the MySQL Query Browser	46
18. A Tour of the MySQL Query Browser	47
18.1. The Main Query Window	47
18.2. The Query Toolbar	48
18.2.1. The Navigation Buttons	48
18.2.2. The Query Area	48
18.2.3. The Action Buttons	49
18.3. The Advanced Toolbar	49
18.4. The Result Area	50
18.5. The Object Browser	50
18.5.1. The Database Browser	50
18.5.2. The Bookmark Browser	52
18.5.3. The History Browser	52
18.6. The Information Browser	52
18.6.1. The Syntax Browser	53
18.6.2. The Function browser	53
18.6.3. The Parameter Browser	53
18.6.4. The Transaction Browser	53
18.7. The Script Editor	53
18.7.1. The Script Debugging Buttons	55
18.7.2. The Script Area	56
19. Using The MySQL Query Browser	57
19.1. Entering and Editing Queries Manually	57
19.2. Building Queries	57
19.3. Creating Views	58
19.4. Navigating Result Sets Within the Result Area	58
19.5. Exporting Result Sets	59
19.6. Handling BLOB and TEXT Columns	59
19.7. Editing Result Sets Within the Result Area	60
19.8. Comparing Result Sets	60

19.9. Creating Master-Detail Views	61
19.10. Managing Stored Procedures and Functions with MySQL Query Browser ..	62
19.11. Editing Queries From a Development Tool	63
III. Common Administrator and Query Browser Components	64
20. Options Dialog	66
20.1. Introduction	66
20.2. General Options	67
20.3. Connections	67
20.4. Editors	69
20.5. The Browser Section	70
20.5.1. Display Options	71
20.5.2. Query Options	71
20.5.3. Various	72
20.6. Administrator	72
20.6.1. User Administration	72
20.6.2. Health Graphs	72
20.6.3. Backup	72
21. The MySQL Table Editor	74
21.1. Introduction	74
21.2. The Main Editor Window	74
21.3. The Columns And Indices Tab	75
21.3.1. The Column Editor	75
21.3.2. The Detail Area	75
21.4. The Table Options Tab	77
21.5. The Advanced Options Tab	77
21.5.1. The Various Section	77
21.5.2. The Row Options Section	77
21.5.3. The Storage Options Section	77
21.5.4. The Merge Table Options Section	78
21.5.5. The Table RAID Settings Section	78
21.6. Applying Your Changes	78
IV. MySQL Migration Toolkit	80
22. Introduction to the MySQL Migration Toolkit	82
23. An Overview of the Migration Process	83
24. The Migration Process In-Depth	86
24.1. Introduction	86
24.2. The Welcome Screen	86
24.3. The Configuration Type Screen	86
24.4. The Source Database Screen	87
24.4.1. Microsoft Access	88
24.4.2. Microsoft SQL Server	88
24.4.3. Oracle	89
24.4.4. MySQL	91
24.4.5. Saving Connection Information	92
24.5. The Target Database Screen	92
24.6. The Connect to Server Screen	93
24.7. The Source Schema Selection Screen	94
24.8. The Reverse Engineering Screen	95
24.9. The Object Type Selection Screen	96
24.9.1. Migrating a Sub-Set of an Object Type	97
24.10. The Object Mapping Screen	98
24.10.1. GRT Object	99
24.10.2. Table Objects	99
24.11. The Meta Migration Screen	100
24.12. The Manual Editing Screen	100
24.13. The Object Creation Options Screen	102
24.14. The Creating Objects Screen	103
24.15. The Data Mapping Options Screen	104

24.16. The Bulk Data Transfer Screen	105
24.17. The Summary Screen	106
24.18. Saving the Current Application State	107
25. Extending The MySQL Migration Toolkit	108
25.1. Introduction	108
25.2. Architecture of the MySQL Migration Toolkit	108
25.3. The Modular Migration Process	108
25.4. Tools Required to Extend the MySQL Migration Toolkit	109
26. Preparing a Microsoft Access Database for Migration	110
V. MySQL Workbench	112
27. Introduction to MySQL Workbench	114
28. MySQL Workbench Overview	115
28.1. The Vertical Toolbar	117
28.1.1. The Standard Mouse Pointer	117
28.1.2. The Hand Tool	117
28.1.3. The Eraser Tool	117
28.1.4. Using the Layer Tool	117
29. Using MySQL Workbench	119
29.1. Creating a New Schema	119
29.2. Creating Tables	119
29.2.1. Adding Columns	120
29.2.2. Adding Indexes	121
29.2.3. Adding Foreign Keys	122
29.2.4. Adding Triggers	122
29.2.5. Adding Table Options	122
29.2.6. Standard Inserts	122
29.3. Creating Views	122
29.4. Creating Routines	123
29.5. Creating Routine Groups	123
29.6. Creating Synonyms	123
29.7. Creating Structured Data Types	123
29.8. Creating Sequences	123
29.9. Saving a MySQL Workbench File	123
30. MySQL Workbench Reverse Engineering	124
VI. Appendices	125
A. Installing From Source	127
A.1. Introduction	127
A.2. Downloading The Source Code	127
A.3. Building from Source under Windows	127
A.3.1. Prerequisites	127
A.3.2. Compiling and Building	127
A.4. Building from Source under Linux	129
A.4.1. Prerequisites	129
A.5. Compiling and Building	130
A.5.1. Building Static Binaries	131
A.5.2. Building RPMs	132
B. Troubleshooting Application Errors	133
B.1. Troubleshooting Connection Errors	133
B.2. Troubleshooting Display Problems	133
C. How Connections Are Stored	134
D. XML Files Common to the MySQL GUI Applications	135
E. XML Files Used by MySQL Query Browser	136
F. XML Files Used by MySQL Administrator	137
G. The MySQL System Tray Monitor	139
H. Third-party Software Used by the MySQL GUI Tools	141
H.1. PCRE Library	141
H.2. PNG Support	141
I. Notes for Translators	142

List of Figures

4.1. Connection dialog	10
4.2. Switching to configure-service mode	11
4.3. Configure-service mode	12
5.1. MySQL Administrator main window	13
5.2. Marked area between sidebar and working area	14
10.1. User accounts	23
10.2. User Information tab	25
11.1. The Graph Settings dialog	29
12.1. Displaying entries from the middle of the General Log	33
14.1. The Restore Content tab	39
15.1. The Replication Status section	40
18.1. The MySQL Query Browser main query window	47
18.2. The query area expanded to ten lines	48
18.3. The MySQL Query Browser status icons	49
18.4. The advanced toolbar	49
18.5. The database browser	50
18.6. The script editor	53
18.7. The file open dialog	54
19.1. The table tool	57
19.2. The BLOB handling icons	60
19.3. The dynamic parameter list	61
19.4. Stored procedure in database browser	62
20.1. Options dialog	66
20.2. Options dialog: Connections	67
20.3. The query browser options dialog	70
21.1. The MySQL Table Editor	74
21.2. The column editor	75
21.3. The index editor	76
21.4. The Confirm Table Edit dialog	78
23.1. The MySQL Migration Toolkit Migration Plan	83
24.1. The MySQL Migration Toolkit welcome screen	86
24.2. The Configuration Type screen	87
24.3. Source database – Microsoft Access	88
24.4. Source database – Microsoft SQL Server	88
24.5. Source database – Oracle	89
24.6. Oracle JDBC driver not attached	90
24.7. Source database – MySQL	91
24.8. Target Database – MySQL	92
24.9. The Connect to Servers screen	93
24.10. The Source Schema Selection screen	94
24.11. The Reverse Engineering screen	95
24.12. The Object Type Selection screen	96
24.13. The detail view of the Object Type Selection screen	97
24.14. The Add Ignore Pattern dialog	98
24.15. The Object Mapping screen	99
24.16. The Meta Migration screen	100
24.17. The Manual Editing screen	100
24.18. The Manual Editing screen – detailed view	101
24.19. The Object Creation Options screen	102
24.20. The Creating Objects screen	103
24.21. The Data Mapping Options screen	104
24.22. The Bulk Data Transfer screen	105
24.23. The Summary screen	106

26.1. The show section	110
26.2. The system objects	110
26.3. Granting access to the system objects	111
28.1. Workbench design screen – Windows	115
28.2. Workbench design screen – OS X	115
28.3. Workbench design screen – Linux	116
29.1. Creating a table	119
29.2. Creating columns	120
B.1. Error message dialog	133
G.1. The MySQL System Tray Monitor	139

List of Tables

A.1. Required Components	130
A.2. Requirements for MySQL Query Browser	130

Preface

The MySQL GUI Tools provide a suite of desktop applications for the administration of MySQL instances and for building and manipulating the data within MySQL databases.

This book provides a combined guide to the following tools:

- **MySQL Administrator** – provides administration tools and utilities to monitor and administer databases, users privileges and data within a MySQL instance.
- **MySQL Query Browser** – provides an interface for building and testing queries interactively.
- **MySQL Migration Toolkit** – used to migrate existing database information to MySQL (Windows only).
- **MySQL Workbench** – a database development environment where you can model your tables and relationships and produce MySQL table definitions.

Note

Not all tools are available for all platforms.

Chapter 1. Installation

1.1. Introduction

MySQL GUI Tools run on Windows, Linux, and Mac OS X. You can find MySQL GUI Tools for the operating system of your choice on the [MySQL GUI Tools Downloads](http://dev.mysql.com/downloads/gui-tools/5.0.html) [http://dev.mysql.com/downloads/gui-tools/5.0.html] page.

1.2. Installing Under Windows

MySQL GUI Tools run on recent 32-bit Windows NT based operating systems, including Windows 2000, XP, and 2003. They don't run on Windows NT 4 and below.

1.2.1. Using the Windows Installer

MySQL GUI Tools can be installed on all Windows operating systems using the Windows Installer (.msi) installation package. The MSI package is contained within a ZIP archive named `mysql-gui-tools-version-win32.msi`, where *version* indicates the MySQL GUI Tools version.

The Microsoft Windows Installer Engine was updated with the release of Windows XP; those using a previous version of Windows can reference [this Microsoft Knowledge Base article](http://support.microsoft.com/default.aspx?scid=kb;EN-US;292539) [http://support.microsoft.com/default.aspx?scid=kb;EN-US;292539] for information on upgrading to the latest version of the Windows Installer Engine.

In addition, Microsoft has recently introduced the WiX (Windows Installer XML) toolkit. This is the first highly acknowledged Open Source project from Microsoft. We have switched to WiX because it is an Open Source project and it allows us to handle the complete Windows installation process in a flexible manner using scripts.

Improving the MySQL Installation Wizard depends on the support and feedback of users like you. If you find that the MySQL Installation Wizard is lacking some feature important to you, or if you discover a bug, please report it in our bugs database using the instructions given in [How to Report Bugs or Problems](http://dev.mysql.com/doc/refman/5.1/en/bug-reports.html) [http://dev.mysql.com/doc/refman/5.1/en/bug-reports.html].

To install MySQL GUI Tools, right click on the MSI file and select Install. After prompting you for your preferences, installation will begin.

Unless you choose otherwise, MySQL GUI Tools are installed in `C:\%PROGRAMFILES%\MySQL\MySQL Tools for version\`, where `%PROGRAMFILES%` is the default directory for programs on your machine and *version* is the version number of MySQL GUI Tools. The `%PROGRAMFILES%` directory might be `C:\Program Files` or `C:\programme`.

Note

Installing MySQL GUI Tools using the Windows installer automatically creates entries in the [Start](#) menu.

1.2.2. Installing From the ZIP File

If you are having problems running the installer, as an alternative, you can download a ZIP file without an installer. That file is called `mysql-gui-tools-noinstall-version-win32.zip`. Using a ZIP program, unpack it to the directory of your choice. You may also want to create shortcuts to `MySQLAdministrator.exe`, `MySQLMigrationTool.exe`, `MySQLQueryBrowser.exe`, and `MySQLWorkbench.exe` for your desktop or the quick launch bar.

1.3. Installing Under Linux

MySQL GUI Tools runs on Linux machines that have a graphical desktop installed. It is designed to run under the Gnome desktop with GTK2 and has been tested on Linux kernel versions 2.4 and 2.6. It should also run on other versions, and even on a number of Unix-like operating systems.

1.3.1. Installing the Generic Tarball

The generic tar archive allows you to install MySQL GUI Tools on most Linux distributions. The tarball file is called `mysql-gui-tools-version.tar.gz`, where *version* indicates the MySQL GUI Tools version (for example, 5.0r3).

To see all files in the tarball, run this command:

```
shell> tar -tzf mysql-gui-tools-version.tar.gz
```

To install MySQL GUI Tools, run this command:

```
shell> tar --directory=/opt -xzf mysql-gui-tools-version.tar.gz
```

This installs the various application binaries in the directory, `/opt/mysql-gui-tools-version`.

If you install MySQL GUI Tools to the `opt` directory, icons for use with desktop shortcuts or for creating menu items are found under the `opt/mysql-gui-tools-version/share/mysql-gui` directory.

1.3.2. RPM Installation

In addition to a generic tarball, some distribution-specific RPMs are available. Currently these include Red Hat Enterprise Linux (RHEL) 3 and 4, Fedora Core 5 (FC5), and SuSE Linux 10.x. For FC5 and SuSE the `gtkmm24` toolkit is a requirement for installing the RPM version of MySQL GUI Tools. On FC5 you may install this toolkit from the command line in the following way:

```
shell> yum -install gtkmm24
```

For installation on SuSE Linux:

```
shell> yast2 -i gtkmm24
```

Note

You may need root privileges to run the `yast2` command.

The Red Hat RPMs are self contained so no additional packages need to be installed.

The RPM downloads are made up of the individual GUI Tools components combined into a single TAR archive. Extract the individual RPMs in the following way:

```
shell> tar -zxf mysql-gui-tools-version.tar.gz
```

This will decompress the RPM files to the current directory.

Install all the RPM files by typing:

```
shell> rpm -iv mysql-*.rpm
```

If you are upgrading to a newer version of MySQL GUI Tools type:

```
shell > rpm -Uv mysql-*.rpm
```

If you install the RPM files individually, you must install the [mysql-gui-tools-version.rpm](#) file first in order to satisfy dependencies.

When upgrading, packages cannot be installed separately because version conflicts will arise.

If you wish, you may install only one of the MySQL GUI Tools. For example, to install MySQL Administrator only, do the following:

```
shell > rpm -iv mysql-gui-toolsversion.rpm mysql-administratorversion.rpm
```

If possible, the RPM installation process creates shortcuts in the start menu of your window manager. For example, SuSE Linux with the KDE window manager adds shortcuts to the MySQL GUI Tools under the Development, Other Tools menu item. Likewise, with FC5, shortcuts are created under the Development menu item.

Icons for use with desktop shortcuts or for creating menu items are found under the [/usr/share/mysql-gui](#) directory.

See the [MySQL GUI Tools Downloads](http://dev.mysql.com/downloads/gui-tools/5.0.html) [http://dev.mysql.com/downloads/gui-tools/5.0.html] page for the most up-to-date listing of the various RPM packages available.

1.4. Installing Under Mac OS X

To install MySQL Administrator under Mac OS X, double-click the downloaded [.dmg](#) file and wait for it to be opened and attached. Once a window containing the MySQL GUI Tools icon pops up, drag it to your Applications folder — or any other location you prefer.

Once the copy is complete, you may eject the disk image.

Chapter 2. Running MySQL GUI Tools

How you start any one of the MySQL GUI Tools depends on the operating system you are using.

2.1. Running Under Windows

Under Windows, the names of the executable files in the MySQL GUI Tools suite are:

- `MySQLAdministrator.exe`
- `MySQLQueryBrowser.exe`
- `MySQLWorkbench.exe`
- `MySQLMigrationTool.exe`

Start any one of the MySQL GUI Tools by double clicking its desktop icon, or by selecting it from the [Start](#) menu. Alternatively, you can open a DOS window and start it from the command line. For example, you could run the MySQL Administrator in the following way:

```
C:\> "C:\%PROGRAMFILES%\MySQL\MySQL Tools for version\MySQLAdministrator.exe"
```

`%PROGRAMFILES%` is the default directory for programs on your machine, for example `C:\Program Files` or `C:\programme`. If your path contains spaces, you must enclose the command within quotation marks as shown above.

2.2. Running Under Linux

Under Linux, the names of the executable files in the MySQL GUI Tools suite are:

- `mysql-administrator`
- `mysql-query-browser`
- `mysql-workbench`

Note

There is no Linux version of MySQL Migration Toolkit.

If you installed MySQL GUI Tools to the `/opt` directory using the tar archive file, change into the `/opt/mysql-gui-tools-version` directory to run any one of the MySQL GUI Tools.

To run MySQL Administrator you would type:

```
shell> ./mysql-administrator
```

However, adding `/opt/mysql-gui-tools-version` to the `PATH` variable makes it much simpler to run MySQL GUI Tools — you need not worry about your present working directory.

When installed from RPM files, the MySQL GUI Tools are found in the `/usr/bin` directory. This directory is usually included in the `PATH` variable, so running any one of the MySQL GUI Tools simply requires typing the executable file name, regardless of your current directory. For example:

```
shell> mysql-administrator
```

For those distributions that create menu items, you may also, of course, start any one of the tools by choosing the menu item.

2.3. Running On Mac OS X

On Mac OS X, navigate to the MySQL GUI Tools installation directory and double-click on the application you wish to start.

Note

There is no OS X version of MySQL Migration Toolkit.

Part I. MySQL Administrator

Table of Contents

3. Introduction to MySQL Administrator	9
4. Connection Dialog	10
4.1. Connection Dialog	10
4.2. Configure-service Mode	11
5. Main Window	13
5.1. Introduction	13
5.2. Sidebar	13
5.3. Working Area	14
5.4. Menus	14
5.4.1. File	15
5.4.2. Edit	15
5.4.3. View	15
5.4.4. Tools	15
5.4.5. Help	16
6. Server Information	17
7. Service Control	18
7.1. Introduction	18
7.2. Start/Stop Service	18
7.3. Configure Service	18
8. Startup Variables	21
8.1. Introduction	21
9. Server Connections	22
9.1. Introduction	22
9.2. Threads	22
9.3. User Connections	22
10. User Administration	23
10.1. Introduction	23
10.2. User Accounts	23
10.3. User Information	25
10.4. Global Privileges	26
10.5. Schema Privileges	27
10.6. Table Column Privileges	27
11. Health	28
11.1. Introduction	28
11.2. Connection Health	28
11.3. Memory Health	28
11.4. Creating Custom Health Graphs	29
11.5. Status Variables	31
11.6. System Variables	31
12. Server Logs	33
12.1. Introduction	33
12.2. Error Log	34
12.3. Slow Log	34
12.4. General Log	34
13. Backup	35
13.1. Introduction	35
13.2. Backup Project	35
13.3. Advanced Options	35
13.3.1. Backup Execution Method	35
13.3.2. Output File Options	36
13.4. Schedule	36
14. Restore	38
14.1. Introduction	38

14.2. The General Tab	38
14.3. Restore Content	38
15. Replication Status	40
15.1. Introduction	40
15.2. Configuring Replication Servers	40
15.3. Using the Replication Status Section	40
16. Catalogs	42
16.1. Introduction	42
16.2. Schema Tables	42
16.2.1. Table Status	43
16.2.2. Row Status	43
16.3. Schema Indexes	43

Chapter 3. Introduction to MySQL Administrator

MySQL Administrator is a program for performing administrative operations, such as configuring, monitoring and starting and stopping a MySQL server, managing users and connections, performing backups, and a number of other administrative tasks.

Most tasks can be performed using a command-line client such as [mysqladmin](http://dev.mysql.com/doc/refman/5.1/en/mysqladmin.html) [http://dev.mysql.com/doc/refman/5.1/en/mysqladmin.html], or [mysql](http://dev.mysql.com/doc/refman/5.1/en/mysql.html) [http://dev.mysql.com/doc/refman/5.1/en/mysql.html], but MySQL Administrator has the following advantages:

- Its graphical user interface makes it more intuitive.
- It provides a better overview of the settings that are crucial for the performance, reliability, and security of your MySQL servers.
- It displays performance indicators graphically, thus making it easier to determine and tune server settings.

MySQL Administrator is designed to work with MySQL versions 4.0 and higher.

MySQL Administrator is, to a large extent, the result of feedback MySQL AB has received from many users over a period of several years. However, if you find it's lacking some feature important to you, or if you discover a bug, please use our [MySQL Bug System](http://bugs.mysql.com) [http://bugs.mysql.com] to request features or report problems.

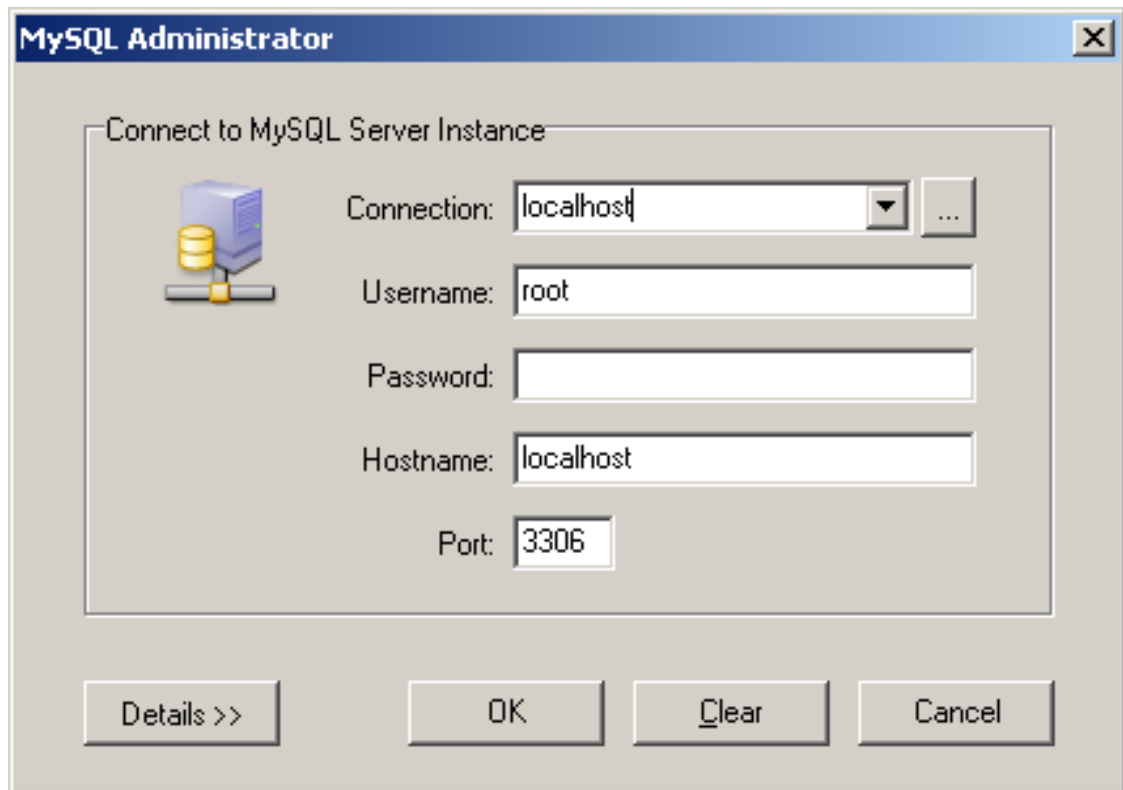
Chapter 4. Connection Dialog

4.1. Connection Dialog

Once MySQL Administrator has been started, it displays a connection dialog box. Specify the MySQL server you would like to connect to, the necessary credentials, and both the host and port. You may also specify a number of other options, if required.

As an alternative to connecting to a running MySQL server, you can run MySQL Administrator in [configure -service mode](#).

Figure 4.1. Connection dialog



Note that MySQL Administrator will not work with MySQL servers prior to version 4.0. If you try to connect to an older server, an error message will appear, telling you that this is not possible.

If the server connection is successfully established, all of the values filled in the fields of the connection dialog will be saved for future connections (see the section that describes [how MySQL Administrator stores connection information](#)). The `Password` field, however, will always be empty: For security reasons, the password is not stored along with the other options, unless you explicitly specify otherwise in the [General Options section of the Options dialog](#).

You can change any of the values in the connection profiles by overwriting the existing values with new ones. Similarly, if you select another connection profile, you can change any of its values. When you click the `OK` button after changing a profile's values, the changes are stored permanently if a successful connection to the MySQL server is established. However, the preferred method for changing a profile is to click the `...` button next to the `Connection` drop-down box. That brings up an `Options` dialog

window that has a [Connections](#) section for modifying connection profiles. You can find the values you can set in the Connection dialog window in the description of the [Connections](#) section of the [Options](#) dialog (see [Options dialog](#)).

4.2. Configure-service Mode


Normally, you will use MySQL Administrator to connect to a running MySQL server. This is done by starting the program and specifying in the [Connection](#) dialog which server to connect to.

Alternatively, you can run MySQL Administrator in configure-service mode. This is useful in the following cases:

- When you want to start a MySQL server from within MySQL Administrator. Obviously, you cannot connect to a server that is not running, so you will have to use the configure-service mode to start that server. Note that, at the moment, you can only start servers on the same machine where MySQL Administrator is running (that is on the local host).
- When you have [multiple servers](http://dev.mysql.com/doc/refman/5.1/en/multiple-servers.html) [http://dev.mysql.com/doc/refman/5.1/en/multiple-servers.html] running on your machine. When running in normal mode on Windows, MySQL Administrator only allows you to configure the Windows service that is listed first in the Windows services management. This means that if you have more than one Windows service, you can only configure the first one of those, no matter which server you connect to. If you want to configure a service other than the first service, you will have to run MySQL Administrator in configure-service mode.

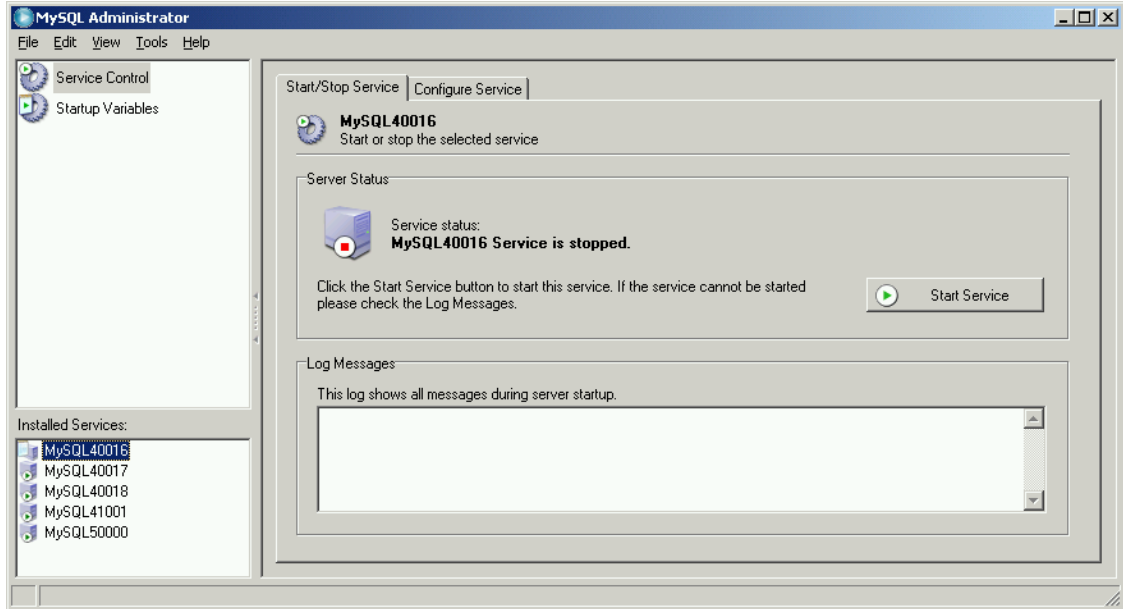
To start MySQL Administrator in configure-service mode, hold down the **CTRL** key in the [Connection](#) dialog. The [Cancel](#) button will toggle to [Skip](#) when that key is pressed.

Figure 4.2. Switching to configure-service mode



Clicking on the [Skip](#) button will skip the normal mode, thus starting MySQL Administrator in configure-service mode.

Figure 4.3. Configure-service mode



Under [Installed Services](#) in the [sidebar](#), you can select the service that you wish to configure, or the server that you want to stop or start. In configure-service mode, not all sidebar options are available:

- The [Service Control](#) section allows you to start and stop a MySQL server, and to configure the Windows service managing that server.
- The [Startup Variables](#) section allows you to set variables for server startup.

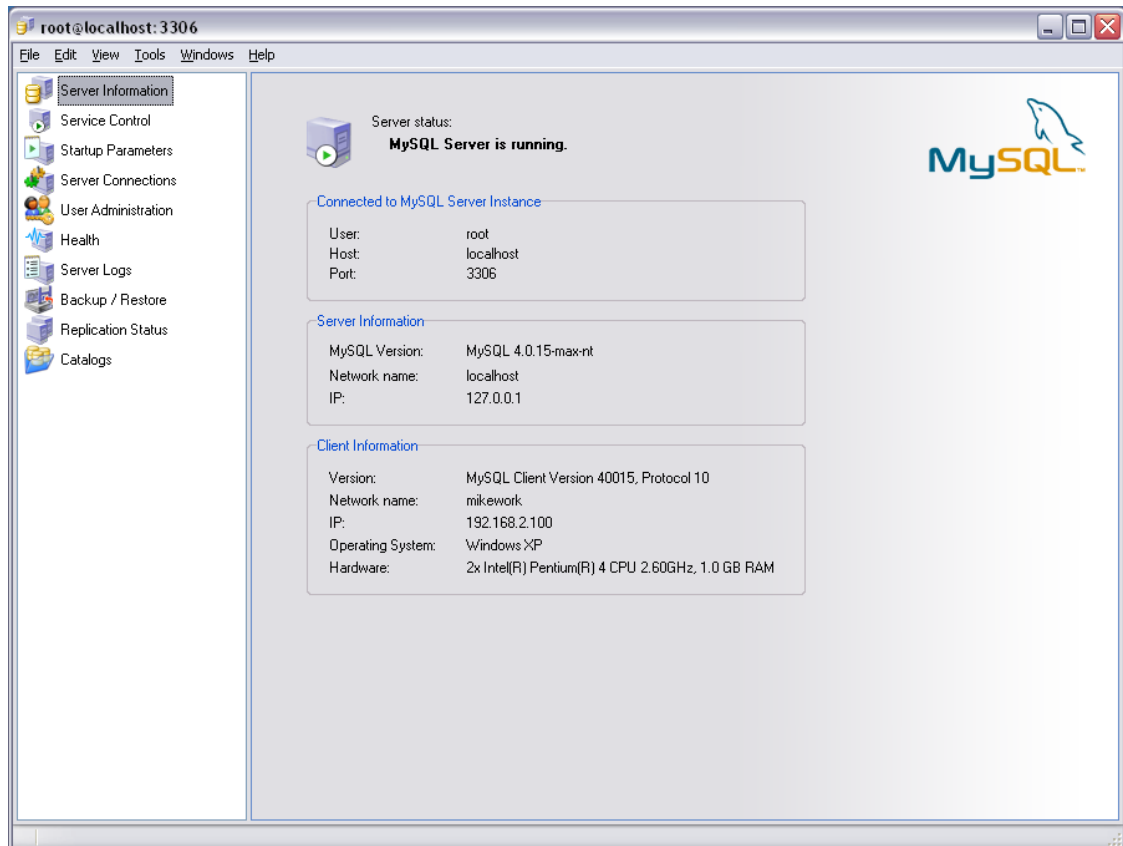
Both of those sections are described in what follows.

Chapter 5. Main Window

5.1. Introduction

After you have successfully connected to your MySQL server, the main window of MySQL Administrator appears (see figure). The main window contains a sidebar, a working area, and menus with sub-menus.

Figure 5.1. MySQL Administrator main window



5.2. Sidebar

The sidebar consists of the following options:

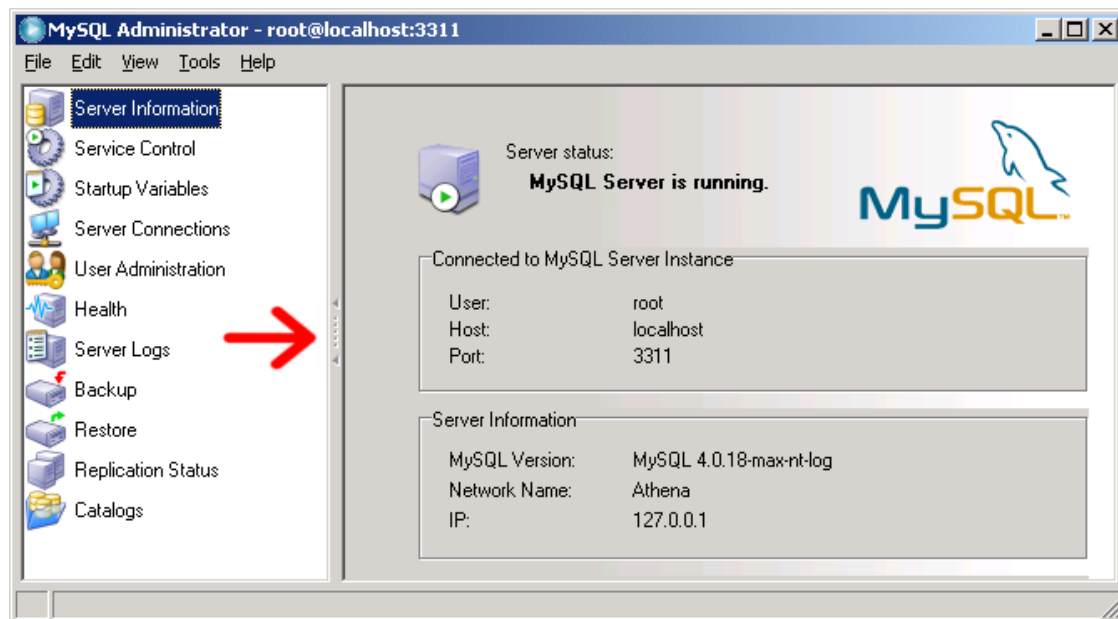
- **Server Information:** Information about the MySQL server you are connected to, about MySQL Administrator and the machine that it runs on, and about your connection.
- **Service Control:** Start and stop the MySQL server. If MySQL Administrator is running on Windows and connected to a MySQL server that is also running on Windows, you can configure the MySQL service.
- **Startup Variables:** Configure the startup variables for the MySQL server.

- **Server Connections:** View or kill current connections to the MySQL server (also called threads).
- **User Administration:** Administer existing users, add new users, and delete existing users.
- **Health:** Graphical displays of a number of usage or hit-rate values that affect server performance, and a hierarchical view of system and status variables.
- **Server Logs:** View log file entries.
- **Backup:** Plan and administer backup projects, select the databases to be backed up, and start the backup process.
- **Restore:** Restore databases from existing backups.
- **Replication Status:** View information about replication master and slave servers.
- **Catalogs:** View information about databases, tables, columns, indexes, and rows; optimize tables.

The following sections describe each of these in detail.

You can change the width of the sidebar in the main window by dragging the separator between the sidebar and the working area. To hide the sidebar completely, click on the marked area in the separator (located in the vertical center of it). To make the sidebar visible again, double click the marked area (which then forms the left margin of the main window).

Figure 5.2. Marked area between sidebar and working area



5.3. Working Area

While the **sidebar** allows you to select options, the working area displays information, or fields, that pertain to the selected section. For some sections, the working area is split into two or more tabs.

5.4. Menus

5.4.1. File

5.4.1.1. New Instance Connection ...

Opens the [Connection dialog](#), allowing you to open a connection to another MySQL server. You may open an arbitrary number of connections. MySQL Administrator opens a new main window for each successfully established connection. Connections may be to the same or different servers. Multiple connections may be opened using the same profile or different profiles.

5.4.1.2. Manage Connections ...

Opens the [Options dialog](#), with the [Connections section](#) preselected.

5.4.1.3. Save current Connection ...

Opens the [Options dialog](#), with the [Connections section](#) preselected. Furthermore, a new connection profile labeled [New Connection](#) is created, which is also preselected. The fields are filled with the values of the most recent connection. You can edit the values and save the settings in the new profile. You may also want to rename the profile from [New Connection](#) to something more appropriate. If you click the [Close](#) button without saving your changes using the [Apply Changes](#) button, you will be prompted to save or discard your changes.

5.4.1.4. Close

Closes the MySQL Administrator window and terminates the connection to the MySQL server that had been established in that window. If you have opened more than one instance (program window, or process) of MySQL Administrator, this command will not close the other windows.

5.4.2. Edit

The Cut, Copy, and Paste commands allow you to cut, copy, or paste selected text in the [working area](#).

5.4.2.1. Options

Opens the [Options dialog](#) window.

5.4.3. View

The View menu selects sections from the [sidebar](#). Selecting an item from this menu is the same as clicking the section with the same name from the sidebar.

5.4.4. Tools

The Tools menu starts other programs.

5.4.4.1. MySQL Command Line Client

This menu item is available only if MySQL Administrator finds the executable file for the MySQL client. At the moment, it searches locations such as `C:\mysql\bin`, `D:\mysql\bin`, and `C:\program directory\mysql\bin` to find `mysql.exe`. If the executable file is found, you can start the command-line client with this command. MySQL Administrator will use the credentials provided in the [Connection dialog](#) to start a server connection using the command-line client.

5.4.4.2. Windows Command Line

Opens a DOS window.

5.4.5. Help

5.4.5.1. Online Help

Opens a window displaying the MySQL Administrator section of the MySQL reference manual.

5.4.5.2. Report Bug to MySQL

Opens a connection to the MySQL bug reporting system in your web browser.

5.4.5.3. Launch MySQL Website

Opens a connection to the MySQL web site in your web browser.

5.4.5.4. About ...

Displays information about MySQL Administrator.

Chapter 6. Server Information

Provides basic information about connection, server, and client.

- Connected to MySQL Server Instance

Information about the parameters used to connect to the MySQL server.

- **User**: Username used for the current connection to the MySQL server.
- **Host**: The name of the machine hosting the MySQL server. This is either identical to the name listed in **Network Name**, or it is `localhost`. The former indicates that you are connecting to a server that either runs on a remote machine, or that you have specified a hostname different from `localhost` when connecting. The latter indicates that the MySQL server runs on the same machine where MySQL Administrator runs.
- **Port**: The port used for connecting to the MySQL server.

- Server Information

Information about the MySQL server and its host.

- **MySQL Version**: The version of the MySQL server.
- **Network Name**: The name of the machine hosting the MySQL server.
- **IP**: The IP number of the MySQL server host.

- Client Information

Information about MySQL Administrator and the host client.

- **Version**: The version of MySQL Administrator, and the version number of the MySQL client/server protocol.
- **Network name**: The name of the host where MySQL Administrator runs.
- **IP**: The IP number of the MySQL Administrator host.
- **Operating System**: Information about the operating system on which MySQL Administrator runs.
- **Hardware**: Information about the hardware on which MySQL Administrator runs.

Chapter 7. Service Control

7.1. Introduction

This section shows how to start and stop the MySQL server you have connected to. Currently, this is only possible if the server is running on the same machine that hosts MySQL Administrator and if that service has been configured. (Future releases will be capable of starting and stopping MySQL servers on remote machines.) If the server runs as a stand-alone application, the Service Control section will not be enabled.

The Configure Service tab is available only for MySQL servers that run under Windows.

7.2. Start/Stop Service

The [Start/Stop Service](#) tab displays the status of the MySQL service. This service starts and stops the MySQL server. Under [Service Status](#), a text message indicates whether or not the service has been started. If the service is running, the button label is [Stop Service](#). Click it to stop the server. If the service is stopped, the button label is [Start Service](#). Click it to start the service.

Under [Log Messages](#), you will find messages that the MySQL server produces when it starts or stops. Messages produced by the server during startup look something like this:

```
Trying to start the server ...
040119 18:33:41 InnoDB: Started
MySql: ready for connections.
Version: '4.0.16-max-nt-log' socket: '' port: 3306
Server was started.
```

When the server stops, it produces messages such as these:

```
Trying to stop the server ...
040119 18:31:45 MySql: Normal shutdown
040119 18:31:46 MySql: Forcing close of thread 11 user:
'superuser'
040119 18:31:46 MySql: Forcing close of thread 8 user: 'superuser'
...
Server was stopped.
```

The details and format of the messages might vary depending on the version of the MySQL server you are using and its configuration.

7.3. Configure Service

Use this tab to configure the Windows service that starts and stops the MySQL server and to set server options. This tab is specific to MySQL servers running under operating systems of the Windows NT family, including Windows NT, Windows 2000, and Windows XP; it is not available for servers running under a different operating system. Obviously, you need to have a Windows service installed before you can configure it. See [Starting MySQL as a Windows Service](#) [<http://dev.mysql.com/doc/refman/5.1/en/windows-start-service.html>], for more information on how to set up a Windows service.

Note that if you have more than one MySQL server running as a Windows service, you should run MySQL Administrator in [configure-service mode](#) to administer these services. In normal mode, you can only administer the Windows service listed first in the Windows service manager, no matter which server you are connected to.

Note that the settings in the [Server Features](#) group take effect only after restarting the MySQL

server.

The following settings are available:

- Service Settings
 - **Launch MySQL server automatically:** If checked, the Windows service starts the MySQL server automatically on operating system startup.
 - **Display Name:** The display name of the Windows service (not the service name!) as it will appear in the Windows services manager when displaying the properties of that service.
 - **Service Description:** The description of the Windows service as it will appear in the Windows services manager.
- Option File
 - **Config Filename:** You can enter the pathname of the server option file. MySQL will enter that information in the Windows registry, and the Windows service starting the MySQL server will use that file, rather than the default file (`C:\windows\installation\directory\my.ini`). *Note that this is available for MySQL servers as of version 4.0.3 only.* See [Creating an Option File](#) [<http://dev.mysql.com/doc/refman/5.1/en/windows-create-option-file.html>], for details about the server option files.
 - **Section Name:** Setting a value other than `mysqld` is most useful when running multiple servers; see [Running Multiple Servers on Windows](#) [<http://dev.mysql.com/doc/refman/5.1/en/multiple-windows-servers.html>].
- Server Features

Note that all server features will only take effect after the MySQL server has been restarted.

- **Support for InnoDB:** If checked, the MySQL server is started with support for InnoDB tables. If unchecked, it doesn't start the InnoDB storage engine; this saves memory, but InnoDB tables cannot be used. If this checkbox is unchecked and greyed out, it means that your MySQL server isn't compiled with support for InnoDB tables. (This is, for example, true for some versions that run under a commercial license.) For more information on InnoDB tables, see [Storage Engines and Table Types](#) [<http://dev.mysql.com/doc/refman/5.1/en/storage-engines.html>].

Be careful when disabling this option! If any databases on the MySQL server contain InnoDB tables, disabling this option makes those tables inaccessible.

- **Support for BDB:** If checked, the MySQL server is started with support for BDB tables. If unchecked, it doesn't start the BDB storage engine; this saves memory, but BDB tables cannot be used. If this checkbox is unchecked and greyed out, it means that your MySQL server isn't compiled with support for BDB tables. (This is, for example, true for non-max versions.) For more information on BDB tables, see [Storage Engines and Table Types](#) [<http://dev.mysql.com/doc/refman/5.1/en/storage-engines.html>].

Be careful when disabling this option! If any databases on the MySQL server contain BDB tables, disabling this option makes those tables inaccessible.

- **Named Pipes:** On operating systems of the Windows NT family, including Windows NT, Windows 2000, and Windows XP, you may use named pipes to connect to a MySQL server running on the local machine, rather than connecting via TCP/IP. If this option is checked, named pipe connections are allowed. Note that named pipes are slower than TCP/IP in many Windows configurations. See [Selecting a MySQL Server type](#) [<http://dev.mysql.com/doc/refman/5.1/en/windows-select-server.html>].

- **Debug Information (slow)**: If checked, the MySQL server is started with an option that allows for easier debugging of problems. Using this option slows down the server significantly, so it is recommended only for debugging purposes. If the MySQL server has not been compiled with debugging support, this option is greyed out. See [Command Options](#) [<http://dev.mysql.com/doc/refman/5.1/en/server-options.html>].
- **Path to Binary**: Contains the pathname of the MySQL server binary file as it is listed in the Windows service manager (for example `C:\mysql\bin\mysqld-max-nt`). Note that the the Windows service manager does not display the `.exe` extension.

Chapter 8. Startup Variables

8.1. Introduction

This section allows you to set variables that the MySQL server reads on startup. These variables are stored in the option file of the server. On Windows, the option file is called `my.ini`, and it is located in the Windows installation directory (for example, `C:\WinNT\my.ini`). Note that it can alternatively be stored in `C:\`, in which case it is named `my.cnf`. Also note that there can be more than one option file. On operating systems other than Windows, the option file is called `my.cnf`, regardless of where it resides. On Windows, MySQL Administrator allows you to store the option file in any location you specify. See [Using Option Files](http://dev.mysql.com/doc/refman/5.1/en/option-files.html) [http://dev.mysql.com/doc/refman/5.1/en/option-files.html] for more information.

Note that the server reads variables stored in the option files only at server startup. After changing an option file, you must restart the server for the changes to take effect. You can do this using the [Stop Server](#) and [Start Server](#) buttons in the [Service Control](#) section.

Chapter 9. Server Connections

9.1. Introduction

This section allows you to view and manage the current connections to the server (also called threads). If you have the [PROCESS](#) privilege, you can see all threads. If you have the [SUPER](#) privilege, you can also terminate threads other than your own. (Note that the [SUPER](#) privilege is available for MySQL servers as of version 4.0.3 only.) Otherwise, you can see your own threads only. See [SHOW PROCESSLIST Syntax](#) [<http://dev.mysql.com/doc/refman/5.1/en/show-processlist.html>], for more information. See [Privileges Provided by MySQL](#) [<http://dev.mysql.com/doc/refman/5.1/en/privileges-provided.html>], for more information on the [PROCESS](#) and the [SUPER](#) privilege (and other available privileges as well).

Note

When you click the [Server Connections](#) section You will see an additional thread for the user you logged in as because the values displayed are retrieved from the MySQL server via an additional separate connection.

9.2. Threads

This tab allows you to view all connections to the server (threads). By default, threads are sorted by their [PID](#) (MySQL process identification number), but you can change that order by clicking on the appropriate headings ([User](#), [Host](#), etc.). You can select one or more threads using the mouse or the keyboard.

To reload the process list from the server, press the [Refresh](#) button.

If you have selected one or more threads, you can terminate them by clicking the [Kill Thread](#) button, or by right-clicking on that thread. Remember that you need the [PROCESS](#) privilege to see threads other than your own, and the [SUPER](#) privilege to kill threads other than your own. Furthermore, note that MySQL Administrator imposes a restriction that prevents you from terminating your own threads. On Windows, using the [Kill Thread](#) button on your own threads has no effect; on Linux, that button is disabled when selecting your own threads.

9.3. User Connections

This tab displays all connections to the server (threads), but in a different way than the [Threads](#) tab does: You will see users rather than individual threads. The number of connections for each user is listed under [Num.](#). Right-clicking on a username and selecting Show User Info displays the [Full Name](#) and the [Description](#), as specified in the [User Information](#) tab of the [User Administration](#) section. Initially, users are sorted by their MySQL [Username](#), but you can change the order by clicking on the appropriate headings ([Num.](#), [Full Name](#), or [Description](#)). Clicking on a username will display the individual threads of that user in the lower box of the tab, ordered by [PID](#). Likewise, you can change the order by clicking on one of the headings.

Clicking on a user will enable the [Kill User](#) button, allowing you to terminate all threads of that user with a single click. This command is also available by right clicking on the username. Clicking on an individual thread in the lower box enables the [Kill Thread](#) button, which is used to terminate a connection. This command is also available by right clicking on a thread.

To reload the user list from the server, press the [Refresh](#) button.

Chapter 10. User Administration

10.1. Introduction

This section shows how to administer existing users, add new users, and delete existing users. For more information on the MySQL user account management and privilege system, see [MySQL User Account Management](http://dev.mysql.com/doc/refman/5.1/en/user-account-management.html) [http://dev.mysql.com/doc/refman/5.1/en/user-account-management.html].

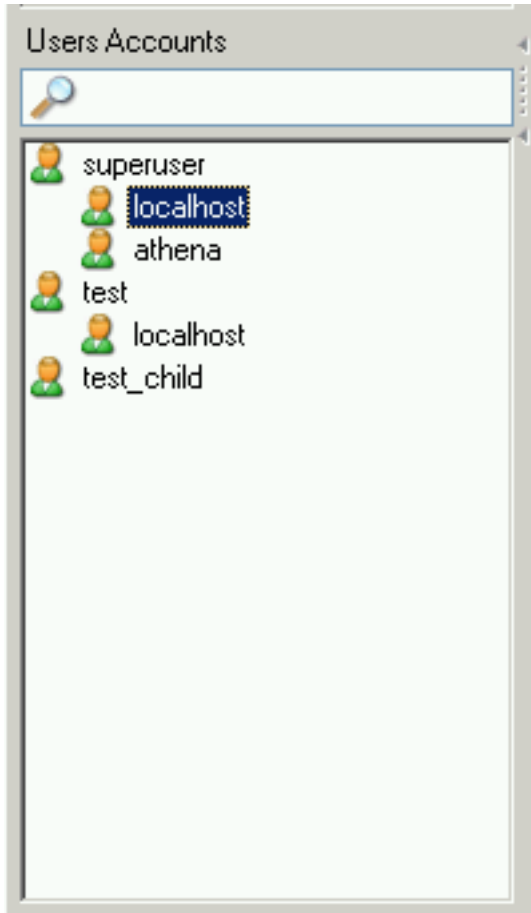
10.2. User Accounts

Existing users are listed in the lower left area of the [sidebar](#). If there are many user accounts on your MySQL server, the field with the magnifying glass icon may be handy for filtering the users you are looking for. Typing in `t` or `T`, for example, will set the filter to users whose names begin with a `t`. The filter works in a case-insensitive fashion.

If you are an experienced MySQL user, you will notice that users are not listed in the typical MySQL way (`user@host`), but rather with their usernames only. If double clicking on a username does not show subcategories for that user, this means that the user's privileges are not restricted to a specific host, or hosts. If there are subcategories, the user's privileges are restricted to a specific host, or set of hosts. You may assign various (and different) sets of privileges, depending on the host from which the user connects to the MySQL server. That concept is described in detail in [Access Control, Stage 1: Connection Verification](#) [http://dev.mysql.com/doc/refman/5.1/en/connection-access.html].

Note that MySQL Administrator has a different concept of what a `user` is than MySQL has. In MySQL, a user is always identified by a `username/host` combination. This means that, for example, `'brian'@'%'` may be a user completely different from `'brian'@'localhost'`. The former might be `Brian Miller`, while the latter might be `Brian Schultz`. That distinction does not hold true for MySQL Administrator: User `brian` is always a particular user, no matter from which host he connects to the MySQL server. That said, `'brian'@'%'` may still have privileges different from `'brian'@'localhost'`.

Figure 10.1. User accounts



For example, you might have a user called `superuser`. If double clicking on that username shows two subcategories, `localhost` and `athena`, this means that the user has a specific set of privileges if he/she connects from `localhost`, and a (probably different) set of privileges if he/she connects from `athena`.

Right clicking on a username, or on one of the subcategories a user might have, provides access to the following commands:

- **Add a new User:** Selecting this command adds a new user with the default name `New User` to the user list. You will have to fill in at least a username in the `MySQL User` field of the `User Information` tab, which, after applying this change, will also rename that user in the user accounts list.
- **Add Host from which the User can connect:** Allows you to specify a host from which the user can connect. This command is unavailable if you highlight a subcategory.
- **Remove Host from which the User can connect:** Allows you to remove a host from which the user can connect. If there are no sub-categories, the user will be removed (you will be prompted to confirm the removal).
- **Clone User:** Makes a copy of the selected user, including all subcategories. The new user is called `New User` until you specify another name in the `MySQL User` field of the `User Information` tab. This command is useful for creating users with identical privileges.
- **Delete User:** Allows you to delete the selected user. You will be prompted to confirm that you really want to delete that user.

- [Refresh User List](#): Rereads the user list from the MySQL server. This is helpful when other users (on other connections) are editing user accounts. Note that selecting that command will collapse all subcategories.

10.3. User Information

The fields in this tab are unavailable unless you either click on a username in the [User Accounts](#) area of the [sidebar](#), or create a new user in that sidebar or using the [New User](#) button on this tab. In the former case, the fields are filled with the information stored for that user, in the latter case, all fields are empty.

Note that all fields will be stored only when you click the [Apply Changes](#) button. If you don't want to save your changes, click the [Discard Changes](#) button.

Figure 10.2. User Information tab

The values you enter in the fields of this tab are stored on the MySQL server. The fields in the [Login Information](#) group of this tab are stored in the `user` table of the `mysql` database, while the fields in the [Additional Information](#) group are stored in the `user-info` table of the `mysql` database. The latter table is created the first time you add a new user, or change an existing user. Its contents may look like this:

```
mysql> SELECT * FROM mysql.user-info \G
***** 1. row *****
      User: superuser
      Full-name: Stefan Hinz
      Description: MySQL Documentation Team
      Email: stefan@mysql.com
      Contact-information: Phone: +49 30 123456789
                          Fax: +49 30 987654321
                          ICQ: 123456789
      Icon:
```

- **Login Information**
 - **MySQL User:** Username that is specified when connecting to the MySQL server. Note that this username has nothing to do with usernames you might use elsewhere on your operating system. Using special characters or spaces for a username is not recommended.
 - **Password:** Password that is specified when connecting to the MySQL server. Note that this password has nothing to do with passwords you might use elsewhere on your operating system. Again, don't use special characters or spaces for the password. It is, however, recommended that you use non-trivial passwords. A very trivial password might be `abcde`, while a non-trivial password might be `lEtusMak3iThartdoGesz`. Note that it is not mandatory to specify a password, although it is highly recommendable to do so.
 - **Confirm Password:** Fill in the password once again to make sure you do not accidentally introduce a typo.
- **Additional Information**

All information specified here is optional.

- **Full Name:** As opposed to the username in **MySQL User**, the full name may contain any characters, including spaces.
- **Description:** Additional description of user.
- **Email:** Email address of user.
- **Contact Information:** More contact information, like postal address, or messenger information.
- **Icon:** Rather than using the default icon, you may select a different icon by pressing the **Load from Disk** button. The icon you use must be exactly 48 by 48 pixels in size, and its file format must be **PNG** (portable network graphics).

10.4. Global Privileges

This tab is available only if the **Show Global Privileges** checkbox in the **Administrator section of the Options dialog** has been checked.

Under **Assigned Privileges**, you find the global privileges assigned to the selected user. For more information about those privileges, see **Privileges Provided by MySQL** [<http://dev.mysql.com/doc/refman/5.1/en/privileges-provided.html>]. *Global* means that those privileges apply to the MySQL server in general (like the **Shutdown** privilege), or to all databases on the server (like the **SELECT** privilege).

Privileges not assigned to the user are listed under **Available Privileges**. You can remove privileges by selecting them in the left box, and then clicking the **>** button. Removed privileges will appear

in the [Available Privileges](#) box. You can assign privileges by selecting them in the [Available Privileges](#) box, and then clicking the < button.

10.5. Schema Privileges

This tab allows you to grant privileges on a schema level. In MySQL, a schema is the same as a database. To grant privileges for a specific database, click on the database name in the left box. (You cannot select more than one database at a time.)

Privileges not assigned to the user are listed under [Available Privileges](#). You can remove privileges by selecting them in the left box, and then clicking the > button. Removed privileges will appear in the [Available Privileges](#) box. You can assign privileges by selecting them in the [Available Privileges](#) box, and then clicking the < button.

10.6. Table Column Privileges

This tab is available only if the [Show Table/Column Privileges](#) checkbox in the [Administrator section of the Options dialog](#) has been checked. It allows you to grant privileges on a table level or on a column level.

To access the tables, double click on the schema (database) name containing that table. You will not see and cannot grant privileges on a schema level in this tab; if you want to do that, use the [Global Privileges](#) tab instead. Also, you cannot select more than one table at a time.

If you want to grant privileges on a column level, double click on the table name; this will give you access to its columns. You cannot select more than one column at a time.

Privileges not assigned to the user are listed under [Available Privileges](#). You can remove privileges by selecting them in the left box, and then clicking the > button. Removed privileges will appear in the [Available Privileges](#) box. You can assign privileges by selecting them in the [Available Privileges](#) box, and then clicking the < button.

Chapter 11. Health

11.1. Introduction

The [Health](#) section of MySQL Administrator shows you how to graphically monitor the real-time status of your MySQL server. A selection of pre-made status graphs are available, and you can easily create your own status graphs and organize them into custom pages and groups.

The default graphs provide graphical information about the MySQL server's current performance, load, and memory usage. Additional pages list all server variables, with tabs for memory health, status variables and system variables. Note that clicking on the status pages will open an additional connection to the MySQL server, which you can see when viewing threads and users in the [Server Connections](#) section.

We'll begin by covering the different pre-made status graphs, and then describe how to create your own status graphs. Finally we'll describe the pages in the [Health](#) section dedicated to the various status and system variables.

11.2. Connection Health

This tab provides information about connections to the MySQL server. The following three pre-made graphs are displayed:

- [Connection Usage](#): Displays a percentage of the maximum number of simultaneous clients allowed (`threads_connected / max_connections`). If this value gets close to 100%, it could indicate that you should increase the value of `max_connections`, which you can do in the [Startup Variables](#) section. The maximum number of connections since MySQL Administrator was loaded is also shown. For more information on server connections, see [SHOW VARIABLES Syntax](#) [<http://dev.mysql.com/doc/refman/5.1/en/show-variables.html>], and [Too many connections](#) [<http://dev.mysql.com/doc/refman/5.1/en/too-many-connections.html>].
- [Traffic](#): This is a representation of the `bytes_sent` status variable. This does not represent total bytes sent since server startup, but instead represents the current number of bytes sent (the `bytes_sent` variable is a cumulative total). The initial maximum value of the graph is 102400, this maximum value will be automatically increased as needed.
- [Number of SQL Queries](#): This is a representation of the `com_select` status variable. This does not represent total queries since server startup, but instead represents the current number of queries (the `com_select` variable is a cumulative total). The initial maximum value of the graph is 10, this maximum value will be automatically increased as needed.

11.3. Memory Health

This tab provides information about memory caches that affect the performance of the MySQL server. For each cache, there is a display that shows the current hit rate or usage of the cache, and a display that shows the hit rate or usage for the last few minutes.

- [Query Cache Hitrate](#): When in use, the query cache of the MySQL server stores the text of a `SELECT` query together with the corresponding result that was sent to the client. If an identical query is later received, the server will retrieve the results from the query cache rather than parsing and executing the same query again. The hit rate indicates the ratio of queries that were cached and queries that were not cached. The higher the hit rate, the better the performance of the server regard-

ing `SELECT` queries. See [The MySQL Query Cache](http://dev.mysql.com/doc/refman/5.1/en/query-cache.html) [http://dev.mysql.com/doc/refman/5.1/en/query-cache.html].

- **Key Efficiency:** The horizontal bar graph indicates how much of the allocated key buffer is in use. If a fairly large percentage of the key buffer is in use you may wish to increase the `key_cache_size` setting in the [Startup Variables](#) section. The line graph is an indication of the number of `key_read_requests` that resulted in actual `key_reads`. The higher the hit rate the more effective your key buffer is. A low hit rate can indicate that you need to increase the size of your keycache.

11.4. Creating Custom Health Graphs

While the pre-made status graphs can provide you with information of the current state of your MySQL server, there may be additionally information that you wish to visually track. The MySQL Administrator [Health](#) section allows you to add additional tabs and graphs to track any information that is contained within the [Status Variables](#).

The first step in creating custom health graphs is to create a new page in the [Health](#) section to store your graphs. To create a new page right click anywhere in the working area and choose the Add a Page option. You will be prompted for a page name and description and a new page will be created.

Once you have created a new page, you will need to create a group. All graphs are organized into groups. For example, the `Key buffer usage` and `Key buffer hit rate` graphs are both members of the group titled `Key Efficiency`. To create a group, right click within your newly created page and choose the Add a Group option.

To create a custom graph, right click within a group and choose the Add a graph option. The following dialog will appear:

Figure 11.1. The Graph Settings dialog

Graph Settings [X]

General

Graph Type: The type of visual representation of the values over time.

Caption: The graph's caption.

Display Caption Whether the captions should be displayed or not.

Refresh Time: Time to next refresh of the value.

Value Settings

Value Formula: Formula to calculate the value. Use [variable] to get values from variables, ^[variable] to get relative values (relative to last evaluation).

Value Unit: The unit of the value.

Value Caption: A caption for the value.

Min. / Max. Values

Min. Value: Lowest possible value. Lower values will be set to this value.

Max. Value: Highest possible value. Higher values will be set to this value.

Autoextend Max. Value The max. value will be overwritten by greater values. This will scale the graph dynamically.

Max. Formula: The max. value is calculated by this formula. See value formula.

Max. Caption: The caption of the max. value.

The [Line-Graph](#) is appropriate for showing trends over time, while the [Bar-Graph](#) option will be more appropriate for showing percentage information. If you choose to create a bar-graph, you will need to choose a caption for your graph. Captions are not allowed for line-graphs.

The data on your graph is set using the [Value Formula](#). You can create a formula using any of the variables available in the [Status Variables](#) and [System Variables](#) tabs. To use the value of a variable, wrap the variable name in square brackets (i.e. `[com_select]`), if you want the relative value of the variable prepend the square brackets with a `^` character.

For example, if we wanted to graph the percentage of temporary tables that were created on disk we could use the following formula:

```
[created_tmp_disk_tables] / [created_tmp_tables]
```

If we wanted to track the number of temporary tables created on a continuous basis, we could use the following as our formula:

```
^[created_tmp_tables]
```

After creating your formula, select the [Value Unit](#) that best represents your data. You can choose from [Percentage](#), [Count](#), [Byte](#), and [Seconds](#). In our examples we would use [Percentage](#) for the first example, and [Count](#) for the second. You can also assign a caption to the value.

Once your formula is assigned you should configure the [Max. Value](#) and [Min. Value](#) for the graph, so that your data is spread evenly across your graph. You can set arbitrary values based on your estimates of how large the values will grow, and check the [Autoextend Max. Value](#) option to allow MySQL Administrator to automatically increase the [Max Value](#) setting automatically to prevent your data from extending off of your graph.

You can also set the [Max. Value](#) option by way of a formula assigned in the [Max Formula](#) field. The same syntax applies in this field as applies in the [Value Formula](#) field. For example, if you were creating a bar graph tracking the number of temporary disk tables created, you could use `[created_tmp_disk_tables]` as the [Value Formula](#) and `[created_tmp_tables]` as the [Max. Formula](#).

11.5. Status Variables

This tab provides information about the status variables of the MySQL server. In the left box of the working area, you find categories and subcategories. Double click on a category name to expand or collapse its subcategories.

For each category, and each subcategory, the status variables are listed in the right box of the working area, together with their values, and with a short explanation. See [SHOW STATUS Syntax](#) [<http://dev.mysql.com/doc/refman/5.1/en/show-status.html>], for more detail on server status variables.

Use the [Refresh](#) button to make MySQL Administrator retrieve the variables and their values once again.

Right clicking in the right box allows you to copy variables to the clipboard.

11.6. System Variables

This tab provides information about the system variables of the MySQL server. In the left box of the working area, you can find categories and subcategories. Double click on a category name to expand or collapse its subcategories.

For each category, and each subcategory, the system variables are listed in the right box of the working area, together with their values and a short explanation. See [SHOW VARIABLES Syntax](#) [<http://dev.mysql.com/doc/refman/5.1/en/show-variables.html>], for more information on system vari-

ables.

System variables that can be set at runtime are marked as editable, that is they are prefixed with an icon different from the variables that cannot be set at runtime. Double clicking on editable variables will open a dialog box allowing you to specify another value for the variable. The changes you make will take effect as soon as you confirm the changes by clicking the **OK** button. Click **Abort** if you want to keep the current value. Changes made in this manner will only persist until the server is restarted, after which the server will revert to its startup settings. In MySQL Administrator, system variables are set on a global level only. Note that you need appropriate privileges to change variables globally. See [SET Syntax](http://dev.mysql.com/doc/refman/5.1/en/set-option.html) [<http://dev.mysql.com/doc/refman/5.1/en/set-option.html>].

Use the **Refresh** button to make MySQL Administrator retrieve the variables and their values once again.

Right clicking in the right box allows you to copy variables to the clipboard.

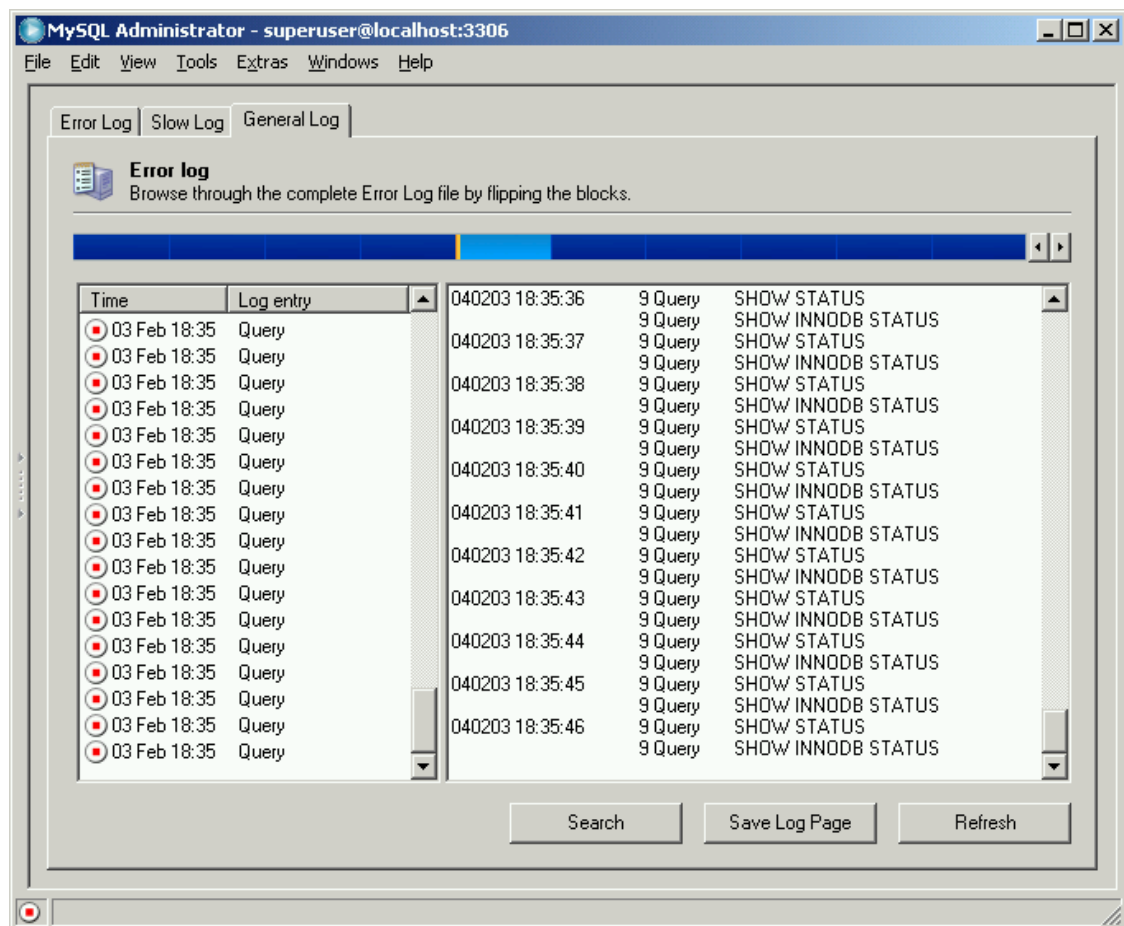
Chapter 12. Server Logs

12.1. Introduction

The [Server Logs](#) section provides easy access to the various log files generated by the MySQL server. You can access the server's error log, slow query log, and general log using this section.

This section has one tab for each of the log files, and each tab is split into two boxes. Above the boxes, there is a bar which allows you to navigate through larger log files by breaking the log file into individual pages. Clicking in the bar will display the individual pages. The light blue area represents the portion of the log file you are currently viewing.

Figure 12.1. Displaying entries from the middle of the General Log



You can only view log files that are stored on the machine where MySQL Administrator is running, you cannot view log files on a remote machine.

Also note that some tabs may be disabled. If a tab is disabled it indicates that the log file in question has not been enabled on the MySQL server. You can enable the various log files in the [Startup Variables](#) section.

The left box (labeled [Events of page x / y](#)) displays the major events that have occurred on the

current page of the log file. The right box (labeled [Page Content](#)) displays the detailed lines of the log file. By clicking on an event in the left box, the first detailed entry of the event will be highlighted in the right box.

In addition to the log file tabs, the [Server Logs](#) section has four action buttons for managing log information:

- [Open Other File](#) will allow you to open a log file other than the currently active log file. This can be useful in situations where you rotate your log files as it allows you to open log files that have been rotated out of active use.
- [Search](#) opens a dialog where you can specify a search term to look for in the log file. If it's found, it will be marked as selected in the right box. This command is also available by right clicking in the right box and choosing the Search ... option.
- [Save Log Page](#) opens a dialog that allows you to specify the path and file name where the current page of the log file should be stored. By default, MySQL Administrator suggests [ErrorLog.txt](#), [SlowLog.txt](#), or [GeneralLog.txt](#) as a filename. Note that executing this command does not affect the actual log file in any way; it just creates a copy of the *current* log file contents that you may store for documenting or analyzing. This command is also available by right clicking in the right box and choosing the Save log Page ... option.
- [Refresh](#) will reload the log files.

12.2. Error Log

Displays the MySQL server's error log. The error log file contains information indicating when the MySQL server was started and stopped and also any critical errors found when running. See [The Error Log](#) [<http://dev.mysql.com/doc/refman/5.1/en/error-log.html>], for more information about that log file.

12.3. Slow Log

Displays the MySQL server's slow query log. If enabled (which you can do in the [Startup Variables](#) section), the MySQL server writes a log file containing all SQL commands that took more than a specific number of seconds to execute. (That number can be configured by the [long-query-time](#) variable; its default value is 10 seconds.) The time to get the initial table locks are not counted as execution time. For more information about that log file, see [The Slow Query Log](#) [<http://dev.mysql.com/doc/refman/5.1/en/slow-query-log.html>].

12.4. General Log

Displays the MySQL server's general query log. This log can be very useful when you suspect an error in a client and want to know exactly what the MySQL server thought the client sent to it. For more information about that log file, see [The General Query Log](#) [<http://dev.mysql.com/doc/refman/5.1/en/query-log.html>].

Chapter 13. Backup

13.1. Introduction

This section shows you how to create backup projects and to perform backups. This section will produce backup files that are plain text and composed of multiple SQL statements that can be executed to restore your database, similar to the output generated by the `mysqldump` utility. For more information on `mysqldump`, see `mysqldump` [<http://dev.mysql.com/doc/refman/5.1/en/mysqldump.html>], in the reference manual.

Existing backup projects are listed in the lower left area of the `sidebar`. If there are many backup projects, the field with the magnifying glass icon can be used to filter the projects you are looking for. Typing in `t` or `T`, for example, will set the filter to projects whose names begin with a `t`. The filter works in a case-insensitive fashion.

You can update the list of backup projects by right clicking on a project and choosing the Refresh Backup Projects option from the drop-down menu. You can delete a backup project by right clicking on the project and choosing the Delete Backup Project(s) option from the drop-down menu.

If you are not working from a previously saved backup project, click the `New Project` button to create a new backup project. You cannot begin configuring a backup operation without either loading an existing project or creating a new project.

Once you have configured your backup project, you can click the `Save Project` button to save your project for future use. To start your backup, click the `Execute Backup Now` button. You will be prompted for a path and filename for the backup file, after which the backup operation will begin.

13.2. Backup Project

The `Backup Project` tab allows you to set the name of your backup project and also determine what data will be backed up.

You can name your project by filling out a name in the `Project Name` field. If you have enabled the `Add Date/Time to Backup Files` option in the `Administrator Options` dialog then the date and time of the backup will be appended to the project name automatically. See [Section 20.6.3, “Backup”](#) for more information on setting this option.

Specify a name that clearly identifies the project, such as `Full Backup` or `World Only`. The project name may contain arbitrary characters, including spaces.

To choose the databases and tables that will be backed up, select the database from the schemata list and click the `>` button. Once the database is in the `Backup Content` list, you can uncheck specific tables that you do not wish to back up.

13.3. Advanced Options

The `Advanced Options` tab allows you to configure advanced options that affect the way the backup operation behaves. These include concurrency controls, backup type, and output formatting.

13.3.1. Backup Execution Method

The `Backup Execution Method` affects concurrency within the backup itself.

`Lock all tables` will cause all tables within the backup to be locked simultaneously to prevent

writing. This will prevent users from writing to the tables during the backup operation and will ensure that the tables are in a consistent state. This option is recommended for MyISAM tables.

`Single transaction` will begin a transaction before reading from the tables. Because of the InnoDB versioning system this will allow all tables to be read in a consistent state. This method is recommended when backing up InnoDB tables.

`Normal Backup` will lock each table individually when performing the backup. While this can have less impact on the actions of users during the backup operation, there is a potential for tables to be backed up in an inconsistent state in relation to one another. This method is not recommended unless the risk of inconsistency is minimal.

`Complete Backup` overrides the table selection made in the `Backup Project` tab and backs up all tables in the selected schema. This option can be useful when scheduling the backup of a schema that changes regularly.

13.3.2. Output File Options

The `Output File Options` control the formatting of the backup file. At present there is only one `Backup Type` available, the `SQL Files` format, which outputs a file similar to the output of the `mysqldump` utility.

`No CREATEs` will suppress `CREATE DATABASE` statements. This can be useful if you need to restore the data to a different database than it was backed up from.

`No EXTENDED INSERTS` will suppress the default behavior of combining multiple inserts into a single statement (for example `INSERT INTO test VALUES ('1',Some data',2234),('2','Some More Data',23444)`). While the `EXTENDED INSERT` syntax is more efficient both in terms of storage and execution, it is not compatible with most RDBMSes.

`Add DROP Statements` adds `DROP TABLE IF EXISTS` statements to the backup file.

`Complete INSERTs` will produce full `INSERT` statements that include a column list (i.e. `INSERT INTO tablename(columnA, columnB) VALUES(valueA, valueB)`). Without this option the column list will be omitted.

`Comment` writes additional backup information to the backup file in the form of SQL comments.

`Don't write full path` causes MySQL Administrator to write only the table name in statements, instead of `schema_name.table_name`.

`Compatibility mode` creates backup files that are compatible with older versions of MySQL Administrator.

`ANSI Quotes` will cause all table and database names to be quoted with ANSI style double quotes instead of backticks.

`Disable keys` will add an `ALTER TABLE ... DISABLE KEYS` statement to the backup file before the `INSERT` statements that populate the tables. This allows for faster recovery by loading all data before rebuilding the index information.

13.4. Schedule

The `Schedule` tab is used to execute backup operations on a daily, weekly, or monthly basis.

To enable scheduling for your backup project, check the `Schedule this backup project` checkbox. You will need to choose a target directory for the backup files and also a name for them. You can browse for a target directory by clicking the `...` button.

If you have enabled the [Add Date/Time to Backup Files](#) option in the [Administrator Options](#) dialogue, then the date and time of the backup will be appended to the file name automatically. See [Section 20.6.3, “Backup”](#) for more information on setting this option.

Once you have decided on a file name and target directory, you can set the interval of the backup operation. Regardless of your choice of daily, weekly, or monthly backups, you can set the time that the backup will occur by setting the [Time](#) option. Set the backup time in a 24 hour format (i.e. 23:30).

When you select the [Execute Backup Weekly](#) option, a set of check boxes representing the days of the week will be displayed. Check the boxes corresponding to the days when you want to perform backups.

When you select the [Execute backup monthly](#) option, a list of dates will appear, allowing you to choose the day of the month for backups.

Click the [Save Project](#) button to save your preferences. After your project is saved it will be executed at the next scheduled interval. To run the backup operation immediately click the [Execute Backup Now](#) button.

Chapter 14. Restore

14.1. Introduction

The [Restore](#) section helps with disaster recover by helping you restore backup files created using MySQL Administrator. See [Chapter 13, Backup](#) for more information on creating backups with MySQL Administrator.

The backup files created by MySQL Administrator are similar but not identical to those created by [mysqldump](#). A consequence of the differences is that MySQL Administrator cannot read dump files created by [mysqldump](#).

14.2. The General Tab

In this tab, you specify which backup file should be restored, the character set and format, and the target location for restored tables.

To set any options on this tab, you must first open a backup file by clicking the [Open Backup File](#) button at the bottom of the window. When all options are set, you can click the [Start Restore](#) button to begin the restore process.

The following options are available:

- [File to restore](#): Specify the path and filename of the backup file. You can browse for that file by clicking the [Open Backup File](#) button.
- [Target Schema](#): You can choose the [Original Schema](#) option to have all tables restored into their original databases. You can also choose the [New Schema ...](#) option to restore all tables into a new database. Finally you can choose from one of the existing databases in the drop-down list to force all tables to be restored to that database.
- [Backup Type](#): At the moment, the only available option is [SQL Files](#). Those are backup files containing SQL statements such as those produced by [mysqldump](#).
- [Force - Continue even if there is an sql error](#): This will allow a restore operation to proceed even if errors are encountered.
- [Create database\(s\) if they don't exist](#): If a non-existent database is referenced it will automatically be created if this option is checked.
- [File Charset](#): Backup files created with MySQL Administrator are encoded in the [utf8](#) character set. Backup files created with other tools such as [winmysqladmin](#) may be encoded with other character sets. Before a backup file can be successfully imported you will need to specify its character set.

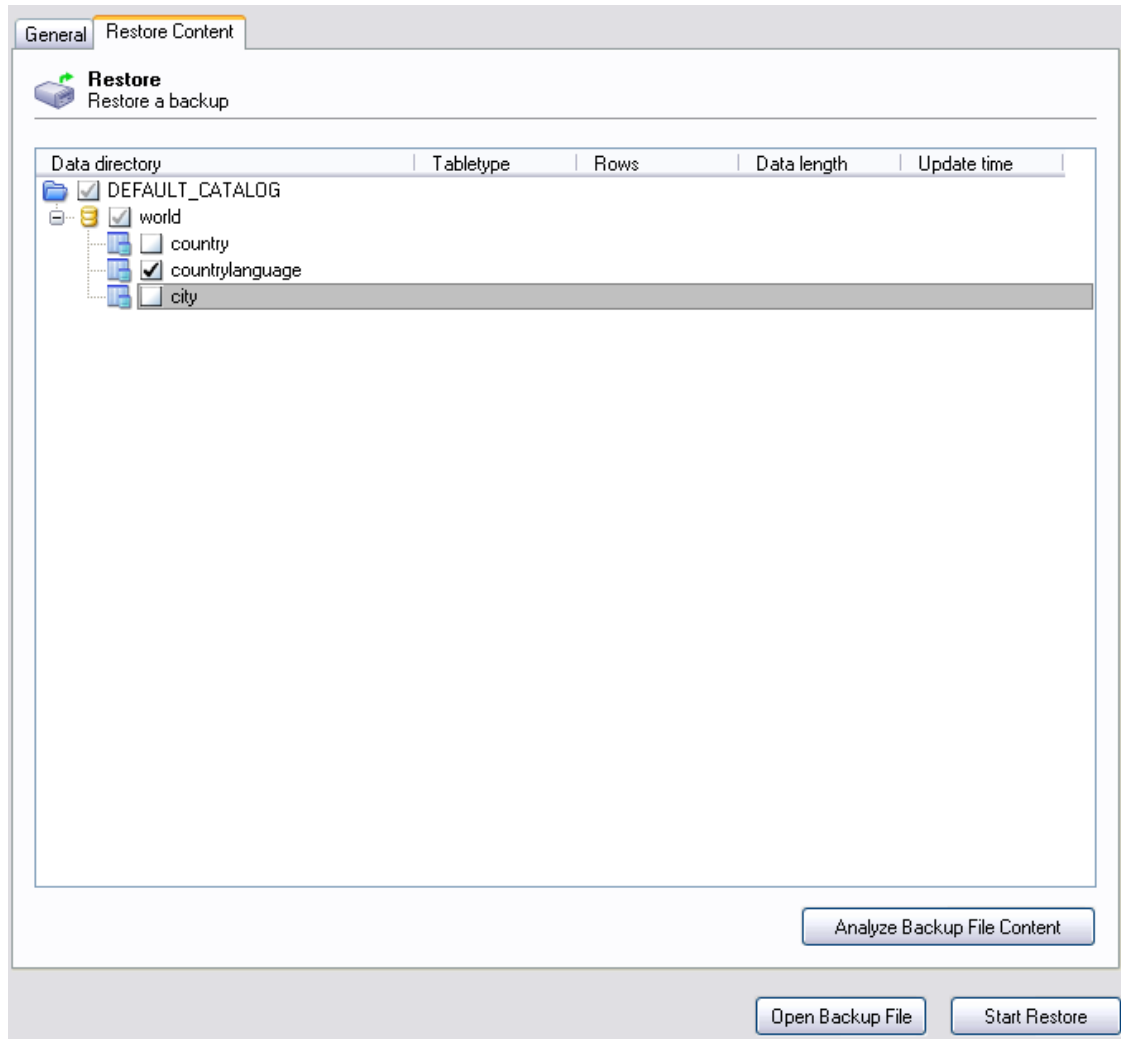
If you do not know the character set of your backup file, click the [Auto-Detect Character Set](#) button to have MySQL Administrator attempt to determine it automatically.

14.3. Restore Content

If you do not wish to restore all the databases or tables listed within your backup file, you can use the [Restore Content](#) tab to limit which databases and tables are restored.

After clicking the [Analyze Backup File Content](#) button, a list of all databases and tables contained within the backup file will be displayed.

Figure 14.1. The Restore Content tab



To prevent a database or table from being restored, uncheck the box next to the object's name.

Chapter 15. Replication Status

15.1. Introduction

The [Replication Status](#) sections give you a visual overview of your MySQL replication master and slaves. The server status icons beside each machine tells you which slave machines are connected to the master and which ones are offline.

15.2. Configuring Replication Servers

To use the [Replication Status](#) section you must be connected to the server that acts as the master in your replication scheme. Your MySQL servers must be version 4 or higher to use this feature.

For a slave server to show in the [Replication Status](#) screen, it must register itself with the master server. To have a slave server register itself with the master, set the `report-host` option in your server configuration (`my.cnf/my.ini`) file. See [Replication Startup Options](#) [<http://dev.mysql.com/doc/refman/5.1/en/replication-options.html>].

You can also set the value in the [Replication](#) tab of the [Startup Variables](#) section of MySQL Administrator.

15.3. Using the Replication Status Section

Once your servers are configured, they will appear in the [Replication Status](#) section, as seen in the following figure:

Figure 15.1. The Replication Status section

Replication Status

Hostname	ID	Port	Kind	Status	Logfile	Log pos
192.168.1.10	1	3306	MASTER	Available	SERVER-bi...	31508506
SLAVE ONE	2	3306	SLAVE	New Instance		
Desktop			SLAVE	Not Available		

All servers are listed under the [Hostname](#) column, along with their server ID, Port, and current status.

New servers that are not currently being monitored will have an exclamation mark as their status symbol, and will be listed as a [New Instance](#) in the [Status](#) column of their entry.

Servers that are currently being monitored and are online will have a green status icon and will be listed as [Available](#) in the [Status](#) column of their entry.

Servers that are currently not in contact with the master server will have a red status icon and will be listed as [Not Available](#).

There will be a slight delay between a slave machine going offline and having its status updated to allow for an adequate timeout in the communications between master and slave.

To start monitoring a new server that is not currently being monitored, click the [Add Instance to Monitoring List](#) button. To stop monitoring a server that is currently being monitored, click the [Remove Instance from monitoring List](#) button.

Chapter 16. Catalogs

16.1. Introduction

This section deals with viewing catalogs, schemata, tables and their properties, creating databases, and performing a number of operations on tables. Existing catalogs and schemata are listed in the lower left area of the [sidebar](#). If there are many schemata on your MySQL server, the field with the magnifying glass icon may be handy for filtering the schemata you are looking for. Typing in `t` or `T`, for example, will set the filter to schemata whose names begin with a `t`. The filter works in a case-insensitive fashion.

Select the database that you wish to view. Right click on any database and choose the Refresh Schemata List option to refresh the database list. Clicking on a database will cause the [Schema Tables](#) and [Schema Indices](#) tabs to be populated with the tables and indexes of the selected database.

To create a new database, right click on an existing database and choose the Create New Schema option.

16.2. Schema Tables

The [Schema Tables](#) tab lists the different tables in the selected database, along with the rowcount, data length, and index length for those tables. For MyISAM tables, you also find the date and time when the table was last modified, under [Update Time](#). You can update the list of tables by pressing the [Refresh](#) button.

The tables are ordered by table name initially, but you may change that sorting by clicking on the appropriate headings ([Type](#), [Row Format](#), etc.).

To drop a table, right click on the table and choose the Drop Table option from the pop-up menu. To create a table you can either click the [Create Table](#) button at the bottom of the window, or right-click on a table and choose the Create Table option from the pop-up menu.

To edit a table's columns and indexes, right-click on the desired table and choose the Edit Table option from the pop-up menu. This will launch the MySQL Table Editor, which you can use to modify the table. For more information see [Chapter 21, The MySQL Table Editor](#).

To edit a table's actual data, right click on the table and choose the Edit Table Data option. This will launch the MySQL Query Browser and load the table's data into the result area of the MySQL Query Browser.

To perform operations on multiple tables select more than one table. For the selected tables, you can perform operations by either clicking the [Maintenance](#) button at the bottom of the table list or by right clicking on one of the selected tables and choosing from the Maintenance sub-menu of the pop-up menu:

The available options are:

- [Optimize Table](#): This corresponds to the `OPTIMIZE TABLE` SQL command and should be used if you have deleted a large part of a table or if you have made many changes to a table with variable-length rows (tables that have `VARCHAR`, `BLOB`, or `TEXT` columns). Deleted records are maintained in a linked list and subsequent `INSERT` operations reuse old record positions. You can use this command to reclaim the unused space and to defragment the datafile. Note that table optimization works for MyISAM and BDB tables only. For more information see [OPTIMIZE TABLE Syntax](#) [<http://dev.mysql.com/doc/refman/5.1/en/optimize-table.html>].
- [Check Table](#): This corresponds to the `CHECK TABLE` SQL command and is used to diagnose table problems. For more information see [CHECK TABLE Syntax](#)

[<http://dev.mysql.com/doc/refman/5.1/en/check-table.html>].

- **Repair Table:** This corresponds to the `REPAIR TABLE` SQL command and should be used in case of table problems. Note that this command works for MyISAM tables only. For more information see [REPAIR TABLE Syntax](http://dev.mysql.com/doc/refman/5.1/en/repair-table.html) [<http://dev.mysql.com/doc/refman/5.1/en/repair-table.html>].

Clicking on a table name will display its properties in the details area when the details area is active. To display the details area, click the `Details >>` button.

16.2.1. Table Status

In this tab, you find detailed information about the selected table. This information could also be retrieved by issuing a `SHOW TABLE STATUS LIKE 'tbl'` SQL command. For more information see [SHOW TABLE STATUS Syntax](http://dev.mysql.com/doc/refman/5.1/en/show-table-status.html) [<http://dev.mysql.com/doc/refman/5.1/en/show-table-status.html>].

16.2.2. Row Status

In this tab, you find detailed information about the rows of the selected table. This information can also be retrieved by issuing a `SHOW TABLE STATUS LIKE 'tbl'` SQL command. For more information see [SHOW TABLE STATUS Syntax](http://dev.mysql.com/doc/refman/5.1/en/show-table-status.html) [<http://dev.mysql.com/doc/refman/5.1/en/show-table-status.html>].

16.3. Schema Indexes

The [Schema Indices](#) tab lists the indexes of the selected database. Besides the index names and the table the index belongs to, you can also see the index type, whether values are unique, and whether NULL values are allowed. You can also get this information by issuing a `SHOW INDEX` SQL command in a command-line client (see [SHOW DATABASES Syntax](http://dev.mysql.com/doc/refman/5.1/en/show-databases.html) [<http://dev.mysql.com/doc/refman/5.1/en/show-databases.html>]). For more information about indexes, see [Column Indexes](http://dev.mysql.com/doc/refman/5.1/en/indexes.html) [<http://dev.mysql.com/doc/refman/5.1/en/indexes.html>].

To see the columns that form a given index, double click on the index.

Part II. MySQL Query Browser

Table of Contents

17. Introduction to the MySQL Query Browser	46
18. A Tour of the MySQL Query Browser	47
18.1. The Main Query Window	47
18.2. The Query Toolbar	48
18.2.1. The Navigation Buttons	48
18.2.2. The Query Area	48
18.2.3. The Action Buttons	49
18.3. The Advanced Toolbar	49
18.4. The Result Area	50
18.5. The Object Browser	50
18.5.1. The Database Browser	50
18.5.2. The Bookmark Browser	52
18.5.3. The History Browser	52
18.6. The Information Browser	52
18.6.1. The Syntax Browser	53
18.6.2. The Function browser	53
18.6.3. The Parameter Browser	53
18.6.4. The Transaction Browser	53
18.7. The Script Editor	53
18.7.1. The Script Debugging Buttons	55
18.7.2. The Script Area	56
19. Using The MySQL Query Browser	57
19.1. Entering and Editing Queries Manually	57
19.2. Building Queries	57
19.3. Creating Views	58
19.4. Navigating Result Sets Within the Result Area	58
19.5. Exporting Result Sets	59
19.6. Handling BLOB and TEXT Columns	59
19.7. Editing Result Sets Within the Result Area	60
19.8. Comparing Result Sets	60
19.9. Creating Master-Detail Views	61
19.10. Managing Stored Procedures and Functions with MySQL Query Browser ..	62
19.11. Editing Queries From a Development Tool	63

Chapter 17. Introduction to the MySQL Query Browser

The MySQL Query Browser is a graphical tool provided by MySQL AB for creating, executing, and optimizing queries in a graphical environment. Where the [MySQL Administrator](http://www.mysql.com/products/administrator/) [http://www.mysql.com/products/administrator/] is designed to administer a MySQL server, the MySQL Query Browser is designed to help you query and analyze data stored within your MySQL database.

While all queries executed in the MySQL Query Browser can also be run from the command-line using the `mysql` utility, the MySQL Query Browser allows for the querying and editing of data in a more intuitive, graphical manner.

MySQL Query Browser is designed to work with MySQL versions 4.0 and higher.

MySQL Query Browser is, to a large extent, the result of feedback MySQL AB has received from many users over a period of several years. However, if you find it's lacking some feature important to you, or if you discover a bug, please use our [MySQL Bug System](http://bugs.mysql.com) [http://bugs.mysql.com] to request features or report problems.


For MySQL Query Browser, there's a dedicated forum available on <http://forums.mysql.com/list.php?108>.

Chapter 18. A Tour of the MySQL Query Browser

18.1. The Main Query Window

Once you successfully connect to a MySQL server you are presented with the main query window of the MySQL Query Browser. All of the application's functionality is available through this window.

Figure 18.1. The MySQL Query Browser main query window



The main query window is divided up into several sections:

- *Query Toolbar:* The query toolbar is where you create and execute your queries. It is composed of three navigation buttons ([Go Back](#), [Next](#), [Refresh](#)), the query area, two action buttons ([Execute](#) and [Stop](#)), and a status indicator.
- *Advanced Toolbar:* The advanced toolbar contains three sets of buttons: the Transaction Buttons ([Start](#), [Commit](#), [Rollback](#)), the query management buttons ([Explain](#), [Compare](#)), and the query building buttons ([Select](#), [From](#), [Where](#), and so on.)

The Advanced Toolbar is only visible when the [Show advanced toolbars](#) and [Show composer buttons toolbar](#) options in the Browser Options screen are checked. See [Section 20.5, “The Browser Section”](#).

- *Result Area:* All query results are displayed in the result area. You can have multiple tabs active at one time, allowing you to work on multiple queries. The result area can be split vertically and horizontally for performing comparisons, and queries in different parts of a split result area can be joined together for master-detail analysis.
- *Object Browser:* The object browser is part of the sidebar and allows you to manage your databases, bookmarks, and history. You can choose which database and tables to query, add commonly-used queries to a collection of bookmarks, and browse through previously issued queries in order to use them again.
- *Information Browser:* The Information Browser is part of the sidebar and is used to look up built-in functions, and to manage queries that contain parameters. One example of this is when establishing a master-detail pair of queries. The data in the master query is available as parameters to the detail query.

Using the View menu, most of these sections can be displayed and hidden by checking and unchecking the elements you wish to view.

Each of these topics are covered in more detail in the sections that follow.

18.2. The Query Toolbar

All queries, whether generated automatically, graphically, or manually, appear in the [Query Toolbar](#). The simplest way to use the MySQL Query Browser is to type a query into the query area and click the execute button. A statement terminator such as `;` or `\G` is not required.

18.2.1. The Navigation Buttons

To the left of the query area are the navigation buttons. The navigation buttons allow you to browse through your query history so that you can review and re-execute a previously executed query. The query history is available through the [object browser](#).

Clicking the [Go Back](#) button loads the previous query in your history, whereas the [Go Next](#) button loads the following query. Only queries that execute without errors are added to your history.


As you navigate with the [Next](#) and [Go Back](#) buttons, the queries you navigate through are not executed unless you explicitly execute them by clicking the [Execute](#) button. The [Refresh](#) button re-executes the last executed query, which may or may not be the current query in the query area.

18.2.2. The Query Area

The query area is where the text of all queries and statements are displayed. By default, the query area is three lines high and automatically expands to a maximum of ten lines in height. For queries longer than ten lines, the query area is scrollable.

For additional space, you can press the **F11** key to maximize the query area. You can also do this by choosing Maximize Query Area from the View menu. When the query is maximized, line numbers are displayed and the query area can be resized by clicking and dragging the line that divides the query area from the result area. To restore the query area, press the **F11** key again.

Figure 18.2. The query area expanded to ten lines



The following commands are available by right clicking the query area: Cut (**Ctrl+X**), Copy (**Ctrl+C**), Paste (**Ctrl+V**), Clear (**Ctrl+N**), Open Query ... (**Ctrl+O**), and Save Query As ... (**Ctrl+S**). The Open and Save commands load the contents of an SQL file into the query area or save the contents of the query area into a text file.

18.2.3. The Action Buttons

To the right of the query area are the action buttons. The **Execute** button executes any queries in the query area, whereas the **Stop** button ceases execution.

If you click the down arrow below the **Execute** button, there are three potential execution options:

- **Execute (Ctrl+Enter)**: This executes the query and displays the results in the currently active result area.
- **Execute in new Tab (Ctrl+Shift+Enter)**: This executes the query and displays the results in a newly created result area.
- **Split Tab and Execute (Ctrl+Alt+Enter)**: This splits the result area horizontally and then displays the query results in the lower half of the result area.

To the right of the action buttons is the status indicator. The status indicator shows the standard MySQL logo when the MySQL Query Browser is idle, and displays an animated logo when the MySQL Query Browser is busy processing a query. Both icons can be seen in the following illustration:

Figure 18.3. The MySQL Query Browser status icons



18.3. The Advanced Toolbar

Below the query bar is the **Advanced Toolbar**. The **Advanced Toolbar** contains a set of buttons for transaction control, query management, and query building.

Figure 18.4. The advanced toolbar



The left panel of the advanced toolbar contains the transaction control buttons. From left to right, the buttons allow you to start, commit, and roll back a transaction. As when using the command-line client,

you can only use transactions with table handlers that support them (InnoDB for example). For more information on transactions, see [InnoDB Transaction Model and Locking](http://dev.mysql.com/doc/refman/5.1/en/innodb-transaction-model.html) [http://dev.mysql.com/doc/refman/5.1/en/innodb-transaction-model.html].

The center panel provides buttons for query management. The [Explain](#) button can be used to get the [EXPLAIN](#) output for the current query from the MySQL server, whereas the [Compare](#) button allows you to compare the results of two queries.

The next panel contains the query building buttons. You can use these buttons to build a query visually by clicking on the tables and columns that you wish to involve in your query, using specialized mouse pointers to indicate which part of the query the different fields and tables occupy.

The right panel contains the [Create View](#) button. More information on the [Create View](#) button can be found in the [Creating Views](#) section.

18.4. The Result Area

All query results are displayed in the result area. Within the result area you can use multiple tabs, and individual result areas can be split either vertically or horizontally. New tabs can be created either by choosing the Execute in new Tab option on the query bar, or by clicking the [New Tab](#) button at the top of the result area. In addition, you can also right click on the current result area and choose the Add new Tabsheet (Ctrl+T) option.

Individual tabs are closed by either right clicking within the result area and choosing the Remove Tab option, or by clicking the [X](#) icon on the tab you wish to close.

To split a result area, right click on it and choose either the Split Tab vertically or the Split Tab horizontally option. After splitting the result area you can then choose one half of the result area and view result sets within it. You can remove sections of the result area by right clicking on the section you wish to remove and clicking Remove Resultset.

When navigating a particularly large result set, you can press the **F12** key to maximize the result area or, alternately, you can choose the Maximize Tabsheets option of the View menu. Pressing **F12** again will restore the result area to its original size.

The result area can be used to review and edit the results of a query, with editing permitted as long as the query is based on a single table and there is sufficient key information to uniquely identify rows. To edit the contents of the result area you must enable edit mode through the use of the [Edit](#) button at the bottom of the result area. Any edits you make are not immediately applied, but instead you need to click the [Apply Changes](#) button next to the [Edit](#) button. Clicking the [Discard Changes](#) button throws away any changes you have made to the data.

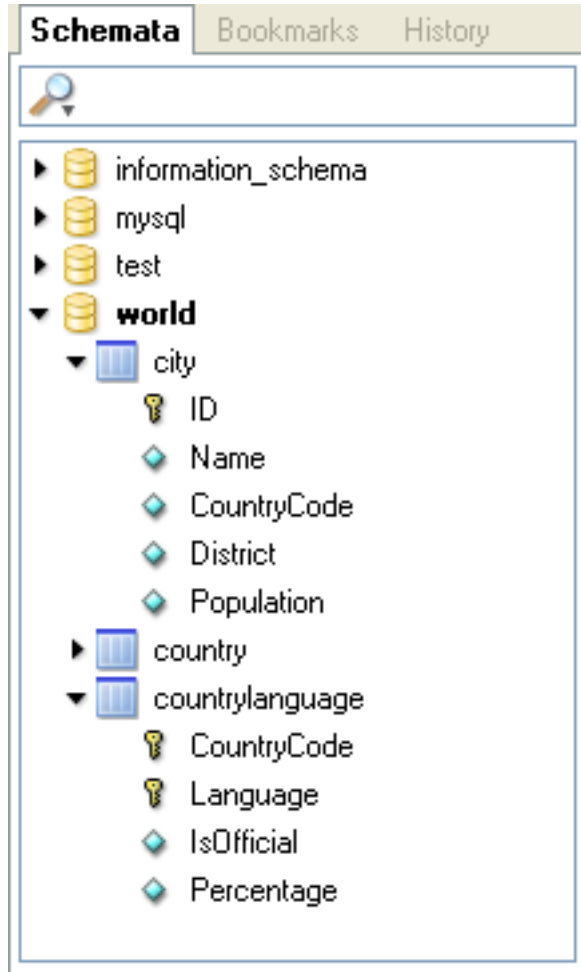
18.5. The Object Browser

The object browser allows you to browse your server's databases, your bookmarks, and your query history.

18.5.1. The Database Browser

The database (schemata) browser is the primary screen of the object browser. You can use the database browser not only to select tables and fields to query, you can also edit tables, create new tables and databases, and drop tables and databases. The database browser can be used to set the default database; this is required before you can issue queries against tables.

Figure 18.5. The database browser



The current default database will always be highlighted with bold text. To change the default database, right click on the desired database and choose the Make Default Schema option or double click on the desired database.

You can filter the databases shown by using the search bar at the top of the database browser. As you fill in the search bar, databases that do not match your search are hidden from view. You can clear the search bar by clicking the **X** button on the right side of the search bar. All databases will once again be displayed. You can specify what objects are filtered by clicking on the hourglass icon at the left side of the search bar and choosing the appropriate option. Options include [Schemata](#), [Schema Assets](#), [Columns/Indices](#) and [Custom Selection](#).

To view a database's tables, stored procedures, and views, click the black arrow on the left of the database name. You can view a table's columns by clicking on the black arrow to the left of the table name. Columns that form part of a primary key have a small key icon to the left of their name. Otherwise they have a blue diamond icon.

When you click the black arrow to the left of a view, the columns that form the view are shown. When you click the black arrow to the left of a stored procedure or stored function, the arguments of the stored procedure or function are shown.

To create a new database, right click within the database browser and choose the Create New Schema option. To create a new table, right click the database you wish to add a table to and choose the Create New Table option. To create a new view, right click within the database you wish to add a view to and choose the Create New View option. To create a new stored procedure or stored function, right click

within the database you wish to add to and choose the Create New Procedure / Function option. You can drop objects by right clicking on the object you wish to drop and choosing the appropriate option (Drop Schema, Drop Table, etc.).

You can edit an object by right clicking on it and choosing the appropriate edit option (Edit Table, Edit View, etc.). When you choose the Edit Table option, the MySQL Table Editor will be displayed with the selected table. For more information on editing tables, see [The MySQL Table Editor](#). When the Edit View, Edit Function, or Edit Procedure options are chosen, the `CREATE VIEW`, `CREATE PROCEDURE`, or `CREATE FUNCTION` statement will be shown in a new tab using the Script Editor. It is possible to edit all server functions/procedures at once by choosing the Edit All Stored Procedures / Functions option from the Script menu.

If you need the CREATE statement for any object, right click on the object and choose the Copy SQL to Clipboard option. The appropriate `CREATE` statement will be copied to the clipboard and can be used to recreate the selected object.

Note

Only the object itself is created with the Copy SQL to Clipboard command, the statements required to populate the object are not created.

18.5.2. The Bookmark Browser

You can place your more commonly used queries in bookmarks so that you can quickly retrieve them and re-use them later. To add a query to your bookmarks, highlight and drag it from the query area into the bookmark browser.

Your bookmarks can be organized into folders and subfolders to help with management of your queries. To add a new subfolder right click on an existing folder and choose the Create Bookmark Folder option. You can remove bookmarks and folders by right clicking on them and choosing the Delete Items option. After you confirm that you wish to delete the item it will be removed from your bookmark list.

Warning

If you delete a folder all items and subfolders within the folder are also deleted.

18.5.3. The History Browser

With the history browser you can browse through all the queries you have previously issued. To expand a given day's queries, double click on the day. To load a history item into the query area, double click on it or drag it to the query area.

You can remove history items by right clicking on them and choosing the Delete Selected History Entries option from the drop-down menu. You can also use the Clear History option to erase all history entries.

You can create bookmarks from history items by right clicking on a selected history item and choosing the Add History Item as Bookmark menu option.

18.6. The Information Browser

The Information Browser provides access to all information that is not directly related to actual data within your database. The parameter browser provides different local, global, and dynamic parameters that can help build your queries, whereas the syntax browser provides a convenient reference to MySQL query syntax. The Function browser is a quick reference to the various functions built into MySQL. The Transaction Browser lists all queries that make up a single transaction and serves as a history for a single transaction.

18.6.1. The Syntax Browser

The syntax browser provides a quick reference tool for determining proper syntax to use when constructing queries. By double clicking on the various types of syntax (SELECT, REPLACE, and so forth) you can bring up reference information using the in-line help system.

Reference information is displayed in a new tab within the result area.

To expand a given syntax category, double click on the category header.

18.6.2. The Function browser

The function browser provides a quick reference tool for usage of the built-in functions of the MySQL server. By double clicking on the various functions you can bring up reference information using the in-line help system.

To expand a given function category, double click on the category header.

Reference information is displayed in a new tab within the result area.

18.6.3. The Parameter Browser

The parameter browser contains all the local, global, and dynamic parameters that can be included in your query. Local parameters affect the query in the current query window only. Global parameters affect all queries. Dynamic parameters are generated automatically from existing queries.

Parameter categories can be displayed and hidden by clicking on the [Global Params](#), [Local Params](#) and [Dynamic Params](#) headings.

To set the value of a parameter, click on the value and press the **F2** key, or double click the value.

To add a new parameter, right click within the parameter browser and select the Add Parameter option.

To delete a parameter, right click on the parameter and choose the Delete Parameter option.

To change a Local parameter into a Global parameter, right click on the parameter and choose the Move Parameter to Global Parameters option.


18.6.4. The Transaction Browser

The Transaction Browser provides a list of all statements that have been executed in the current transaction. The Transaction browser will be displayed when you start a transaction using the transaction control buttons.

18.7. The Script Editor

The [Script Editor](#) provides an alternate interface for use when dealing with large SQL scripts that involve multiple SQL statements.

Figure 18.6. The script editor



The Script Editor displays within an individual tab in the [Result Area](#). When a Script Editor tab is active, the [Advanced Toolbar](#) buttons are replaced with script debugging buttons.


To open a script in the Script Editor, choose the Open Script ... option from the File menu. You can also press the **Ctrl + O** keys to open a SQL script.

To create a new Script Editor window without loading a script, choose the New Script Tab option from the File menu.

If you're experiencing problems with character sets (for example, Spanish accented characters, Russian, or Chinese characters aren't displayed properly), here's some advice:

Avoid ANSI encoding in your script files. Rather than using ANSI, save scripts in *UTF-8* encoding. Take care when loading script files generated by other tools. Many tools save files in ANSI encoding, which in turn means you have to open them as ANSI in Query Browser (see the [Files of type](#) combo box in the file open dialog):

Figure 18.7. The file open dialog



There are significant limitations to ANSI formatted files; the current system locale must be the same as the machine where the file originated. For example, you cannot store a script in ANSI on a French version of Windows and expect it to load correctly on a Spanish version of Windows. Because of this and many other limitations we strongly recommend always using Unicode for scripts. It doesn't matter whether you use UTF-8 or UTF-16, since both are Unicode transformation formats and can be read properly by any of the MySQL GUI tools. Using UTF, you could even write Spanish text with Chinese comments on a German version of Windows.

The MySQL Query Browser forum provides many contributions by users of non-Latin1 characters sets; see <http://forums.mysql.com/list.php?108>.

18.7.1. The Script Debugging Buttons

The following buttons are available when using the Script Editor:

- **Execute**: Execute the script from the beginning and do not stop execution until the end of the script is reached.
- **Continue**: Execute the script from the beginning (or current position) and stop for errors or break points.
- **Step**: Execute the next statement, functions will be executed but not stepped into.
- **Pause**: Stops script execution and highlights on the next statement to be executed.

- `Stop`: Stop execution of the script.

18.7.2. The Script Area

The current script being edited in the Script Editor is displayed in the `Script Area`. The Script Area features line numbering and syntax highlighting to make it easier to work with scripts.

Script lines that are marked with a blue dot to the left of the line are eligible for use as breakpoints. To set a line as a breakpoint, click the blue dot. You can also set a breakpoint by positioning the cursor in the line you wish to break at and clicking the `Toggle Breakpoint` option in the Script menu. To remove all breakpoints, choose the `Remove All Breakpoints` option of the Script menu.

To execute a selection of statements, highlight the statements you wish to execute and choose the `Execute Selection` option of the Script menu.

Chapter 19. Using The MySQL Query Browser

In this chapter we further describe the tools available in the MySQL Query Browser by using practical examples.

19.1. Entering and Editing Queries Manually

The query browser is most commonly used for executing queries and analyzing their results. The most direct way to create a query is to type it directly into the query area. As you type in your query, the SQL syntax portions of the query (SELECT, FROM, WHERE, and so on) are highlighted in blue.

As you enter your query, the query area will expand from an initial three lines in height to a maximum ten lines in height. For additional space, you can press the **F11** key to maximize the query area. You can also choose the Maximize Query Area of the View menu to maximize the query area. When the query is maximized, line numbers are displayed for the query, and the query area can be resized by clicking and dragging the line that divides the query area from the result area. To restore the query area, press the **F11** key again.

Once you have entered a query, click on the **Execute** button and your query results are displayed in the result area. You can also press **Ctrl+Enter** to execute the query. If there is an error in your query, an error area appears at the bottom of the result area and displays the relevant error message and error number.

In addition to loading the query results into the current active result area, you can also create a new result area for the results of your query or split your current result area and load results into the new section.

To execute the query and load the results into a new result area click the down-arrow below the **Execute** button and choose the Execute in new Tab option or press **Ctrl+Shift+Enter**.

To split the active result area and display the query results click the down-arrow below the **Execute** button and choose the Split Tab and Execute option or press **Ctrl+Alt+Enter**.


You must set a default database before you can query specific tables (though you can still perform queries that are not database-specific). You can set the default database at the connection screen, or by right-clicking on a database in the database browser and choosing Make Default Schema, or by choosing the Change Default Schema option from the File menu.

19.2. Building Queries

One feature of the MySQL Query Browser is the ability to build queries. Use the database browser to select the columns and tables you wish to query and have the query created automatically based on your choices.

The first step to building a query is to choose a table to query. Click and drag a table to the query area to start a query. For example, by dragging the **City** table to the query area, **SELECT * FROM City C** is added to the query area. You can also double-click a table to start a new **SELECT** query.

Figure 19.1. The table tool



When you select a table from the database browser and drag it over the query area, a table tool with some query composition actions is displayed. Drop the table you're dragging on the desired action and the query is modified accordingly. The following actions are possible:

- `SELECT` replaces the current statement with a `SELECT` query containing the dragged table.
- `Add Table` adds the dragged table to the list of tables in the current `SELECT` query
- `JOIN Table`: If a `SELECT` query is already in the query box, with a table in it, the new dragged table is added and the appropriate `WHERE` clauses to perform a `JOIN` will be added
- `LEFT OUTER JOIN` is the same as the previous, but does a `LEFT OUTER JOIN` instead of a `JOIN`
- `UPDATE` replaces the current statement with an `UPDATE` statement containing the dragged table
- `INSERT` replaces the current statement with an `INSERT` statement containing the dragged table
- `DELETE` replaces the current statement with a `DELETE` statement containing the dragged table

Tables are joined based on identical column names for MyISAM tables and foreign key information for InnoDB tables.

Once a table is selected, you can choose specific columns to query; click the `Select` button from the query building buttons on the button bar. Your mouse pointer is changed to a `Select` pointer, which you can use to choose columns from the database browser. For example, by clicking on the `Id`, `Name`, and `Country` fields of the `world` sample database, the query `SELECT C.Id, C.Name, C.Country FROM City C` is built in the query area.

Once you have chosen the columns you wish to query, you can use the other query building buttons to complete your query with `WHERE`, `GROUP BY`, and `ORDER BY` clauses. When a new section of the query is added with the query building buttons, the cursor in the query area is placed in position for editing; if you click a field with the `WHERE` pointer, the cursor is in position for you to type in the details of the `WHERE` clause.

You can change between the different query building pointers by clicking on the query building buttons in the button bar, or by using a combination of `Ctrl+Alt` and the first letter of the pointer you wish to use (`Ctrl+Alt+S` for `SELECT`, `Ctrl+Alt+W` for `WHERE`, and so forth.)

If the query building buttons are not visible, they can be displayed through the Browser Options. See [Section 20.5.1, "Display Options"](#) for more information.

19.3. Creating Views

The simplest way to create a view using the MySQL Query Browser is to use the `Create View` button. Execute a query that represents the view you wish to create. Once the query is executed, click the `Create View` button, provide a name for the view, and the view is created.

You can also create a view by right-clicking on the database you wish to add a view to in the database browser and choosing the `Create New View` option. After you enter a name for the view, a view template will be displayed in the script editor.

19.4. Navigating Result Sets Within the Result Area

Once you have successfully executed a query you can then view and manipulate the result set within the

result area.

You can navigate the result area using the arrow keys, tab key, and PageUp/PageDown keys. The Home and End keys can be used to move to the first and last column within a given row. Your current position within the result set is shown in the bottom-left corner of the application window. The [First](#) and [Last](#) buttons at the bottom of the result area can be used to move to the first and last rows of the result set.

NULL data will be indicated with a special [NULL](#) icon to differentiate NULL data from empty strings. BLOB fields will be empty and will have a special [BLOB](#) icon.

When navigating a particularly large result set, you can press the **F12** key to maximize the result area. You can also choose the Maximize Tabsheets option of the View menu to maximize the result area. Pressing **F12** again will restore the result area to its original size.

To search for a particular value within the result set, click the [Search](#) button. To see additional search options, click the [Details >>](#) button. The following options are available:

- [Case Sensitive](#): The search is performed in a case-sensitive manner. By default, searches are not case-sensitive.
- [Whole Words Only](#): The search does not allow partial matches. By default, partial matches are allowed (i.e. [Edm](#) will match [Edmonton](#)).
- [Search From Top](#): The search begins with the first row in the result set.
- [Search From Cursor](#): The search begins from the currently selected row.
- [Search all Text / Columns](#): The search will involve all columns of the result set.
- [Search Only in Selected Text / Column](#): The search will only involve the currently selected column(s) of the result set.
- [Search Up](#): The search will scan the result set, starting at the designated start location and moving upwards.
- [Search Down](#): The search will scan the result set, starting at the designated start location and moving down.


To perform a search and replace operation on a result set, click the [Replace](#) tab in the search dialog window. The options for performing a replace operation are the same as for a search. To replace a single instance of a string, click the [Replace](#) button. To replace all instances of a string, click the [Replace All](#) button.

19.5. Exporting Result Sets

You can export any result set from MySQL Query Browser by right-clicking within the result set and choosing an option from the Export Resultset sub-menu. You can choose to export the result set in [CSV](#), [XML](#), [HTML](#), Microsoft Excel [XLS](#) or [PLIST](#) formats.

19.6. Handling BLOB and TEXT Columns

The MySQL Query Browser provides functionality for dealing with [BLOB](#) and [TEXT](#) columns through a series of special icons.

Figure 19.2. The BLOB handling icons

These icons appear in any **BLOB** or **TEXT** columns in your result set. From left to right the following icons are available:

- **Open File:** This icon looks like a file folder and is used to open a file for loading into the field.
- **View:** This icon looks like a magnifying glass and is used to open the field viewer to view the contents of the field. The field viewer can be used to view **TEXT** fields and **BLOB** fields that contain images.
- **Edit:** This icon looks like a pencil and opens the field viewer in edit mode, allowing you to make changes to the data and apply the changes to the result set.
- **Save:** This icon looks like a floppy disk and is used to save the contents of a **TEXT** or **BLOB** field into a file.
- **Clear:** This icon looks like an X within a black box and is used to clear the contents of a **TEXT** or **BLOB** field.

Only the **View** and **Save** icons are visible if you have not enabled edit mode. See [Section 19.7, “Editing Result Sets Within the Result Area”](#) for information on editing result sets.

The functionality represented by the icons is also available by right-clicking on the field. The blob handling icons can be displayed and hidden through the display options. See [Section 20.5.1, “Display Options”](#) for more information.

19.7. Editing Result Sets Within the Result Area

When a query is based on a single table and there is sufficient key information, the result set can be edited from within the result area. To edit a result set click the **Edit** button. If the **Edit** button is not active, your result set is not editable.

Once in edit mode you can insert, update, and delete rows from the result set. Navigate the fields using the **Tab** and arrow keys, and press **Enter** to edit the content of a field. You can also double-click a field to make it editable. When editing a field, use the tab key to move to the next editable field. All edited fields are highlighted in blue for easy identification.

To add rows to the result set, scroll to the blank row at the bottom of the result area and fill in the fields. All new rows are highlighted in green.

To delete a row right-click on the row and choose the Delete Row option. All deleted rows are highlighted in red.

Changes made to the result set are not applied immediately, but instead are cached until the **Apply Changes** button is pressed. You can abort your edits with the **Discard Changes** button. Exiting edit mode without choosing to apply or discard your changes prompts you to apply or discard your work.

19.8. Comparing Result Sets

You can compare result sets graphically with the MySQL Query Browser, allowing you to easily determine where rows have been inserted, updated, or deleted.

To compare two result sets, execute the first of the queries you wish to compare. Once the result set has loaded, right-click on the result set and choose the Split Tab Horizontally option. Load your second query into the new section of the result area and click the [Compare](#) button to compare the two result sets.

When you have activated the compare mode, both result sets will scroll in unison, both vertically and horizontally. Rows are matched for comparison, with blank rows added when one set has a row that the other set lacks.

If one result set has a row that the other result set does not have, that row is highlighted in green. The other result set has a blank row inserted that is highlighted in red. If both result sets have a matching row, but individual fields are different, those fields are highlighted in blue.

In order to successfully compare two result sets, you need two queries with matching column names and column order. For the MySQL Query Browser to match rows, primary keys must be defined in the tables used.

19.9. Creating Master-Detail Views


MySQL Query Browser makes it easy to view data that is in a master-detail (one to many) relationship. This can be useful for viewing customer/order data, group/member data, and so forth.

To display a master-detail view, first query the master table in a new result area. For example, using the `world` sample database, you could issue a query like `SELECT Code, Name, Continent FROM Country`.

Once you have created the master result set, right-click on the result area and choose the Split Tab Vertically option. For the detail query, create a query that uses the dynamic parameters available from the master query, as seen in the parameter browser. In this example, we could query the `City` table in the following way: `SELECT Id, Name FROM City WHERE Country = :Code`.

The colon character indicates to the MySQL Query Browser that you are adding a dynamic parameter, and a pop-up list of available parameters should appear within the query area, allowing you to use the arrow keys to choose a parameter. You can also type the name of the parameter that you would like to use to link the detail query.

Figure 19.3. The dynamic parameter list



Once you execute the detail query, it will automatically refresh any time you change the active row in the master query, allowing you to quickly see all detail rows as you navigate the master result set.

You can repeat this process, splitting the result area yet again, allowing you to have detail result sets *for* detail result sets.

Note


You can only continue splitting a result set in the same direction as the previous split but you cannot, for example, split vertically and then horizontally.

19.10. Managing Stored Procedures and Functions with MySQL Query Browser

When used with MySQL version 5 and higher, the MySQL Query Browser supports creating, editing, and viewing stored procedures and functions.

Stored procedures and stored functions are displayed in the database browser with a special icon to distinguish them from regular tables and views. Click the arrow to the left of a stored procedure or stored function to display the parameter list for that procedure or function.

Figure 19.4. Stored procedure in database browser



To edit a stored procedure or stored function, right-click on it in the database browser and choose the Edit Procedure or Edit Function option. This opens a new script editor tab with the selected procedure/function displayed. Once you have finished editing, click the [Execute](#) button above the script area or click the Execute option of the Script menu to update the procedure/function.

To create a new stored procedure or function, choose the Create Stored Procedure / Function option from the Script menu. Enter the desired procedure/function name and click either the [Create PROCEDURE](#) or [Create FUNCTION](#) button. A template similar to the following will be created:

```
DELIMITER \\  
  
DROP PROCEDURE IF EXISTS `test`.`new_proc`\  
CREATE PROCEDURE `test`.`new_proc` ()  
BEGIN  
  
END\  
  
DELIMITER ;
```

After typing the procedure/function, click the [Execute](#) button above the script area or click the Execute option of the Script menu to create it.

To edit all stored procedures/functions at once, choose the Edit All Stored Procedures / Functions option of the Script menu. A new script editing tab will be created, containing all the stored procedures and

stored functions for the current default database.

To remove an existing stored procedure or stored function, right-click on it in the database browser and choose the Drop Procedure or Drop Function option.

19.11. Editing Queries From a Development Tool

In order to help programmers optimize and troubleshoot their queries more efficiently, the MySQL Query Browser can copy queries from application code using your favorite integrated development environment (IDE).

Note

This functionality is only available for the Windows version of MySQL Query Browser.

The following PHP code will be used as an example:

```
$$SQL = "SELECT Id, Name, Country FROM City" .  
        "WHERE Name LIKE $cityname";
```

To copy the listing into the MySQL Query Browser, copy the block of code (including the assignment portion), right click within the query area of the MySQL Query Browser, and choose the Paste Clipboard Content as PHP Code option. The non-query portions of the highlighted area will be stripped and the query will be pasted into the query area.

The dynamic elements of the query are converted into local parameters, visible in the [parameter browser](#):

```
SELECT Id, Name, Country FROM City  
WHERE Name LIKE :cityname
```

To set the value of a local parameter, highlight the value in the parameter browser and press **F2**. You can also double-click on the value to edit it. The value you assign will be used when the query is executed.

After editing the query, right-click within the query area and choose the Copy Query as PHP Code option. The surrounding PHP code will be re-inserted along with the modified query. This functionality allows you to edit queries quickly while programming.

Part III. Common Administrator and Query Browser Components

Table of Contents

20. Options Dialog	66
20.1. Introduction	66
20.2. General Options	67
20.3. Connections	67
20.4. Editors	69
20.5. The Browser Section	70
20.5.1. Display Options	71
20.5.2. Query Options	71
20.5.3. Various	72
20.6. Administrator	72
20.6.1. User Administration	72
20.6.2. Health Graphs	72
20.6.3. Backup	72
21. The MySQL Table Editor	74
21.1. Introduction	74
21.2. The Main Editor Window	74
21.3. The Columns And Indices Tab	75
21.3.1. The Column Editor	75
21.3.2. The Detail Area	75
21.4. The Table Options Tab	77
21.5. The Advanced Options Tab	77
21.5.1. The Various Section	77
21.5.2. The Row Options Section	77
21.5.3. The Storage Options Section	77
21.5.4. The Merge Table Options Section	78
21.5.5. The Table RAID Settings Section	78
21.6. Applying Your Changes	78

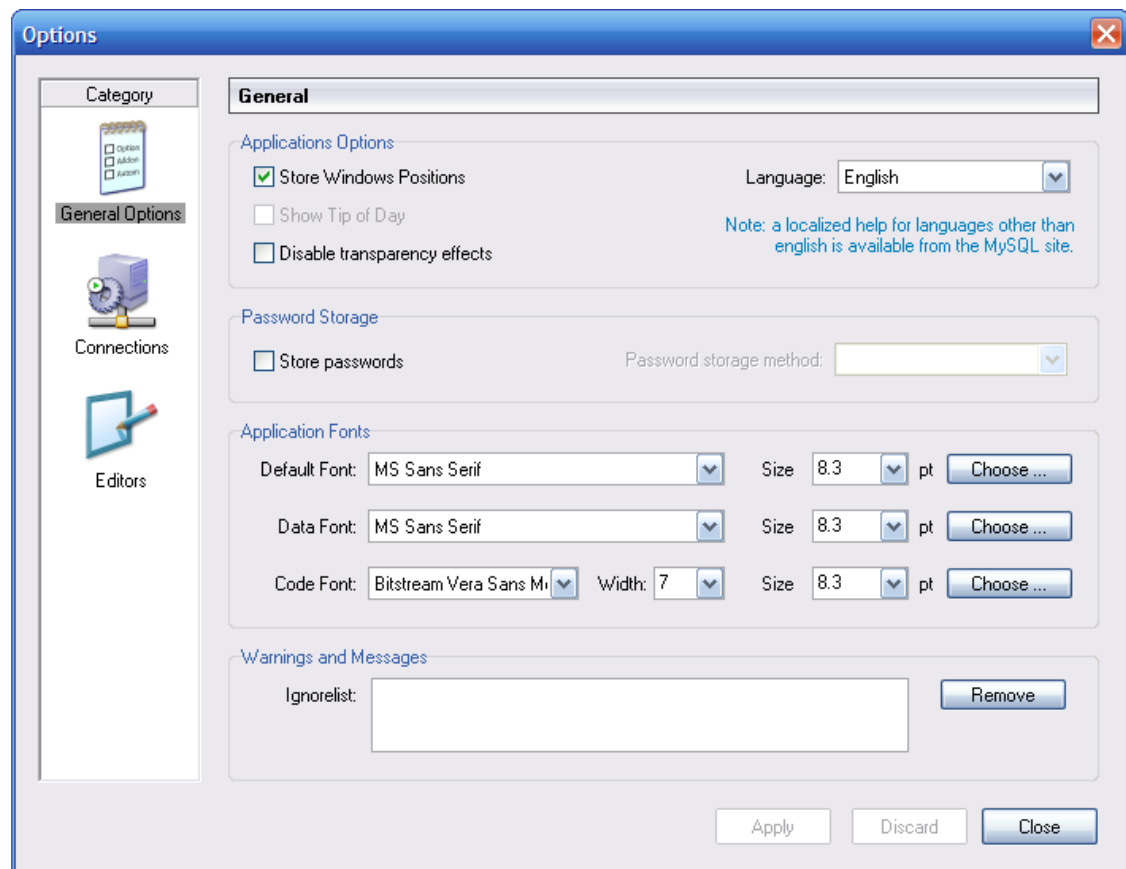
Chapter 20. Options Dialog

20.1. Introduction

The **Options** dialog allows you to configure connection profiles, general program settings, and more. You can open the **Options** dialog window using one of the following methods:

- In the connection dialog window, click the **...** button.
- In the main application window select **Options ...** from the Tools menu.

Figure 20.1. Options dialog



In the sidebar of the dialog, you can select the section you wish to configure. There are three action buttons in the lower right corner of the window:

- **Apply**: Applies and saves changes.
- **Discard**: Discards any changes you have made.
- **Close**: Closes the **Options** dialog window. If you have not applied or discarded your changes, you will be prompted to do so.

20.2. General Options

The [General Options](#) section allows you to specify a number of settings that are valid for all graphical MySQL applications.

- [Store Window Positions](#): When an application is started the next time, its latest window position will be reused.
- [Show Tip of Day](#): If checked, a pop-up window with the tip of the day appears at program startup.
- [Store Passwords](#): If checked, passwords are stored in the user's connection profile. You can specify the password storage method:

[Plaintext](#): Passwords are stored without encryption; this can be insecure. [Obscured](#): Passwords are encrypted using a weak algorithm. This encryption method is operating system independent. [OS Specific](#): Use the default encryption method provided by your operating system. The default option is [Plaintext](#).

- [Language](#): Select the interface language. The default is [English](#).
- [Default Font](#): The font used for all application text.
- [Data Font](#): The font used for all query and table data displayed.
- [Code Font](#): The font used for all queries entered by the user..
- [Ignorelist](#): Whenever you check the [Do Not Show This Message Again](#) option on error and message prompts, they are added to this list. If you would like a particular message to be shown again, click the [Remove](#) button after selecting the message from the list.

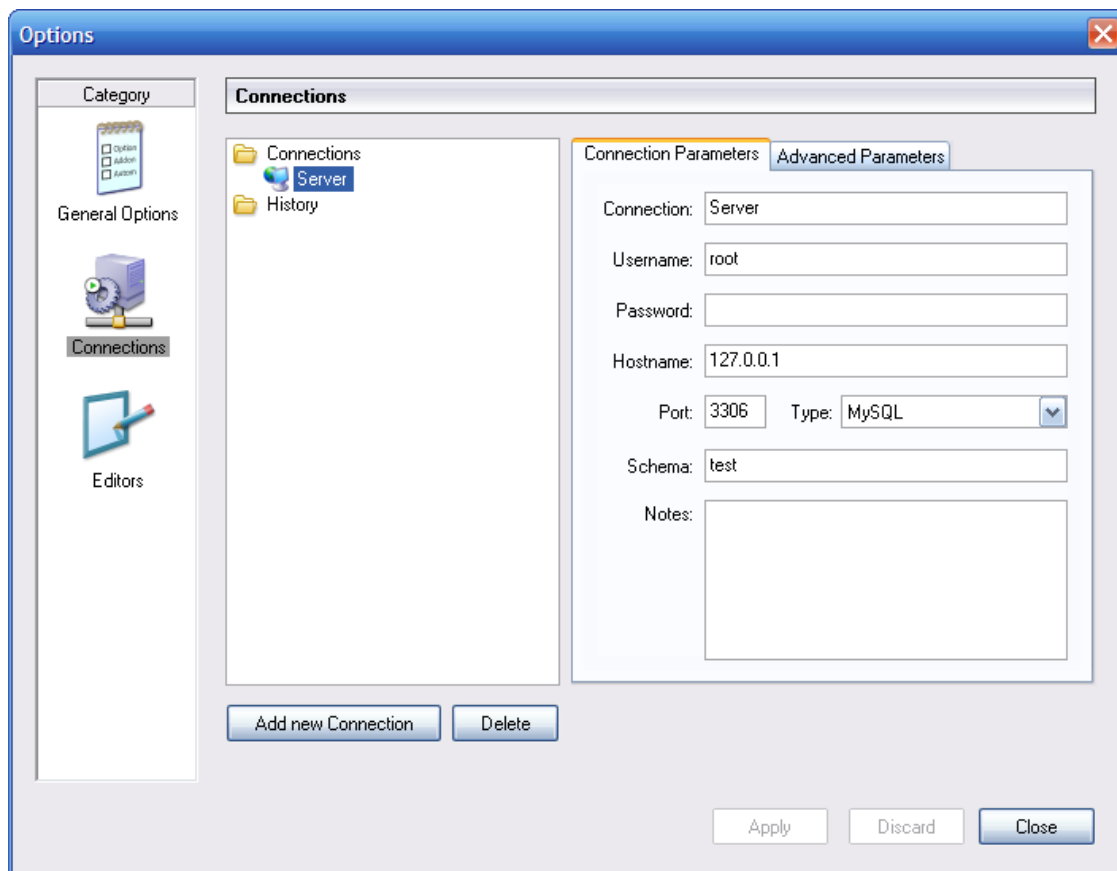
You may change the font and font size of any of the application fonts by clicking the [Choose ...](#) button to the right of the font.

20.3. Connections

The [Connections](#) section allows you to create, edit, and delete connection profiles. The center box displays a list of currently available profiles, together with a history of connections that were made without being stored in a profile. You can collapse or expand both the [Connections](#) and [History](#) trees by double clicking them.

Connections are automatically added to the [History](#) tree whenever you establish a connection to a MySQL server without using one of the profiles stored under the [Connections](#) tree. They do not appear in the drop-down box of the [Connection](#) dialog, but you can use any of them by manually typing their name into the [Connection](#) box of the [Connection](#) dialog.

Figure 20.2. Options dialog: Connections



To edit an existing connection profile, click on its name and change the values that appear in the **Connection Parameters** and **Advanced Parameters** tabs, then click on the **Apply Changes** button to save your changes.

When you select a connection profile from either the **Connections** or **History** trees, the **Connection Parameters** tab displays the following fields:

- **Connection:** The connection profile label. This is the name by which you refer to the profile and that appears in the **Connection** drop-down box of the Connection dialog. It may contain any characters, including spaces. Choose distinctive names so that you can easily tell which profiles they refer to. The names can help you distinguish connections to different MySQL servers, or connections as different MySQL users to a given server.
- **Username:** The username used to connect to the MySQL server.
- **Password:** The password used to connect to the MySQL server. Note that passwords are not stored in the connection profile, unless you specify otherwise in the **General Options** section.
- **Hostname:** The name of the host machine where the MySQL server runs, or its IP address.
- **Port:** The TCP/IP port that the MySQL server listens to on the host machine.
- **Type:** Specifies the protocol used to connect to the database server. The default protocol is **MySQL** (which uses the native MySQL protocol).
- **Schema:** The default database for a connection when using the MySQL Query Browser.

- **Notes:** You can use this field to enter comments or additional information describing the connection profile.

Note

The [Advanced Parameters](#) tab is not available on all platforms. Advanced parameters can still be configured in the Connection dialog. Use the [Details ...](#) button to display the [Advanced Connection Options](#).

When you select a connection profile from either the [Connections](#) or [History](#) list, the [Advanced Parameters](#) tab displays the following checkboxes:

- [Use compressed protocol](#): If checked, the communication between the application and the MySQL server will be compressed, which may increase transfer rates. This corresponds to starting a MySQL command-line tool with the `--compress` option.
- [Return number of found rows, not number of affected rows](#): By default, MySQL returns the number of rows changed by the last `UPDATE`, deleted by the last `DELETE` or inserted by the last `INSERT` statement. When this option is checked, the server returns the number of rows matched by the `WHERE` statement for `UPDATE` statements.
- [Ignore spaces after function names, make them reserved words](#): Normally, any reference to a function name in an SQL statement must be followed immediately by an opening parenthesis. If this option is checked, spaces may appear between the function name and the parenthesis, like this:

```
COUNT (*)
```

Enabling this option has the effect that function names become reserved words. This option corresponds to starting a MySQL command-line tool with the `--ignore-spaces` option.

- [Allow interactive_timeout seconds of inactivity before disconnect](#): Normally, the connection is closed by the MySQL server after a certain period of inactivity on the client side. This period can be set with the `interactive_timeout` variable. If checked, the server will not close the connection unless the period of inactivity exceeds the value set by `interactive_timeout`. This corresponds to starting a MySQL command-line tool with the `--connect-timeout=seconds` option.
- [Enable LOAD DATA LOCAL handling](#): By default, the `LOCAL` option of the `LOAD DATA` statement is disabled for security reasons. Enabling this option will allow you to load data from the local machine (the machine where the client GUI application is running). This option corresponds to starting a MySQL command-line tool with the `--local-infile=1` option. (Note that this option is ineffective unless the MySQL server allows `LOCAL` handling.)

20.4. Editors

The [Editors](#) section is used to configure options specific to the different editors available within the MySQL GUI Suite. At the time of writing this is limited to the [MySQL Table Editor](#).

- [Show SQL command before applying changes](#): Toggles whether the MySQL Table Editor will show you the `CREATE TABLE` or `ALTER TABLE` statement it is about to execute for confirmation when you click [Apply Changes](#).
- [All columns Not Null per default](#): Determines whether the MySQL Table Editor will


designate columns as being `NOT NULL` by default when creating new columns.

- `All integer columns unsigned per default`: Sets whether integer columns are declared `UNSIGNED` by default when creating new columns.
- `Default storage engine`: Sets the storage engine to be assigned to newly created tables. This value is independent of the default storage engine of the MySQL server.
- `PK Naming`: Dictates the format that should be used to name `PRIMARY KEY` columns that are automatically generated. The `%tablename%` portion will be replaced with the name of the appropriate table.
- `Index Naming`: Sets the name automatically generated for new indexes. The `%nr%` string will be replaced with an automatically incrementing number.
- `FK Naming`: Configures the name used when creating new foreign keys.
- `PK Datatype`: The datatype used when creating a `PRIMARY KEY` column.
- `Def. data type`: The default datatype assigned to all new columns that are not part of a `PRIMARY KEY`.

20.5. The Browser Section

The `Browser` section allows you to specify options that relate to the MySQL Query Browser. The `Browser` section is divided into three areas: `Display Options`, `Query Options`, and `Vari-`
`ous`.

Figure 20.3. The query browser options dialog



20.5.1. Display Options

The display options affect the appearance of the MySQL Query Browser.

- **Show advanced toolbars:** Display or hide the [The Advanced Toolbar](#).
- **Hide the tab when only one page is open:** Determines whether or not to display a tab icon at the top of the [Result Area](#) if there is only one result area.
- **Toolbars use gradient background:** Enables and disables use of gradient effects in the background of the [Query Toolbar](#).
- **Show field overlay images for long VARCHAR, TEXT, and BLOB fields:** Toggles on and off the BLOB management icons.
- **Show composer buttons toolbar:** Toggles on and off the query composition buttons on the Advanced Toolbar.

20.5.2. Query Options

Enforce queries to be editable by adding primary key columns to the select: Adds `PRIMARY KEY` columns to a query that does not explicitly `SELECT` them to ensure that the results of a query can be edited. The `PRIMARY KEY` columns will not be displayed in this case, but will still be stored.

[Open resultset in associated application after export](#): Causes the associated application to be opened after choosing the Export Resultset option from the right click menu of the [Result Area](#).

20.5.3. Various

[Associate sql/query files with Query Browser](#): Configures Windows to associate all .sql files with the MySQL Query Browser. As a result, all such files will be opened with the MySQL Query Browser by default.

20.6. Administrator

The [Administrator](#) section allows you to specify options that affect the behavior of the MySQL Administrator. The [Administrator](#) section is divided into three sub-sections that will be discussed further in the sections that follow.

20.6.1. User Administration

These options affect the interface of the [User Administration](#) section of MySQL Administrator.

- [Show Global Privileges](#): If checked, the [Global Privileges](#) tab will be visible in the [User Administration](#) section. In that tab, you can grant users global privileges (as opposed to privileges on a database, table, or column level only). This may affect security. See [The MySQL Access Privilege System](#) [<http://dev.mysql.com/doc/refman/5.1/en/privilege-system.html>]. If unchecked, the [Global Privileges](#) tab will not be visible in the [User Administration](#) section.
- [Show Table/Column Privileges](#): In most situations, it is sufficient to grant either global or database privileges for users. You can hide the [Table/Column Privileges](#) tab in the [User Administration](#) section by enabling this checkbox.

20.6.2. Health Graphs

These options affect the appearance of the [Health Graphs](#) section of MySQL Administrator.

- [Use Peak Level Indicator](#): This will display a thin yellow line on the health graphs indicating the highest value recorded by the graph.
- [Reset peak Level after a number of ticks](#): If this box is not checked, the [Peak Level Indicator](#) will be permanently positioned at the highest value ever recorded on the health graph. With this box checked, the [Peak Value Indicator](#) will be reset after a period of time indicated by the [Time after the Peak level resets](#) option. Use this option if you would like to know the peak level over an interval rather than the peak level overall.
- [Time after the Peak level resets](#): When the [Reset peak Level after a number of ticks](#) option is set, this determines the amount of time (in ticks) that will pass before the [Peak Level Indicator](#) will be reset.

20.6.3. Backup

The backup options configure how MySQL Administrator performs backup operations.

- [Add Date/Time to Backup Files](#): This determines whether the date and time of the backup operation are added to the name of the backup file. See [the backup section](#) for more information.

Chapter 21. The MySQL Table Editor

21.1. Introduction

The MySQL Table Editor is a component of both the MySQL Query Browser and MySQL Administrator, enabling the visual creation and modification of tables.

The MySQL Table Editor can be accessed from the MySQL Query Browser by right clicking on a table within the database browser and choosing the Edit Table option, or by right clicking on a database within the database browser and choosing the Create New Table option.

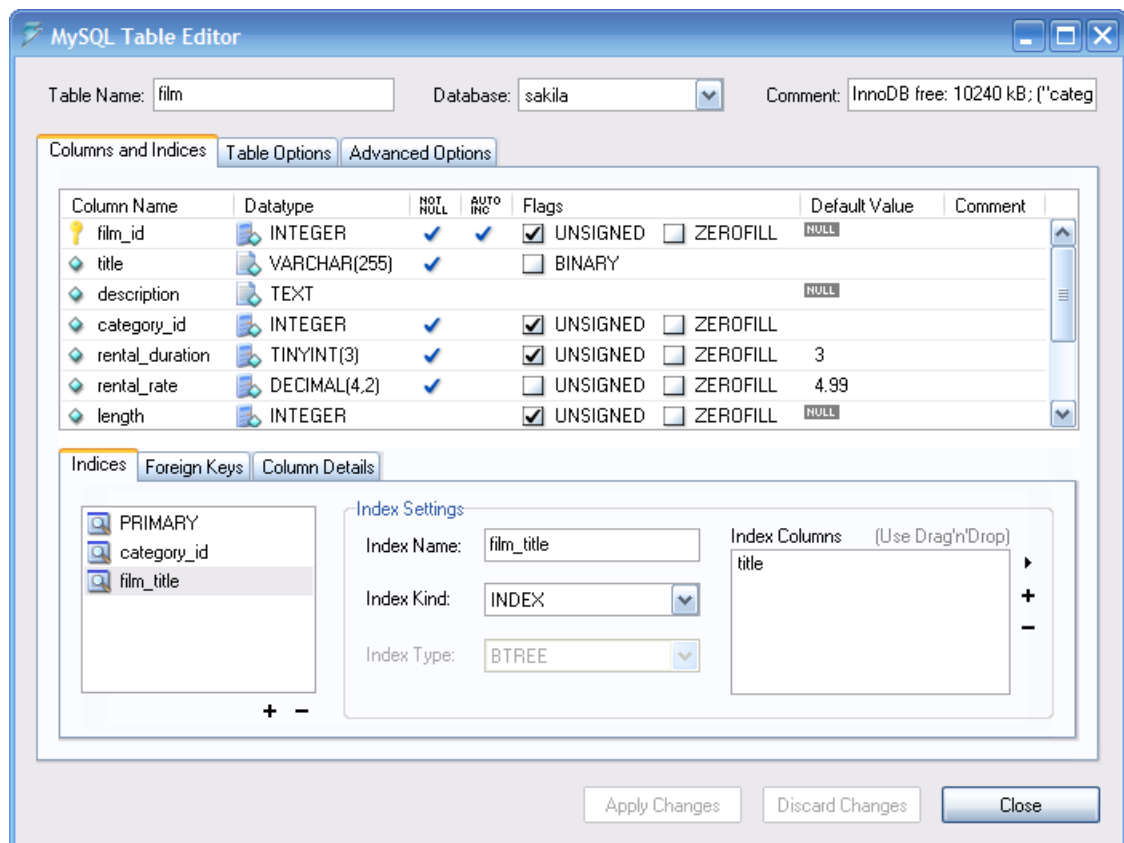
The MySQL Table Editor can be accessed from MySQL Administrator through the Catalogs screen. Once you have selected a database, right click on a table and choose the Edit Table option from the drop-down menu. You can also select a table and click the [Edit Table](#) button to access the MySQL Table Editor.

21.2. The Main Editor Window

The MySQL Table Editor consists of a work space divided into three tabs, some general information prompts, and three action buttons.

Regardless of the active tab, you can always edit the table name and the table comment.

Figure 21.1. The MySQL Table Editor



The tabbed area is divided into three sections:

- *Columns and Indices:* Use the Columns and Indices tab to create and modify the table's column and index information. You can also create FOREIGN KEY relationships using this tab.
- *Table Options:* Use the Table Options tab to choose the storage engine and default character set used by the table.
- *Advanced Options:* Use the Advanced Options tab to configure options such as the per-table storage directory, MERGE and RAID table options, and table/row length options.

Each of these areas are discussed in further detail in the following sections.















21.3. The Columns And Indices Tab

The [Columns](#) and [Indices](#) tab can be used to display and edit all column and index information for your table. Using this tab, you can add, drop, and alter columns and indexes.

21.3.1. The Column Editor

You can use the column editor to change the name, data type, default value, and other properties of your table's columns.

Figure 21.2. The column editor

Column Name	Datatype	NOT NULL	AUTO INC	Flags	Default Value	Comment
 film_id	 INTEGER	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> UNSIGNED <input type="checkbox"/> ZEROFILL	<input type="text" value="NULL"/>	
 title	 VARCHAR(255)	<input checked="" type="checkbox"/>		<input type="checkbox"/> BINARY		
 description	 TEXT				<input type="text" value="NULL"/>	
 category_id	 INTEGER	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/> UNSIGNED <input type="checkbox"/> ZEROFILL		
 rental_duration	 TINYINT(3)	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/> UNSIGNED <input type="checkbox"/> ZEROFILL	<input type="text" value="3"/>	
 rental_rate	 DECIMAL(4,2)	<input checked="" type="checkbox"/>		<input type="checkbox"/> UNSIGNED <input type="checkbox"/> ZEROFILL	<input type="text" value="4.99"/>	
 length	 INTEGER			<input checked="" type="checkbox"/> UNSIGNED <input type="checkbox"/> ZEROFILL	<input type="text" value="NULL"/>	

To change the name, data type, default value, or comment of a column, double click on the value you wish to change. The value becomes editable and you can complete your changes by pressing the **Enter** key.

To modify the flags on a column ([UNSIGNED](#), [BINARY](#), [ASCII](#), and so on) check and uncheck the boxes corresponding to the flag you wish to change.

To the left of the column name is an icon that indicates whether the column is a member of the primary key. If the icon is a small key, that column belongs to the primary key, otherwise the icon is a blue diamond. To add or remove a column from the primary key, you can click on the icon.

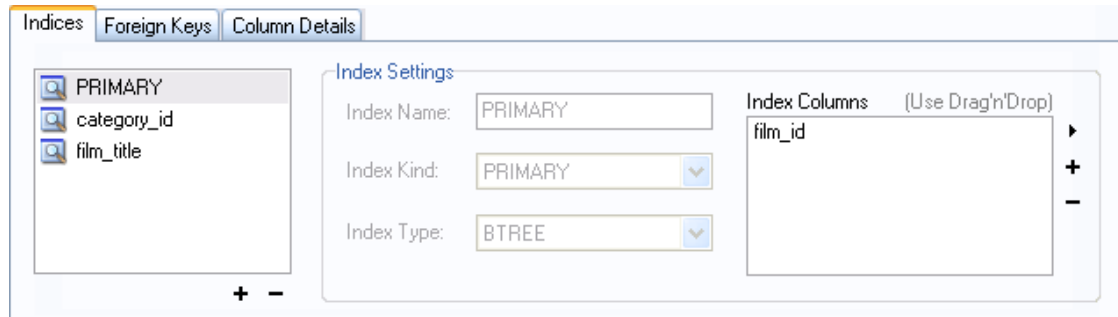
21.3.2. The Detail Area

The detail area of the [Columns](#) and [Indices](#) tab is divided into three tabs used for modifying the index, foreign key, and miscellaneous information regarding your columns.

21.3.2.1. The Indices Tab

The **Indices** tab holds all index information for your table. You can add, drop, and modify indexes using the indices tab, in combination with the column editor.

Figure 21.3. The index editor



To add an index, click the + (plus) button below the index list. The MySQL Query Browser prompts for an index name and the new index is created. To drop an index, select the index and click the - (minus) button.

Use the **Index Name** and **Index Kind** dialog boxes to modify the name and type (**UNIQUE**, **FULL-TEXT**, and so on) of the index.

To add columns to an index, either click and drag the column to the **Index Columns** box or select the column you wish to add and click the + (plus) button to the right of the **Index Columns** box. You can remove a column from the index by selecting the column and clicking the - (minus) button to drop the column from the index.

To change an index to only refer to a column prefix (such as with the CHAR and VARCHAR string data types), select the index column you wish to prefix in the **Index Columns** box and then click the arrow icon to the right of the **Index Columns** box. Select the **Set Index Column Length** option from the drop-down menu that appears.

21.3.2.2. The Foreign Keys Tab

The **Foreign Keys** tab is divided into two sections, one with a list of foreign keys and one with various dialog boxes for foreign key settings.

To add a foreign key, click the + (plus) button below the foreign key list. The MySQL Query Browser prompts for a foreign key name and the new foreign key is created.

To drop a foreign key, select the foreign key and click the - (minus) button below the foreign key list.

You can modify the name of the foreign key, its **ON DELETE**, and its **ON UPDATE** actions using the dialog boxes provided in the **Foreign Key Settings** section of the tab.

To establish a foreign key relationship, choose a table from the **Ref. Table** drop-down list. The columns that can be referenced as foreign keys are listed in the area below, which has two sections. To add a column (of the table you are editing) you can either double click the area below **Column** and then select a column from the drop-down list, or drag a column from the column editor to the **Column** section. To add a column (of the reference table) double click the area below **Reference Column** and then select a column from the drop-down list.

21.3.2.3. The Column Details Tab

The **Column Details** tab provides an interface for setting the parameters of a column without using

the table interface of the column editor.

All settings that are available in the [Column Editor](#) are also available in the Column Details tab, and in addition you can also configure the column character set and column default collation from the Column Details tab.

21.4. The Table Options Tab

The [Table Options](#) tab allows you to change the storage engine and default character set of your table. The potential storage engines are listed, along with a brief summary of each storage engine's features and strengths.

To change the storage engine for your table, click on the radio button next to the desired storage engine.

To change the default character set or collation of your table, choose a new option from the drop-down list of available character sets.

21.5. The Advanced Options Tab

The [Advanced Options](#) tab is used to configure table options that would be considered outside the standard set of options that most users designate when creating and modifying tables.

The [Advanced Options](#) tab is divided into several sub-sections, each of which is described in the upcoming sections of this manual.

For descriptions of most options set using the Advanced Options tab, see [Figure 29.1, “Creating a table”](#).

21.5.1. The Various Section

The [Various](#) section of the [Advanced Options](#) tab contains options for you to set the [PACK KEYS](#) behavior, the table password, the initial [AUTO_INCREMENT](#) value, and the delayed key update behavior.

The [AUTO_INCREMENT](#) and delayed key update behaviors apply only to MyISAM tables.

21.5.2. The Row Options Section

The [Row Options](#) section can be used to configure options such as the row format, checksum use, and the row size parameters needed for large tables.

To set the row format, choose the desired row format from the drop-down list. See [MyISAM Table Storage Formats](#) [<http://dev.mysql.com/doc/refman/5.1/en/myisam-table-formats.html>], for more information on the different row formats that are available. This option only applies to MyISAM tables.

When you expect a table to be particularly large, use the [Avg Row Length](#), [Min Rows](#), and [Max Rows](#) options to enable the MySQL server to better accommodate your data. See [Figure 29.1, “Creating a table”](#), for more information on how to use these options.

21.5.3. The Storage Options Section

The [Storage Options](#) section is used to configure a custom path to the table storage and data files. This option can help improve data integrity and server performance by locating different tables on different hard drives.

This option is only available for MyISAM tables and is not available for servers running under the Windows operating system.

21.5.4. The Merge Table Options Section

The [Merge Table Options](#) section is used to configure MERGE tables in MyISAM. To create a MERGE table, select MERGE as your storage engine in the [Table Options Tab](#) and then specify the tables you wish to MERGE in the [Union Tables](#) dialog.

You can also specify the action the server should take when users attempt to perform INSERT statements on the merge table. See [The MERGE Storage Engine](#) [<http://dev.mysql.com/doc/refman/5.1/en/merge-storage-engine.html>], for more information on MERGE tables.

21.5.5. The Table RAID Settings Section

The [Table RAID Settings](#) section allows you to configure RAID support for MyISAM tables. RAID allows MyISAM table data files to grow larger than the 2GB/4GB size limit imposed by some operating systems.

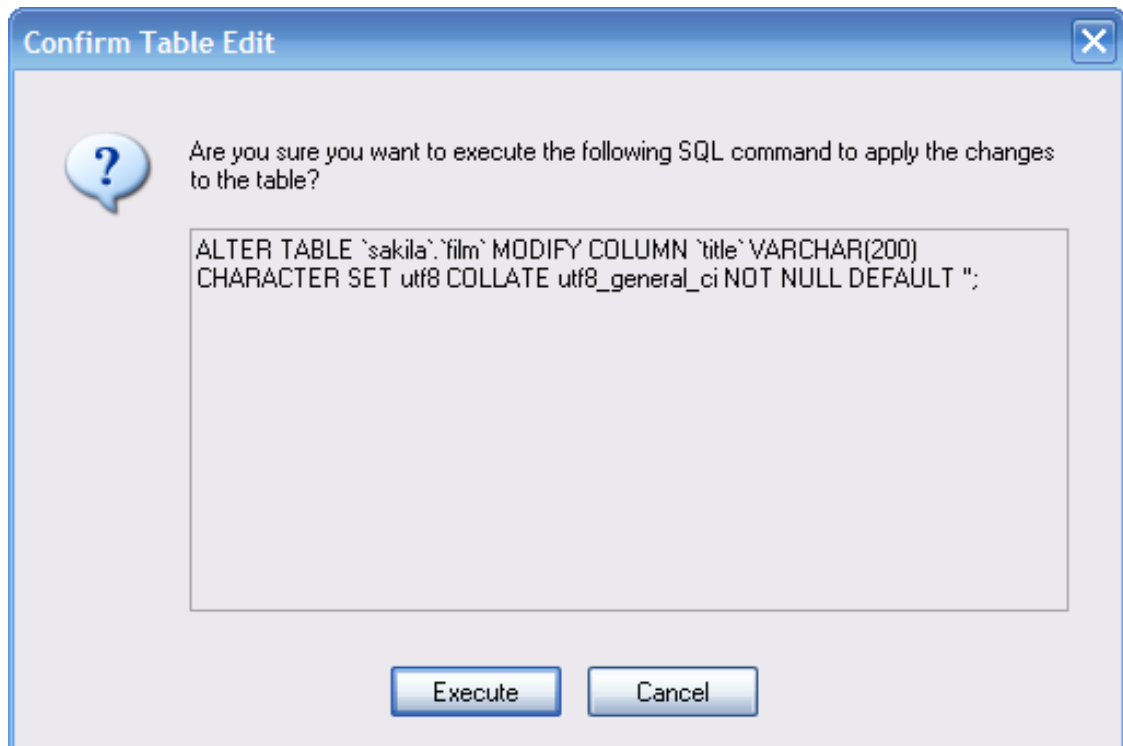
For more information on using RAID support with MyISAM, see [Figure 29.1, “Creating a table”](#).

21.6. Applying Your Changes

The changes you make with the MySQL Table Editor are not immediately applied but are instead queued to be applied in batches after you have made all your edits.

To apply the changes you have made, click the [Apply Changes](#) button. The [Confirm Table Edit](#) dialog will appear.

Figure 21.4. The Confirm Table Edit dialog



You can click the [Execute](#) button to confirm the changes and have them applied, or click the [Cancel](#) button to cancel the changes (the table editor window is redisplayed with your changes intact). You can also click the [Discard Changes](#) button in the main MySQL Table Editor window to discard all changes you have made.

You can also copy the proposed changes to the clipboard for further editing by highlighting the ALTER TABLE or CREATE TABLE statement, right-clicking and choosing Copy from the drop-down menu.

Part IV. MySQL Migration Toolkit

Table of Contents

22. Introduction to the MySQL Migration Toolkit	82
23. An Overview of the Migration Process	83
24. The Migration Process In-Depth	86
24.1. Introduction	86
24.2. The Welcome Screen	86
24.3. The Configuration Type Screen	86
24.4. The Source Database Screen	87
24.4.1. Microsoft Access	88
24.4.2. Microsoft SQL Server	88
24.4.3. Oracle	89
24.4.4. MySQL	91
24.4.5. Saving Connection Information	92
24.5. The Target Database Screen	92
24.6. The Connect to Server Screen	93
24.7. The Source Schema Selection Screen	94
24.8. The Reverse Engineering Screen	95
24.9. The Object Type Selection Screen	96
24.9.1. Migrating a Sub-Set of an Object Type	97
24.10. The Object Mapping Screen	98
24.10.1. GRT Object	99
24.10.2. Table Objects	99
24.11. The Meta Migration Screen	100
24.12. The Manual Editing Screen	100
24.13. The Object Creation Options Screen	102
24.14. The Creating Objects Screen	103
24.15. The Data Mapping Options Screen	104
24.16. The Bulk Data Transfer Screen	105
24.17. The Summary Screen	106
24.18. Saving the Current Application State	107
25. Extending The MySQL Migration Toolkit	108
25.1. Introduction	108
25.2. Architecture of the MySQL Migration Toolkit	108
25.3. The Modular Migration Process	108
25.4. Tools Required to Extend the MySQL Migration Toolkit	109
26. Preparing a Microsoft Access Database for Migration	110

Chapter 22. Introduction to the MySQL Migration Toolkit

The MySQL Migration Toolkit is a graphical tool provided by MySQL AB for migrating schema and data from various relational database systems to MySQL.

MySQL Migration Toolkit is designed to work with MySQL versions 5.0 and higher.

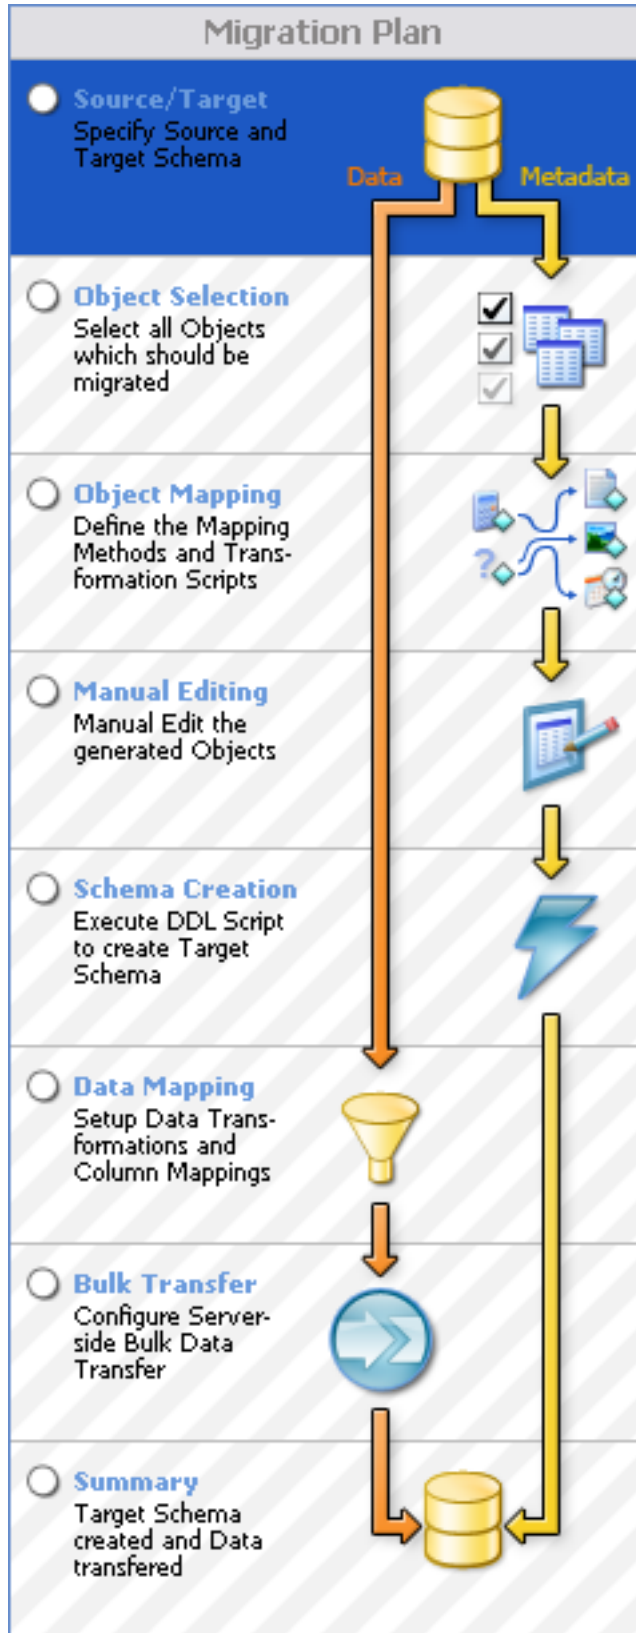
The current version of MySQL Migration Toolkit is a release candidate. While every effort has been made to ensure it is free of bugs, it should not be used in a production environment. Always back up your data before using the MySQL Migration Toolkit. If you find it's lacking some feature important to you, or if you discover a bug, please use our [MySQL Bug System](http://bugs.mysql.com) [http://bugs.mysql.com] to request features or report problems.

At the moment, MySQL Migration Toolkit is only available for Windows.

Chapter 23. An Overview of the Migration Process

Migrating data from an external RDBMS to MySQL is an eight step process :

Figure 23.1. The MySQL Migration Toolkit Migration Plan



- **Source/Target Selection:** In the first step you specify the connection parameters for the source and target database servers.
- **Object Selection:** In the second step you select the objects (tables, views, stored procedures) that will be migrated.
- **Object Mapping:** In the third step you choose the method used for mapping and transforming the objects.
- **Manual Editing:** In the fourth step you can manually edit the new objects to ensure a proper transformation.
- **Schema Creation:** In the fifth step the MySQL Migration Toolkit creates the transformed object on the target MySQL server.
- **Data Mapping:** In the sixth step you specify any changes that need to be made to the data as it is migrated.
- **Bulk Transfer:** In the seventh step the MySQL Migration Toolkit transfers the data from the source server to the target server.
- **Summary:** In the eighth and final step the MySQL Migration Toolkit creates a summary report of the migration process for you to review.

Each of these sections will be covered in more detail in the coming chapters.

Chapter 24. The Migration Process In-Depth

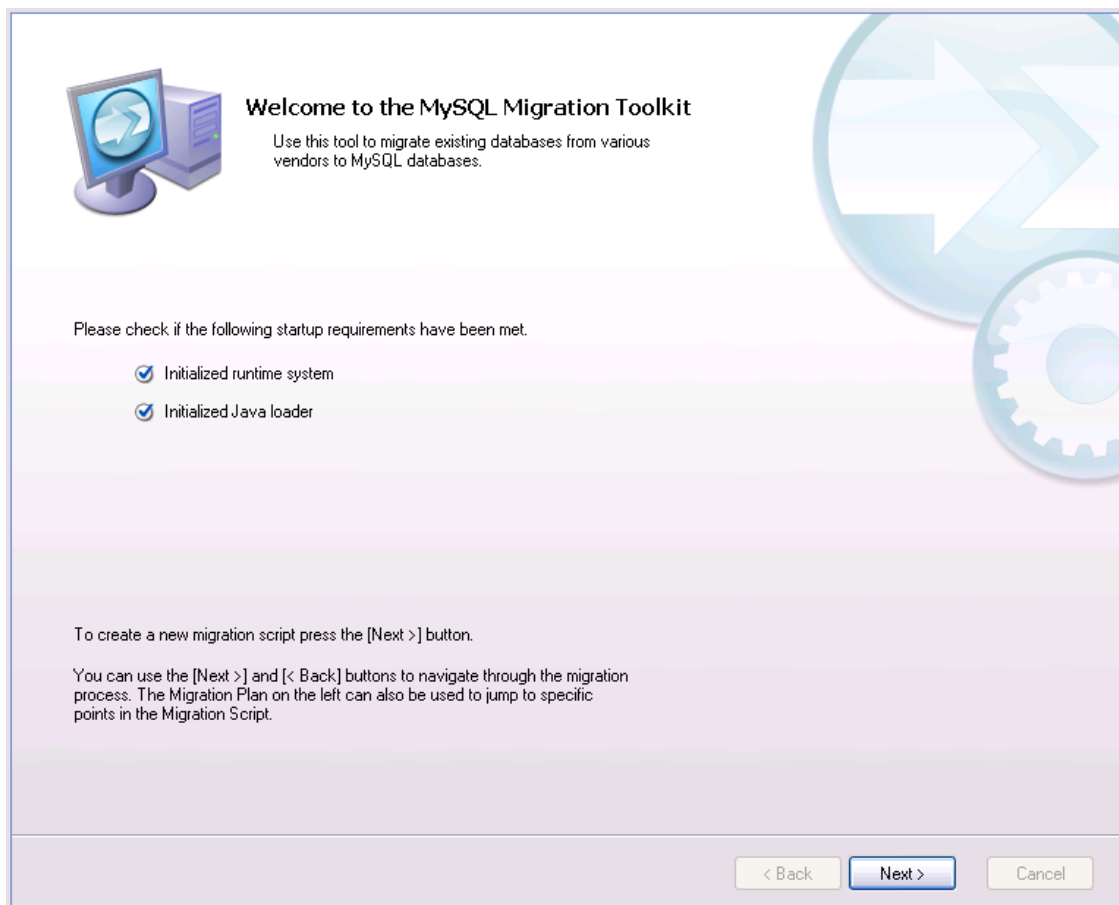
24.1. Introduction

In this chapter we will cover the different steps of the MySQL Migration Toolkit in depth. The different steps will be covered in the order that they appear in the MySQL Migration Toolkit.

24.2. The Welcome Screen

The first screen of the MySQL Migration Toolkit is the Welcome Screen:

Figure 24.1. The MySQL Migration Toolkit welcome screen

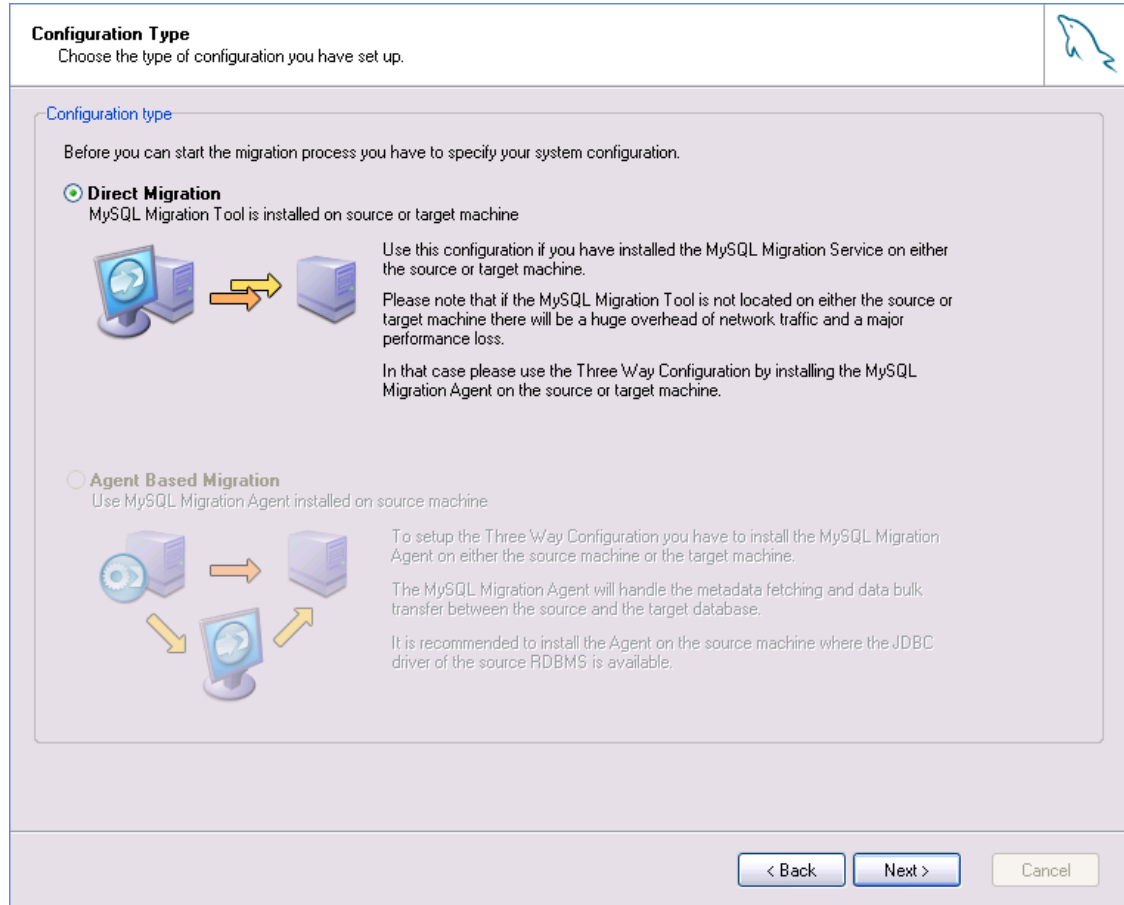


The welcome screen indicates the progress of loading the various components of the MySQL Migration Toolkit. If you encounter any errors on the welcome screen you should close the MySQL Migration Toolkit and confirm that you have properly installed the Java Runtime Environment. See [Introduction](http://dev.mysql.com/doc/migration-toolkit/en/mysql-migration-toolkit-installation-introduction.html) [http://dev.mysql.com/doc/migration-toolkit/en/mysql-migration-toolkit-installation-introduction.html] for information on downloading and installing the Java Runtime Environment.

24.3. The Configuration Type Screen

The Configuration Type screen allows you to choose between a Direct Migration and a Agent-Based Migration:

Figure 24.2. The Configuration Type screen



Use the [Direct Migration](#) if the MySQL Migration Toolkit is installed on either the source or target machine. The Direct Migration should not be used when the MySQL Migration Toolkit is not located on either the source or target machine, as it will create a large amount of network traffic and will result in decreased performance.

Use the [Agent-Based Migration](#) when migrating between two server machines that do not support the use of the MySQL Migration Toolkit. The MySQL Migration agent should be installed on the source machine before using the Agent-Based Migration.

24.4. The Source Database Screen

Use the Source Database screen to select the source RDBMS used in the migration and to specify the connection parameters.

The Source Database screen's appearance will vary depending on the type of source database selected.

All Data Source screens will feature a [Details](#) button that can be used to expose the [Advanced Settings](#) panel. The Advanced Settings panel can be used to manually specify a JDBC driver and JDBC connection string for your migration session.

24.4.1. Microsoft Access

The Source Database screen appears as follows when you select Microsoft Access as the source database:


Figure 24.3. Source database – Microsoft Access

Source Database
Select the source database you want to migrate from.

Source Database Connection

Database System: Select the target RDBMS you want to migrate from.

Connection Parameters

 **MS Access Server**
MS Access over JDBC-ODBC bridge

Connection: + - Select a stored connection or use [+] to store or [-] to remove.

Database File: ... MS Access database file.

Username: Name of the user to connect with.

Password: The user's password.

Specify the path to the `.mdb` database file in the **Database File** field and specify the database username and password information if applicable.

You must make special modifications to your Access database file before it can be used with MySQL Migration Toolkit. Please see [Chapter 26, Preparing a Microsoft Access Database for Migration](#) for further information.

24.4.2. Microsoft SQL Server

The Source Database screen appears as follows when you select **MS SQL Server** as the source database:

Figure 24.4. Source database – Microsoft SQL Server

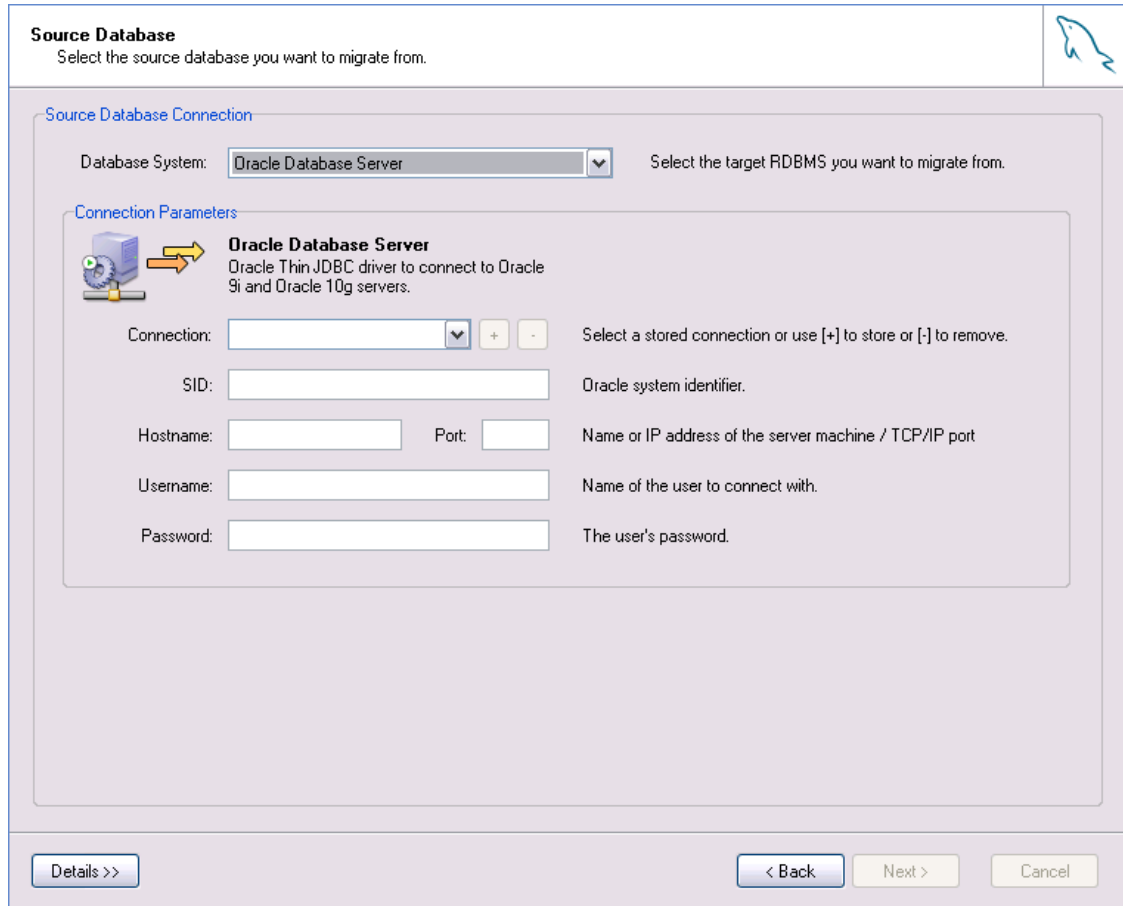
The screenshot shows a 'Source Database Connection' dialog box. At the top, 'Database System' is set to 'MS SQL Server' with a dropdown arrow and the instruction 'Select the target RDBMS you want to migrate from.' Below this is a 'Connection Parameters' section. It features an icon of a server with a yellow arrow pointing right, followed by the text 'MS SQL Server' and 'JDBC driver to connect to MS SQL Server 2000.' There are four input fields: 'Connection:' with a dropdown and '+' '-' buttons; 'Hostname:' with 'localhost' entered; 'Port:' with '1433' entered; and 'Password:' which is empty. To the right of each field is a descriptive label: 'Select a stored connection or use [+] to store or [-] to remove.', 'Name or IP address of the server machine / TCP/IP port', 'Name of the user to connect with.', and 'The user's password.'

Specify the hostname, username, and password to connect to the source Microsoft SQL Server to connect and click [Next](#).

24.4.3. Oracle

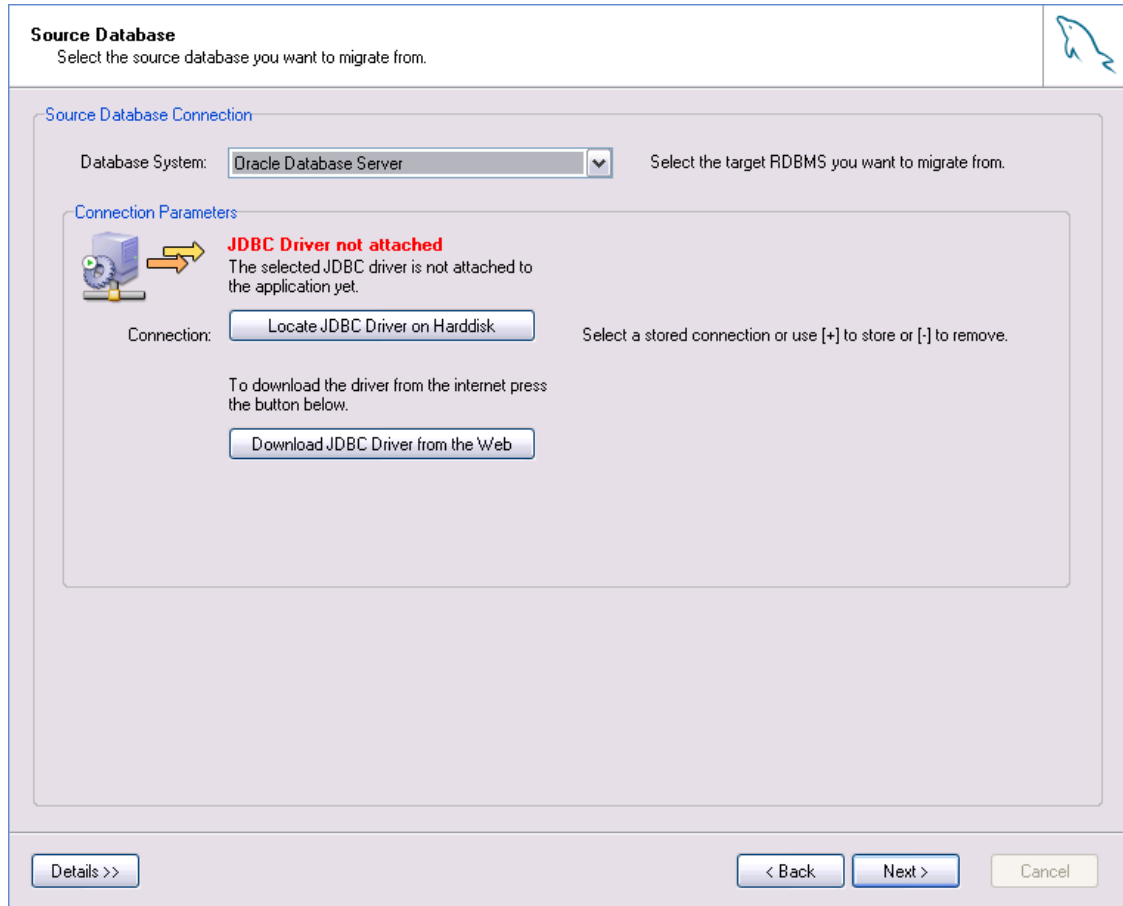
The Source Database screen appears as follows when you select Oracle as the source database:

Figure 24.5. Source database – Oracle



If you encounter the following database selection screen, it means that you do not have the appropriate JDBC driver for Oracle installed:

Figure 24.6. Oracle JDBC driver not attached



If the Oracle JDBC driver is present on your system, click the [Locate JDBC Driver on Hard-disk](#) button to attach the driver.

If the Oracle JDBC driver is not present on your system, click the [Download JDBC Driver from the Web](#) button to download it. Download the `ojdbc14.jar` file and then attach it by clicking on the [Locate JDBC Driver on Harddisk](#) button.

After attaching the Oracle JDBC driver you need to restart the MySQL Migration Toolkit.

24.4.4. MySQL

The Source Database screen appears as follows when you select MySQL as the source database:


Figure 24.7. Source database – MySQL

Source Database
Select the source database you want to migrate from.

Source Database Connection

Database System: Select the target RDBMS you want to migrate from.

Connection Parameters

 **MySQL Server**
MySQL JDBC driver to connect to MySQL 4.0, 4.1 and 5.0 servers.

Connection: + - Select a stored connection or use [+] to store or [-] to remove.

Hostname: Port: Name or IP address of the server machine / TCP/IP port

Username: Name of the user to connect with.

Password: The user's password.

24.4.5. Saving Connection Information

After entering the connection information for the source database, click the **+** button to save the connection information.

When prompted, enter a name for the connection information and click the **Ok** button to save the connection information for later re-use.

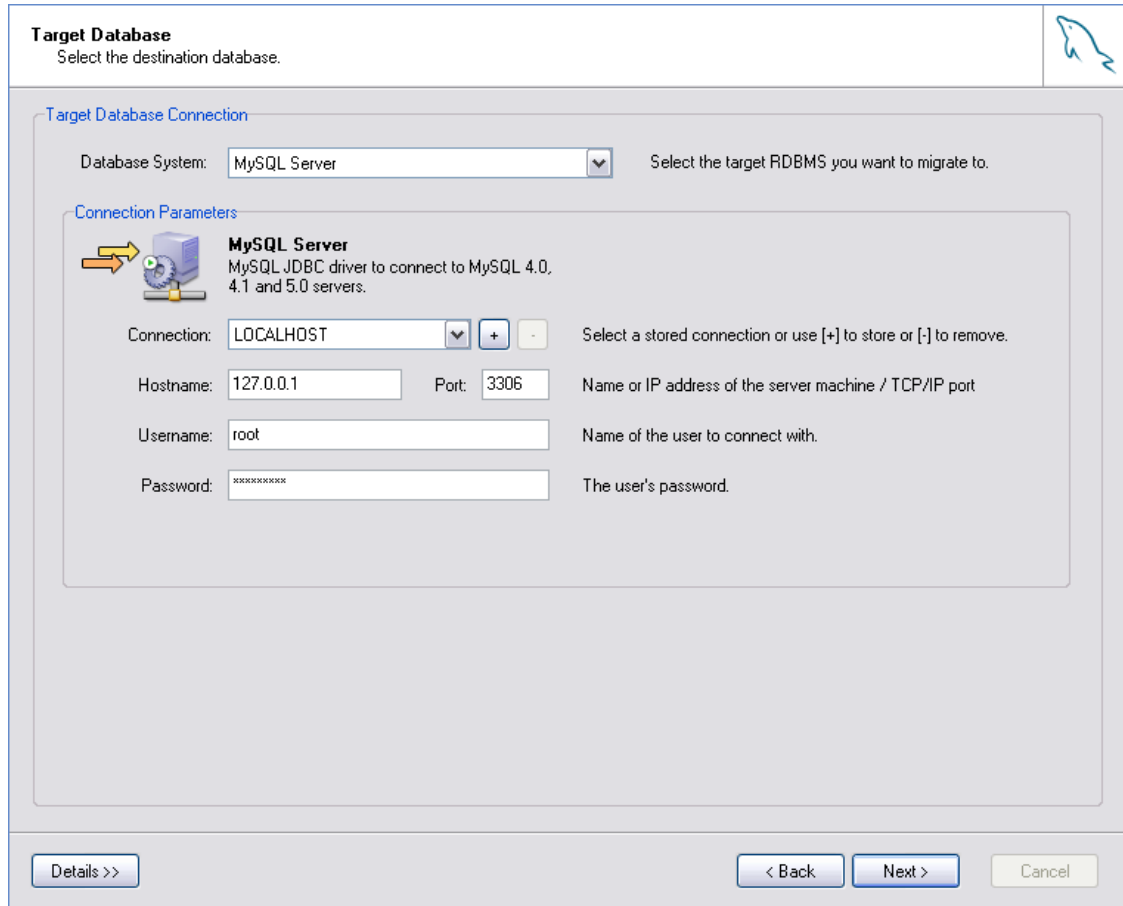
You can discard saved connection information by selecting the saved connection from the **Connection** drop-down list and clicking the **-** button.

24.5. The Target Database Screen

Use the Target Database screen to select the target RDBMS used in the migration and to specify the connection parameters.

The target Database screen uses an interface that is identical to that of the [Source Database screen](#):

Figure 24.8. Target Database – MySQL

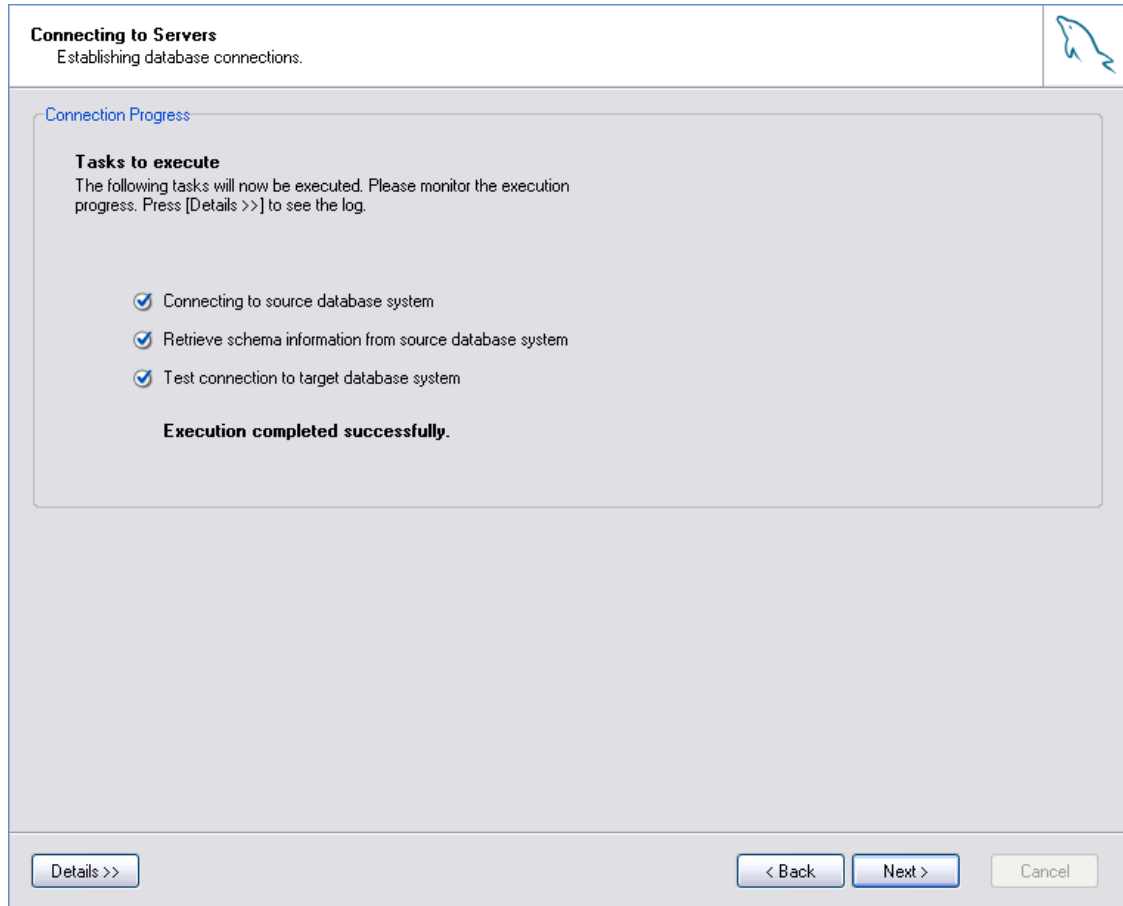


Target database support for the MySQL Migration Toolkit is currently limited to MySQL 4.1 and MySQL 5.0.

24.6. The Connect to Server Screen

After you specify your source and target database servers, the MySQL Migration Toolkit will connect to each server and retrieve the schema information from the source server:

Figure 24.9. The Connect to Servers screen



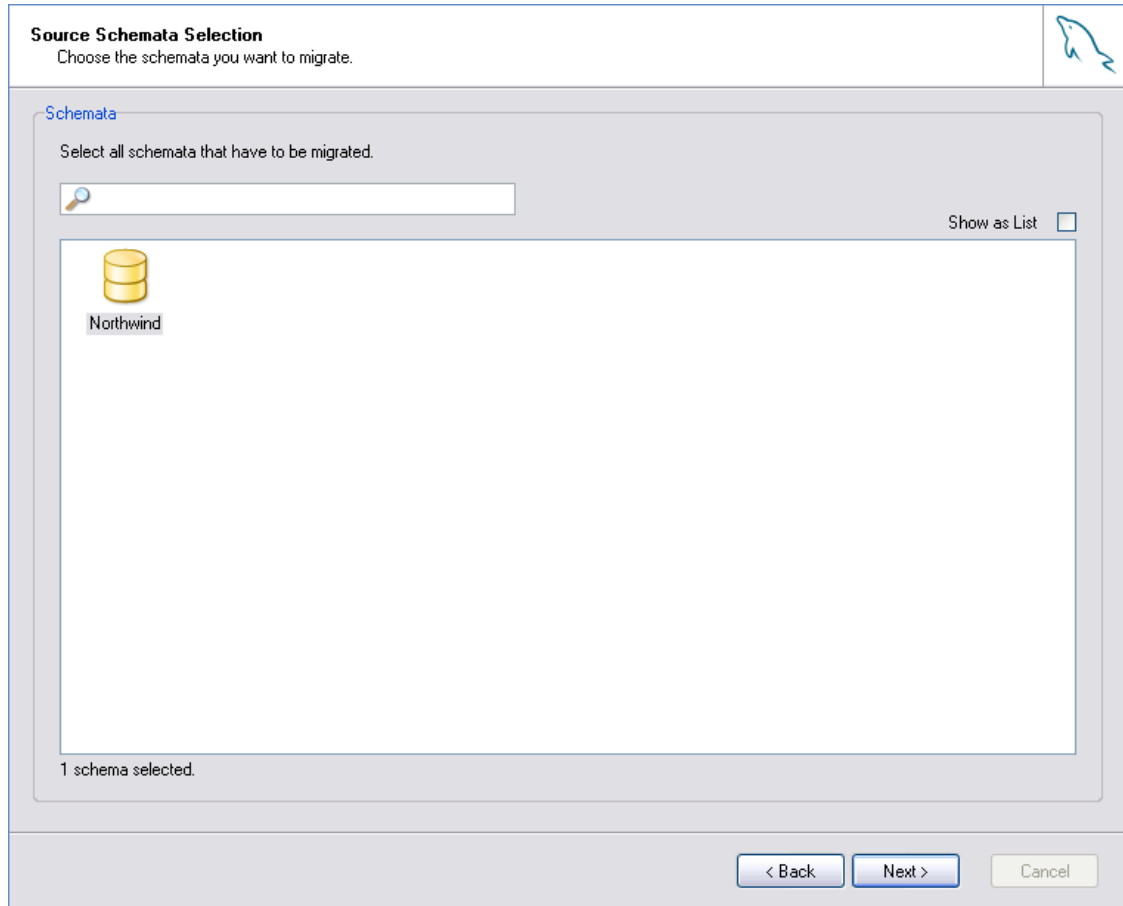
Click the [Details](#) button to see a more detailed log of the connection process.

24.7. The Source Schema Selection Screen

Use the Source Schema Selection screen to choose which databases from the source server you would like to migrate.

If there are a large number of databases to choose from, you can search for a specific database by entering the database name in the [schemata](#) textbox as shown in the following image.

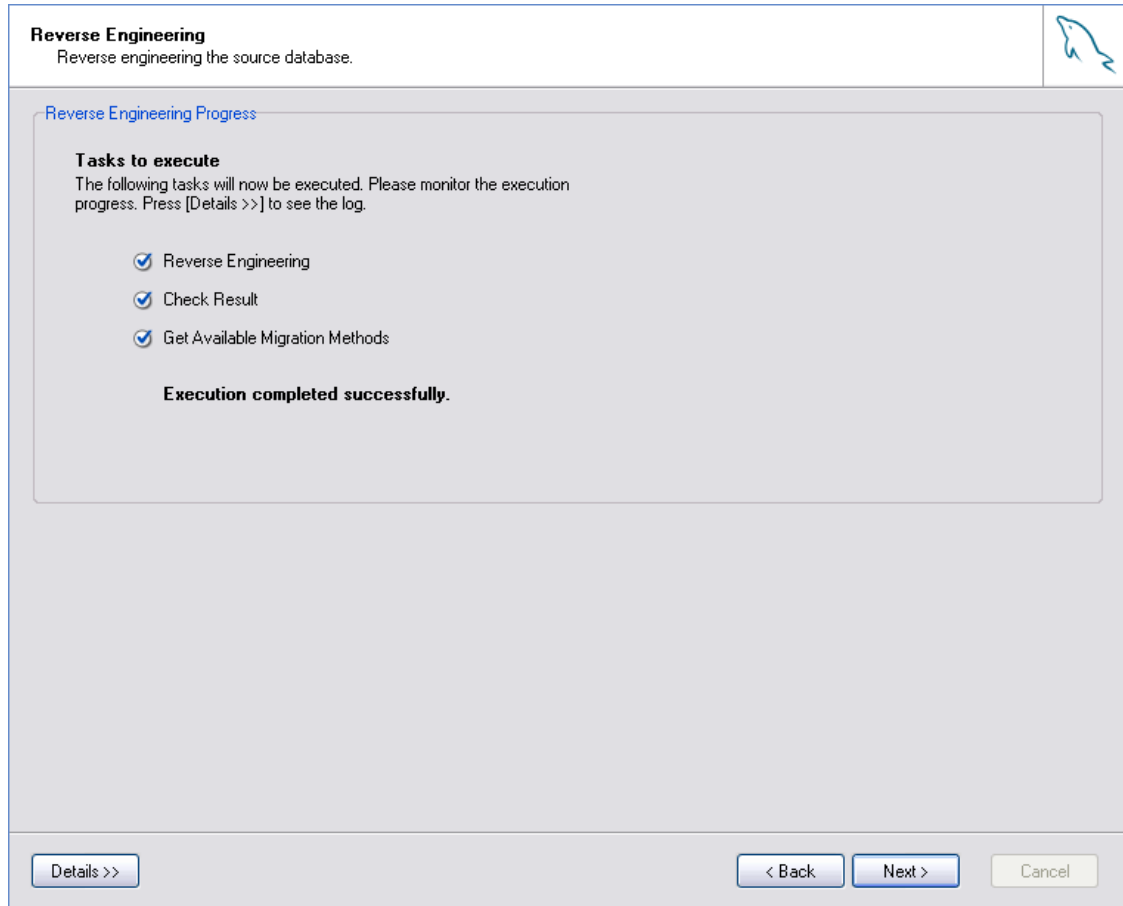
Figure 24.10. The Source Schema Selection screen



24.8. The Reverse Engineering Screen

Once you select the databases you wish to migrate, the MySQL Migration Toolkit begins the process of reverse engineering the source database:

Figure 24.11. The Reverse Engineering screen



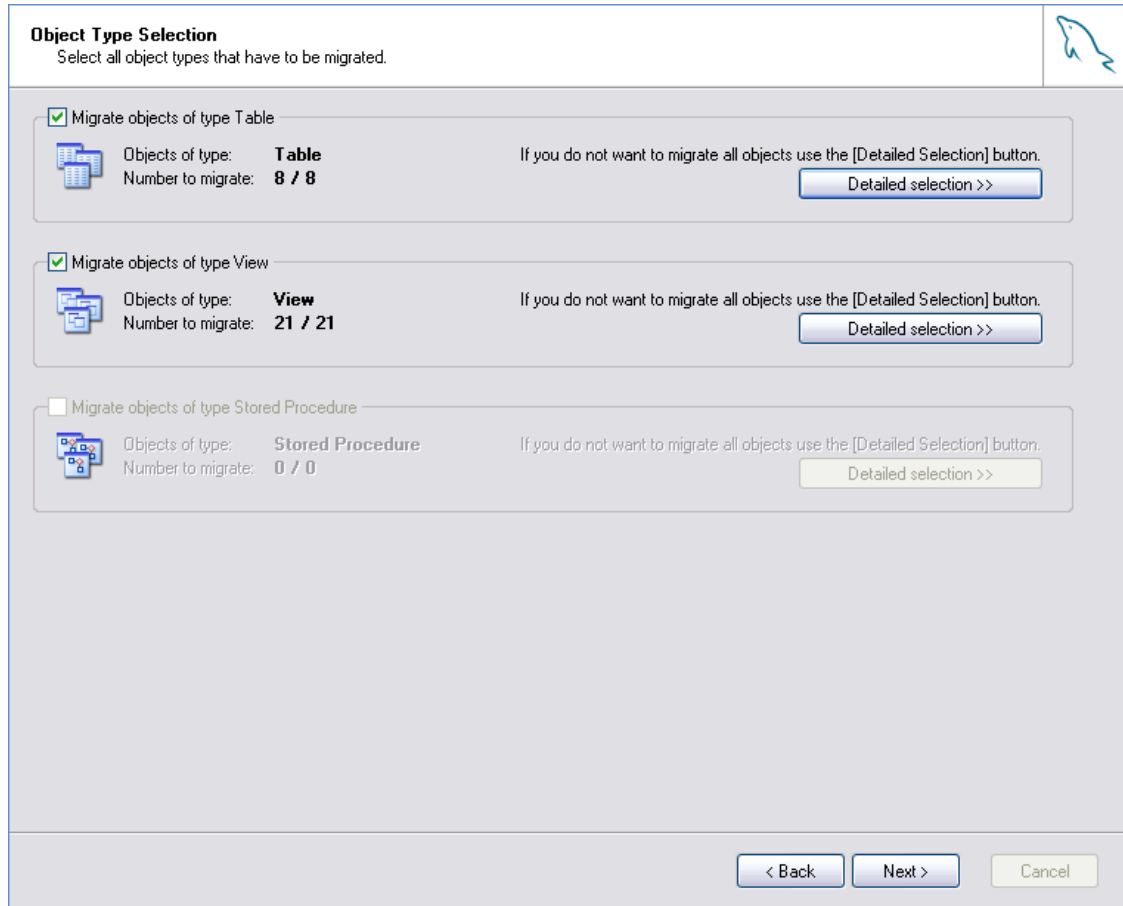
The MySQL Migration Toolkit collects column and index information for each table in the source database, along with information on stored procedures and views.

Click the [Details](#) button to see a detailed log of the reverse engineering process.

24.9. The Object Type Selection Screen

Use the Object Type Selection screen to choose which objects you wish to migrate:

Figure 24.12. The Object Type Selection screen

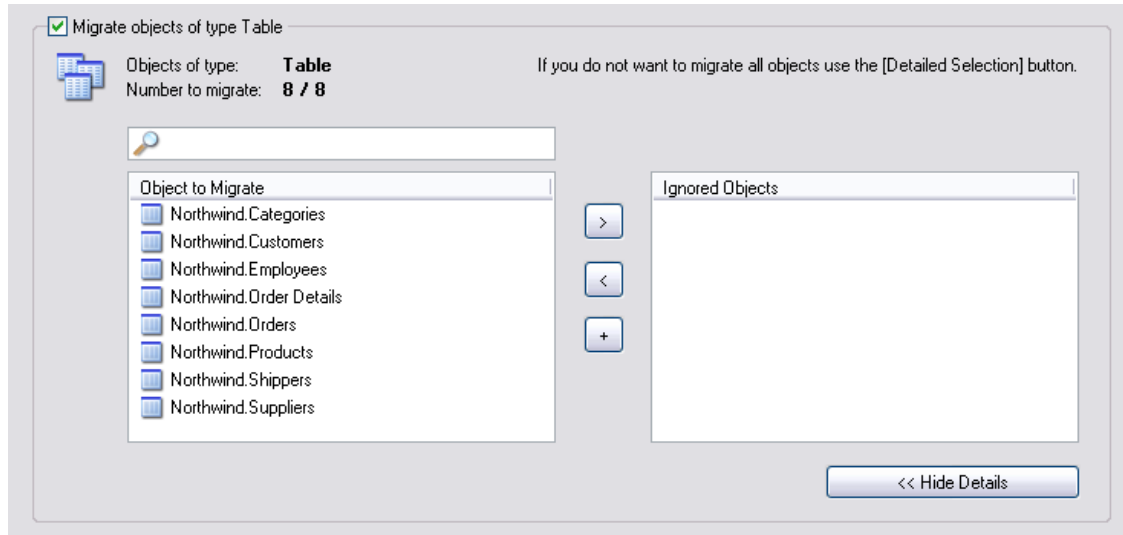


Check the box next to the objects types (Tables, Views, Stored Procedures) that you wish to migrate.

24.9.1. Migrating a Sub-Set of an Object Type

If you only wish to migrate a sub-set of the available object types, click the [Detailed Selection](#) button next to the object type:

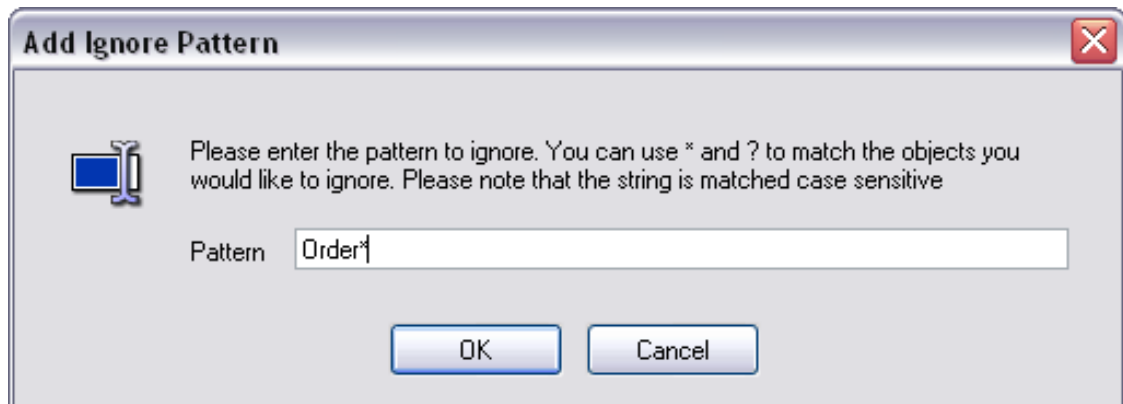
Figure 24.13. The detail view of the Object Type Selection screen



The objects in the left window will be migrated. To ignore an object, select it and click the > button. To move an object out of the ignore list, select it and click the < button.

To exclude objects based on pattern matching, click the + button:

Figure 24.14. The Add Ignore Pattern dialog



Patterns can include the * and ? characters, with * representing multiple characters (including no characters at all) and ? representing a single character.

24.10. The Object Mapping Screen

Use the Object Mapping screen to configure the methods used to migrate the database objects to MySQL. The MySQL Migration Toolkit comes with multiple migration methods that can be used to convert the database objects of an external RDBMS to MySQL.

In most cases the default settings will be adequate.

There are four areas of migration to be addressed: Generic RunTime (GRT) Object, Table, View, and Stored Procedures:

Figure 24.15. The Object Mapping screen

Object Mapping
Please define how to map the database objects.

Migration of type GRT Object
Migration method: **Generic** Generic method to migrate a schema to MySQL. [Set Parameter >>](#)

Migration of type Table
Migration method: **Generic** Generic method to migrate a table to MySQL. [Set Parameter >>](#)

Migration of type View
Migration method: **Access Standard** Standard method to migrate an Access view to MySQL. [Set Parameter >>](#)

Migration of type Stored Procedure
Migration method: No migration method available for this source type.

[Details >>](#) [< Back](#) [Next >](#) [Cancel](#)

24.10.1. GRT Object

The GRT Object section of the Object Mapping screen dictates the properties of the database itself. By default a generic profile is used, with a `Latin1` character set.

To modify the character set used on the database level, click the [Set Parameter](#) button. Choose from `Latin1`, `Multilanguage`, or `User defined`.

24.10.2. Table Objects

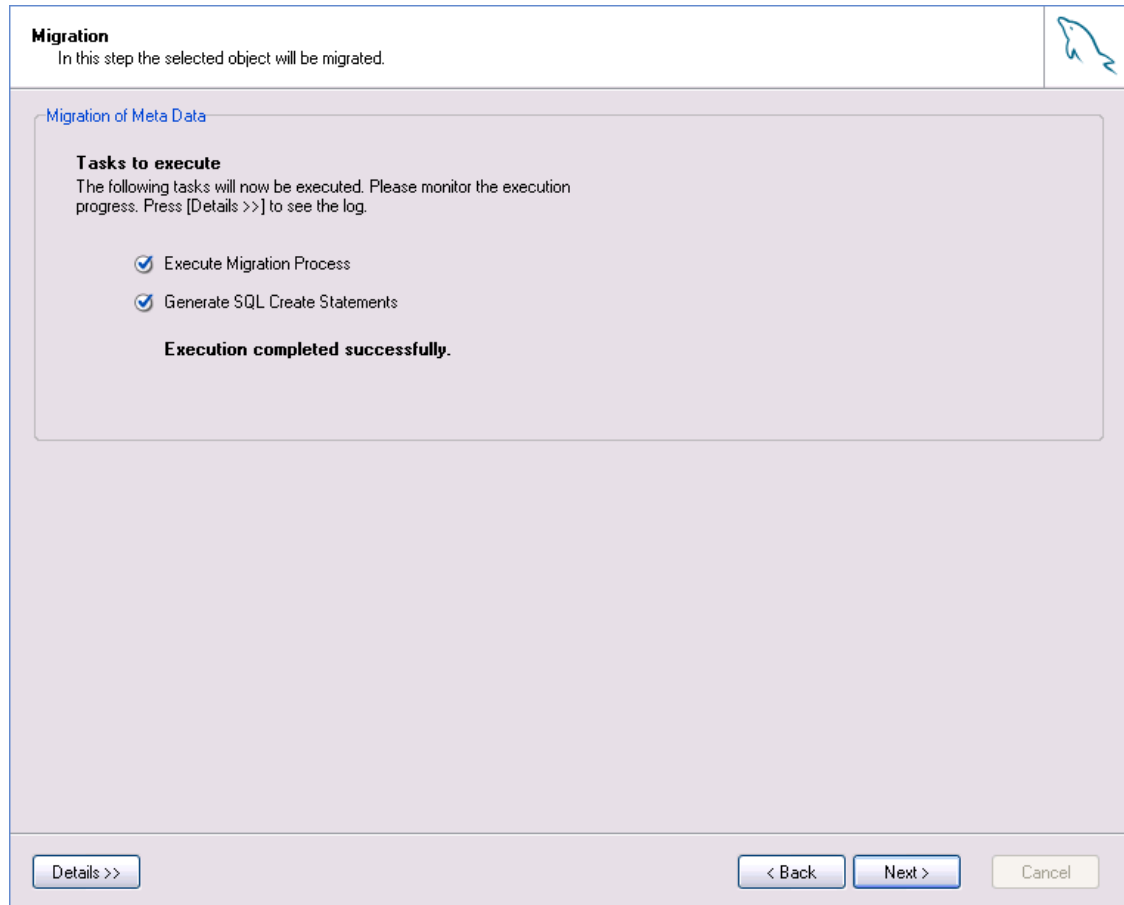
The Table section of the Object Mapping screen dictates the properties of the individual tables. By default a generic profile is used, with an `InnoDB` storage engine.

To modify the storage engine used with the migrated tables, click the [Set Parameter](#) button. Choose the `Data consistency` option to use the InnoDB storage engine for transactional and foreign key support. Choose the `Statistical data` option to use the MyISAM storage engine with increased performance but no transaction safety. Choose the `Data consistency / multilanguage` option to use the InnoDB storage engine with UTF8 as the default charset. If none of the provided options meet your needs, choose the `User defined option and provide your own settings`.

24.11. The Meta Migration Screen

After you configure data object mapping, the MySQL Migration Toolkit performs the conversion of the database objects and generates SQL `CREATE` statements.

Figure 24.16. The Meta Migration screen



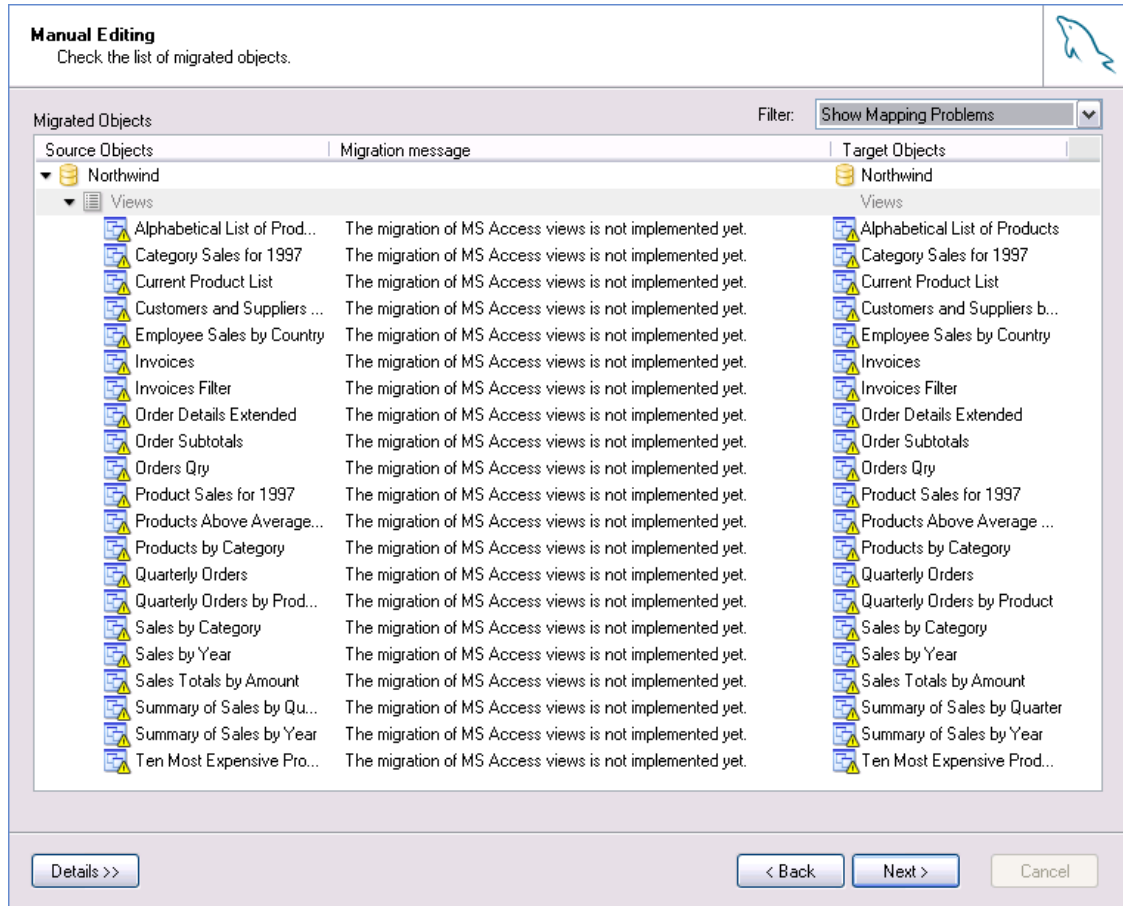
The SQL `CREATE` statements are not executed on the target server at this stage of the migration, but will be executed later.

Click the [Details](#) button to view a detailed log of this stage of the migration process.

24.12. The Manual Editing Screen

Use the Manual Editing screen to review the SQL `CREATE` statements generated by the MySQL Migration Toolkit:

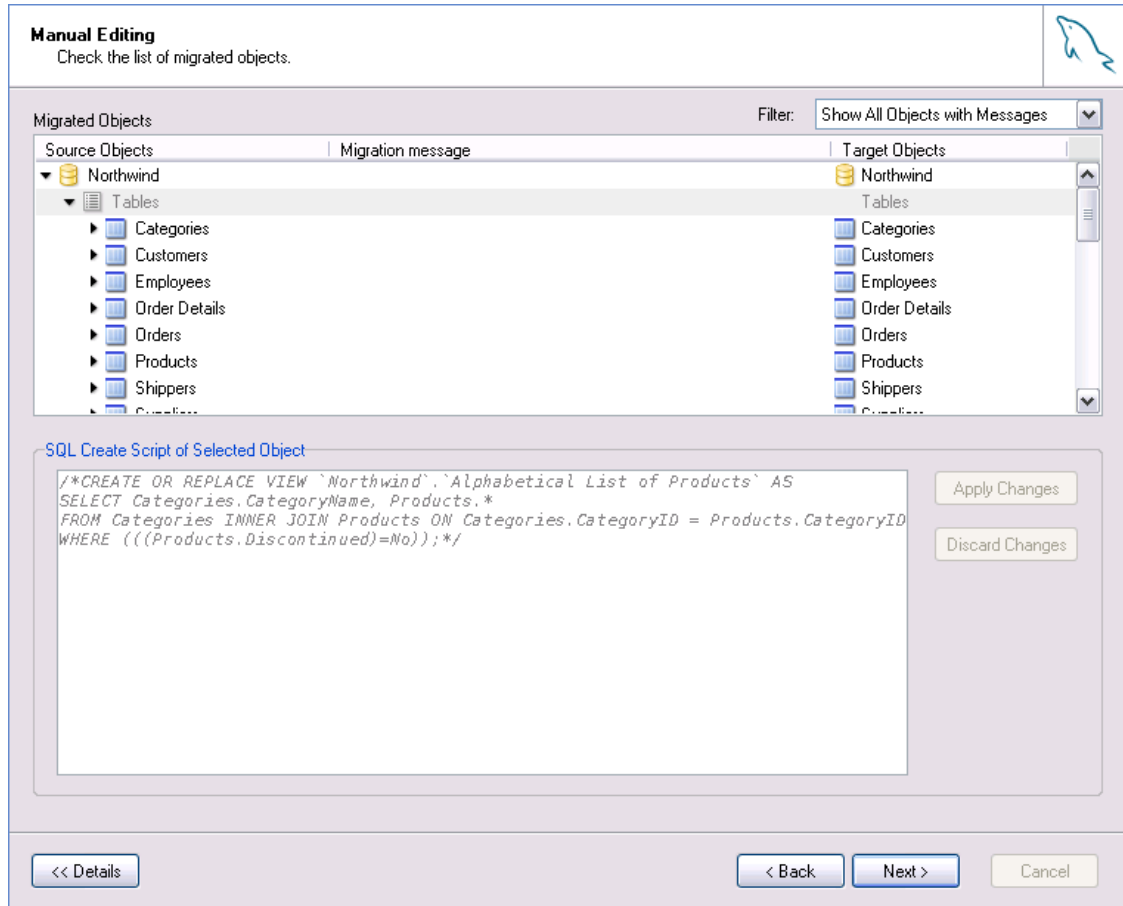
Figure 24.17. The Manual Editing screen



By default only objects that were not successfully converted are displayed in the Manual Editing screen. To view all objects select the Show All Objects entry of the **Filter** drop-down list. Choose the Show All Objects with Messages entry of the **Filter** drop-down list to view all object with status messages.

To edit the SQL **CREATE** statements created by the MySQL Migration Toolkit, select the object and click the **Details** button:

Figure 24.18. The Manual Editing screen – detailed view

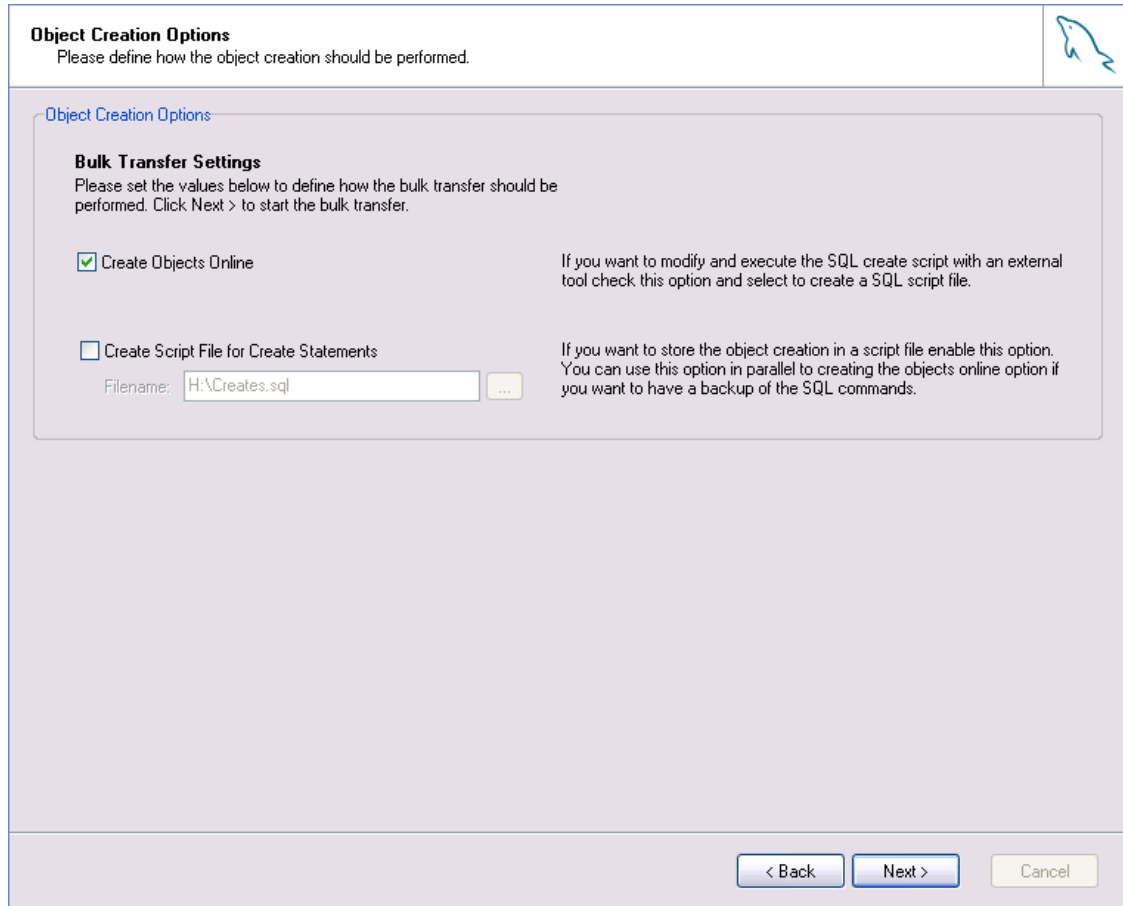


Make changes to the `CREATE` statement and click the `Apply Changes` button. If you make a mistake while editing, click the `Discard Changes` button to undo the `CREATE` statement.

24.13. The Object Creation Options Screen

After performing manual object editing the MySQL Migration Toolkit is ready to create the database objects on the target server. You have the option of either creating the database objects directly on the target server or to create a script file of the `CREATE` statements for later execution:

Figure 24.19. The Object Creation Options screen

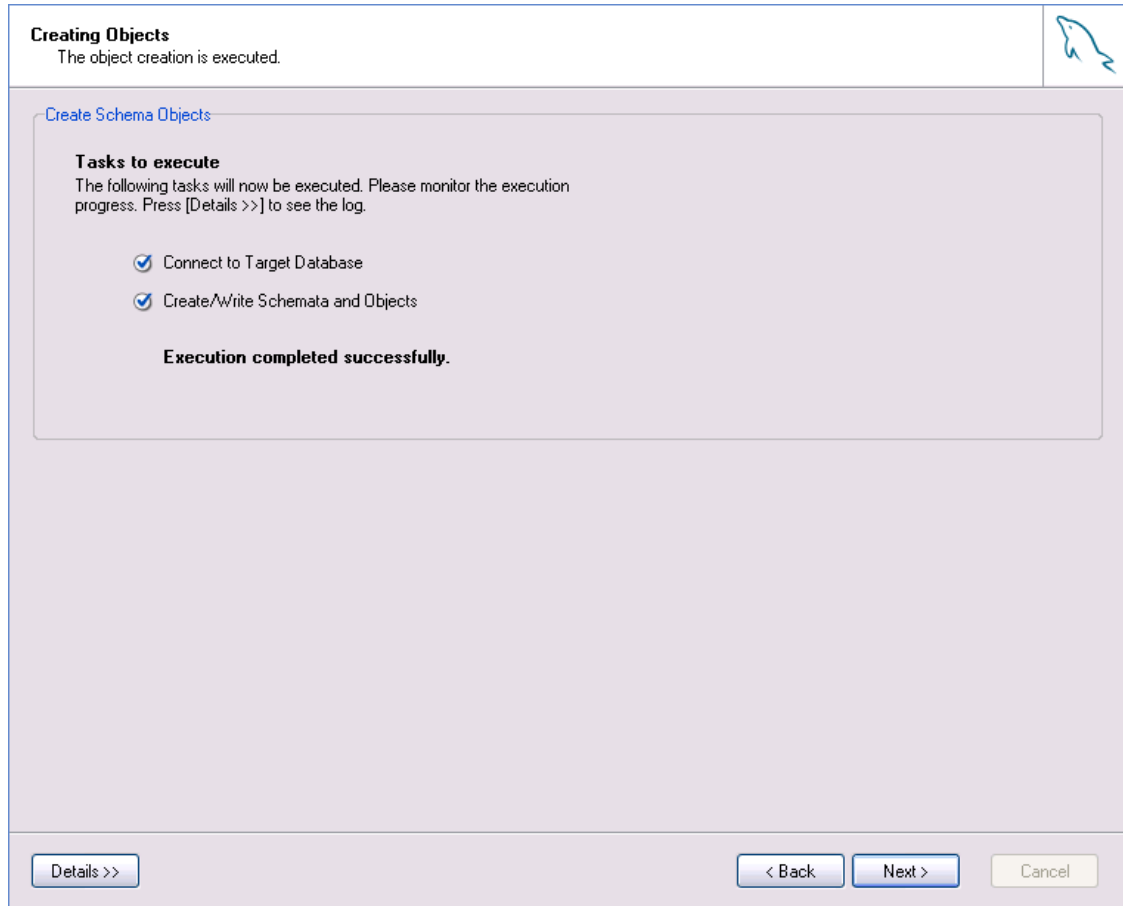


Select both options to create the target database objects and preserve a backup copy of the `CREATE` statements for later use.

24.14. The Creating Objects Screen

Once you choose object creation objects, the MySQL Migration Toolkit connects to the target server and creates the database objects (assuming you chose to have the MySQL Migration Toolkit connect to the target server to create the database objects). If you chose to have a copy of the `CREATE` statements saved to disk this will be performed at this time.

Figure 24.20. The Creating Objects screen

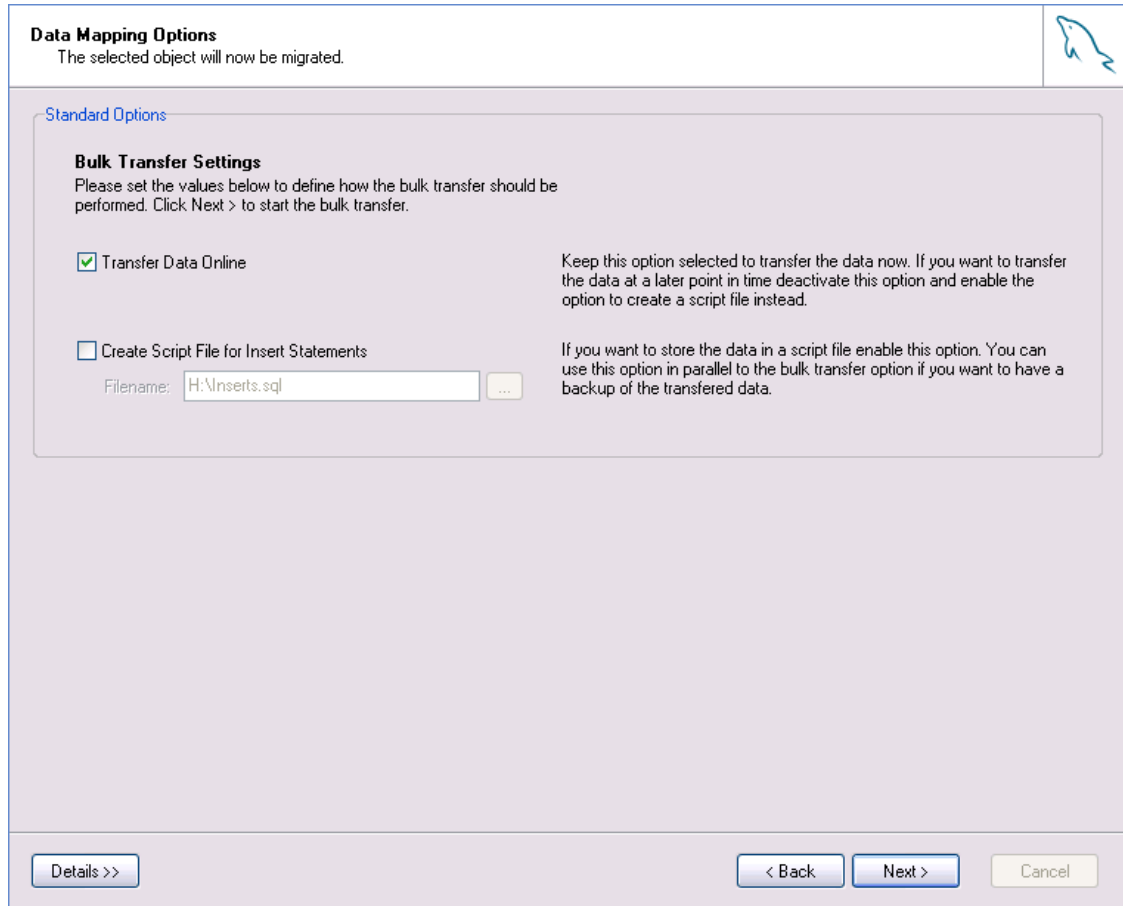


Click the [Details](#) button to view a detailed log of this stage of the migration process.

24.15. The Data Mapping Options Screen

After the database objects are created on the target database server the MySQL Migration Toolkit is ready to move the server data to the target server. You have the option of either inserting the data directly on the target server or to create a script file of the `INSERT` statements for later execution:

Figure 24.21. The Data Mapping Options screen



Select both options to move the data and preserve a backup copy of the `INSERT` statements for later use.

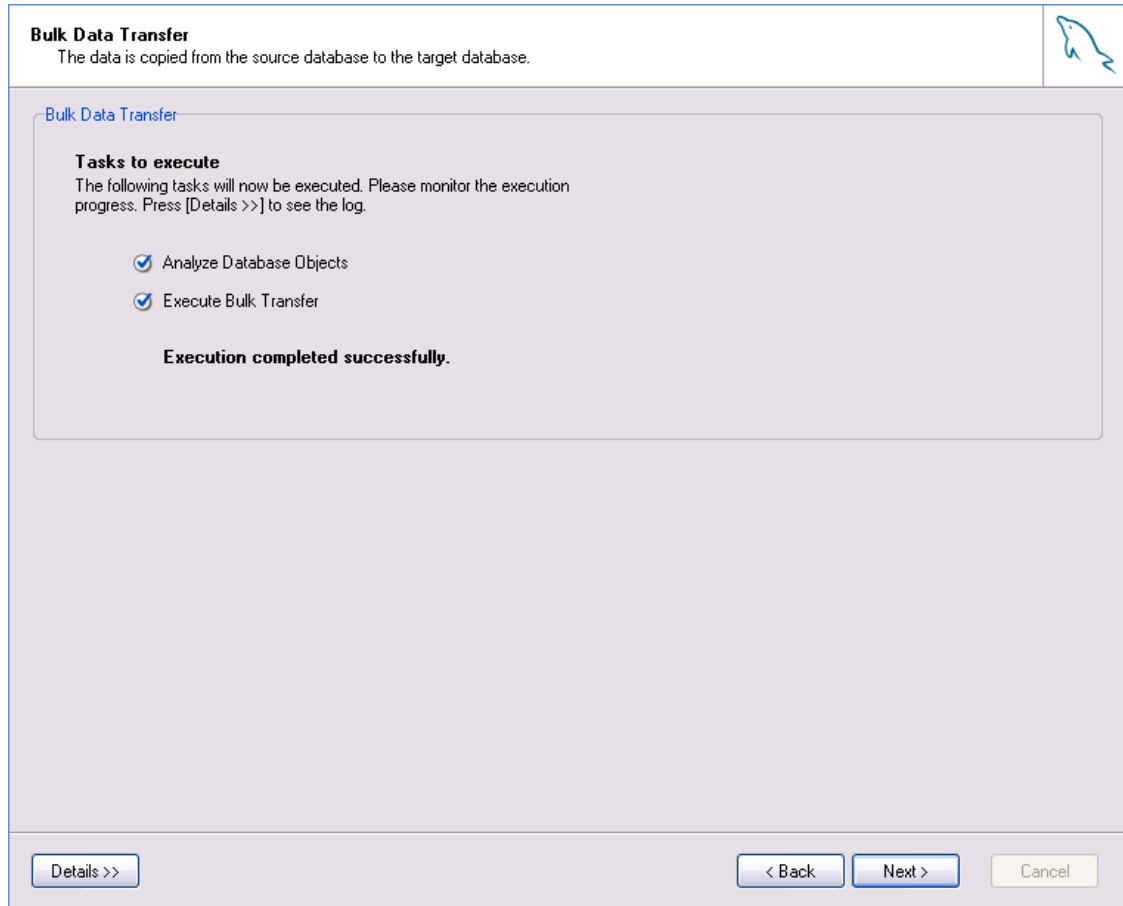
You can access additional options by clicking the [Details](#) button. If you do not wish to move BLOB data to the target server, check the box next to the [Exclude BLOB values](#) option. The BLOB data will not be moved to the target server and will not be written to the script file. If you do not wish to move BLOB data to the target server, but wish to have the BLOB data written to the script file, also check the box next to the [Write BLOBs to Insert Script](#) option.

If you wish to limit the number of rows transferred to the target server, check the box next to the [Maximum Numbers of Rows to Transfer for a Table](#) option and enter the desired number of rows. This option can be useful when generating test data.

24.16. The Bulk Data Transfer Screen

Once the data mapping options are set, the MySQL Migration Toolkit will begin the bulk data transfer process. Data will be converted to a MySQL compatible format and inserted into the target database server using bulk `INSERT` statements. Data is typically inserted in batches of 15,000 rows at a time to maximize insertion speed.

Figure 24.22. The Bulk Data Transfer screen

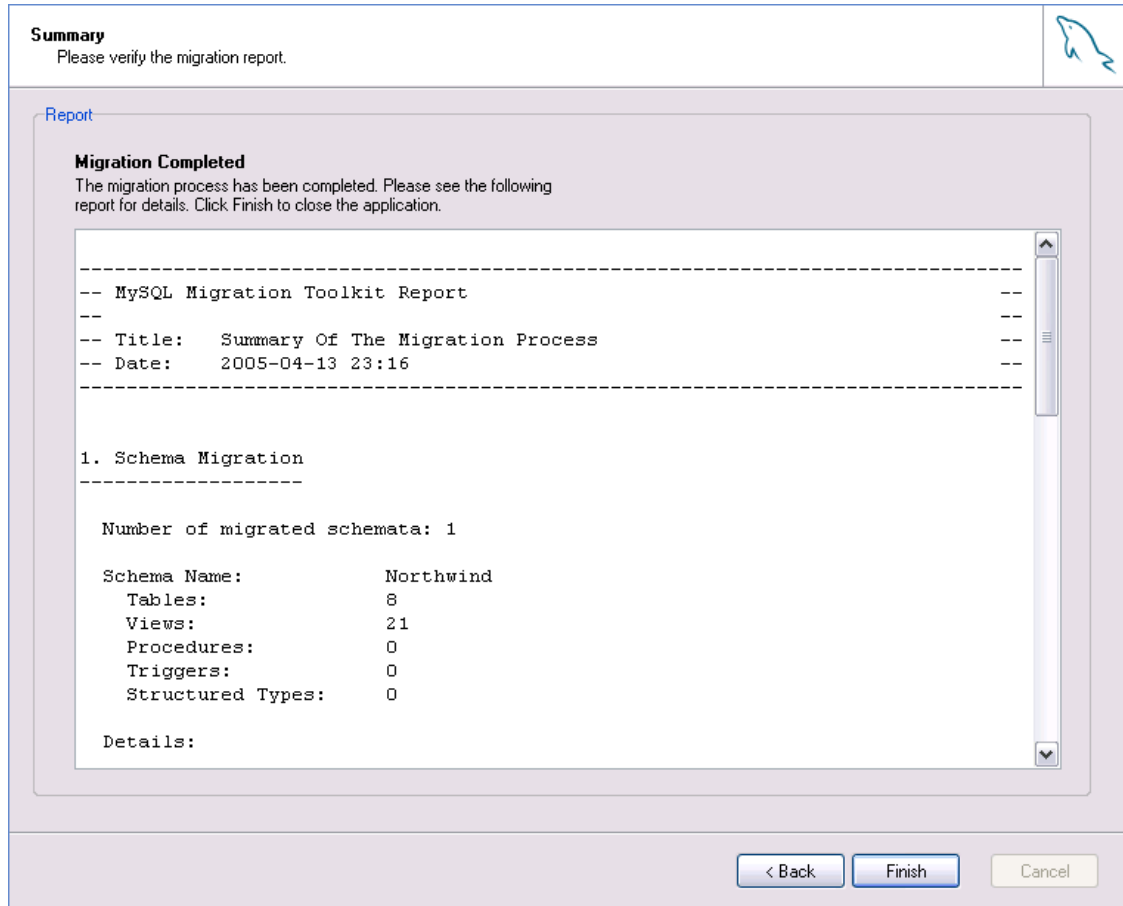


Click the [Details](#) button to view a detailed log of this stage of the migration process.

24.17. The Summary Screen

Once the bulk data transfer is complete, a summary report of the migration will be displayed:

Figure 24.23. The Summary screen



Click the [Finish](#) button to exit the MySQL Migration Toolkit.

24.18. Saving the Current Application State

If you need to close the MySQL Migration Toolkit part way through the migration process, you can save the state of the MySQL Migration Toolkit to disk and continue later.

To save the current state of the MySQL Migration Toolkit, choose the Store current application state option of the File menu.

To later retrieve the application state and resume configuring the migration, choose the Reload stored application state option of the File menu.

Chapter 25. Extending The MySQL Migration Toolkit

25.1. Introduction

One of the key features of the MySQL Migration Toolkit is that it can be easily extended and customized to support new data sources. Through the use of its Generic RunTime (GRT) Environment, the MySQL Migration Toolkit can support new behavior and data sources using code written in languages such as Java, PHP, and C.

Extending the MySQL Migration Toolkit to support a new RDBMS requires the creation of two new modules: one module that reverse engineers (ie., retrieves schema information) the source database objects and converts them into GRT objects, and one module that migrates the source GRT objects into MySQL GRT objects. The MySQL Migration Toolkit then converts the GRT MySQL objects into SQL statements that create the the target MySQL server objects.

Existing modules can be easily expanded and customized to achieve the perfect migration for individual requirements. New migration “methods” that define how the migration is performed can be added easily. The new methods are listed and can be selected from the Wizard interface or used in migration scripts.

25.2. Architecture of the MySQL Migration Toolkit

The GRT is a thin C layer, inspired by Objective C, which allows for dynamic typing and dynamic data objects. The MySQL Migration Toolkit is written on top of this layer and supports the creation of new modules written in C, Java, PHP, and Lua, with support for Python and Mono forthcoming.

The MySQL Migration Toolkit is built on three primary modules: reverse engineer modules, migration modules, and transformation modules.

Reverse engineering modules retrieve the schema information from the source database and return GRT objects that describe the schema. Reverse engineering modules will have a name similar to `ReverseEngineeringAccess`.

Migration modules convert the source database GRT objects to MySQL GRT objects and then handle the bulk data transfer between the source and MySQL databases. Migration modules will have a name similar to `MigrationAccess`.

Transformation modules convert the MySQL GRT objects into the actual SQL statements used to create objects such as tables and views on the target MySQL server. Transformation modules will have a name similar to `TransformationMySQL`. Transformation modules are supplied by the MySQL GUI team and need not be created to add support for a new source database.

All modules are derived from base classes whose methods can be re-written to match the new source database.

25.3. The Modular Migration Process

From the point of view of modular development, the migration process is as follows:

1. The MySQL Migration Toolkit prompts the user for a source database and connection parameters. The user selection determines which modules will be used for the source database.
2. The MySQL Migration Toolkit calls the `testConnection` method of the reverse engineering

module. The `testConnection` method returns success or failure to the MySQL Migration Toolkit.

3. The MySQL Migration Toolkit calls the `getSchemata` method of the reverse engineering module. The `getSchemata` method returns a list of the schemas available on the source RDBMS for the user to choose from.
4. Once the user has selected one or more schema to migrate, the MySQL Migration Toolkit calls the `reverseEngineer` method of the reverse engineering module. The `reverseEngineer` method converts all objects in the source RDBMS (tables, views, procedures) into a collection of GRT objects.
5. After the source database has been reverse engineered, the MySQL Migration Toolkit calls the `migrationMethods` method of the migration module. The `migrationMethods` method generates a list of available methods than can be selected by the user in the [Object Mapping Screen](#).
6. Once the user has selected the migration methods, the MySQL Migration Toolkit calls the `migrate` method of the migration module. The `migrate` method converts the source RDBMS GRT objects into MySQL GRT objects by calling the selected migration method for each source schema object. If no explicit method is assigned, the default migration method for the object type will be used.
7. After the MySQL GRT objects have been created, the MySQL Migration Toolkit calls the transformation module to convert the MySQL GRT objects into SQL statements that will create the objects on the target MySQL server. The MySQL Migration Toolkit then either executes the SQL statements on the target MySQL server or writes them to a script file.
8. In the final step, the MySQL Migration Toolkit calls the `dataBulkTransfer` method of the migration module. The `dataBulkTransfer` method loops through the selected tables in the selected schema and migrates the rows of the tables to the target MySQL database or loads them into a script file, depending on the user preference.

25.4. Tools Required to Extend the MySQL Migration Toolkit


The MySQL AB GUI Team uses Eclipse for development of the MySQL Migration Toolkit modules and recommends Elcipse for use when developing modules for the MySQL Migration Toolkit. See <http://www.eclipse.org/> for more information.

Chapter 26. Preparing a Microsoft Access Database for Migration

The MySQL Migration Toolkit requires access to the system tables of a Microsoft Access database for the purpose of reverse-engineering. By default, read access to the system tables of an Access database is restricted for external applications.

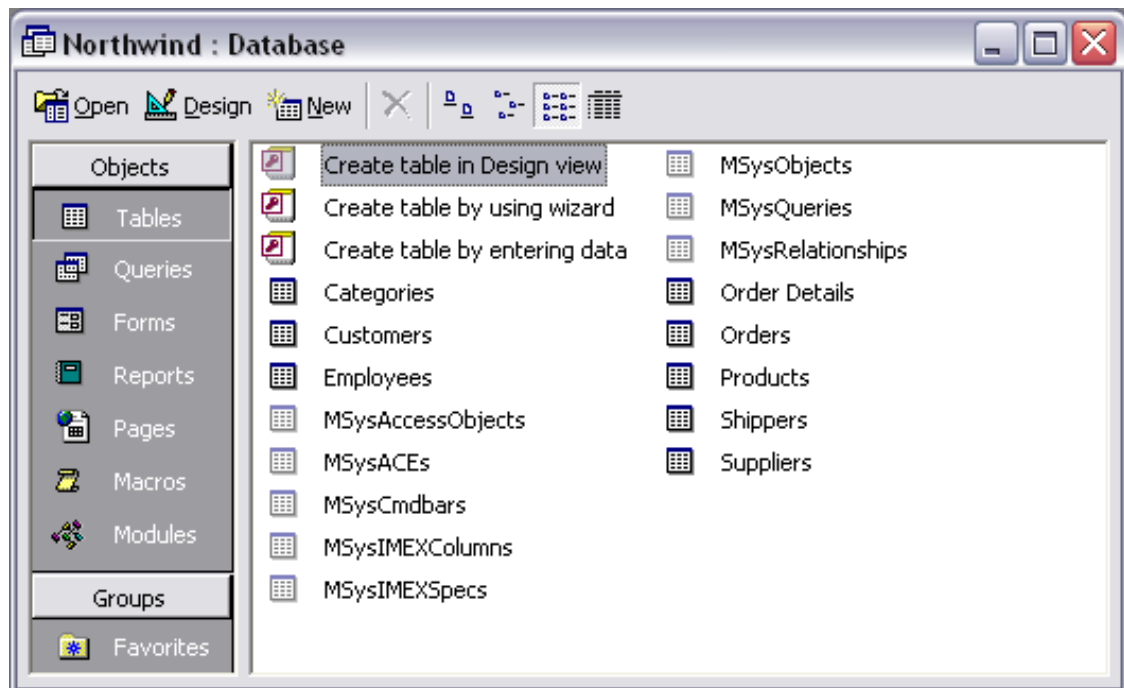
To enable access to the system tables by the MySQL Migration Toolkit, open the database in Microsoft Access and choose the Options entry of the Tools menu. Within the `Options` dialog, select the `View` tab and look for the `Show` section:

Figure 26.1. The show section



Check the box next to the `System objects` option and close the options dialog. System tables for the database should now be present:

Figure 26.2. The system objects

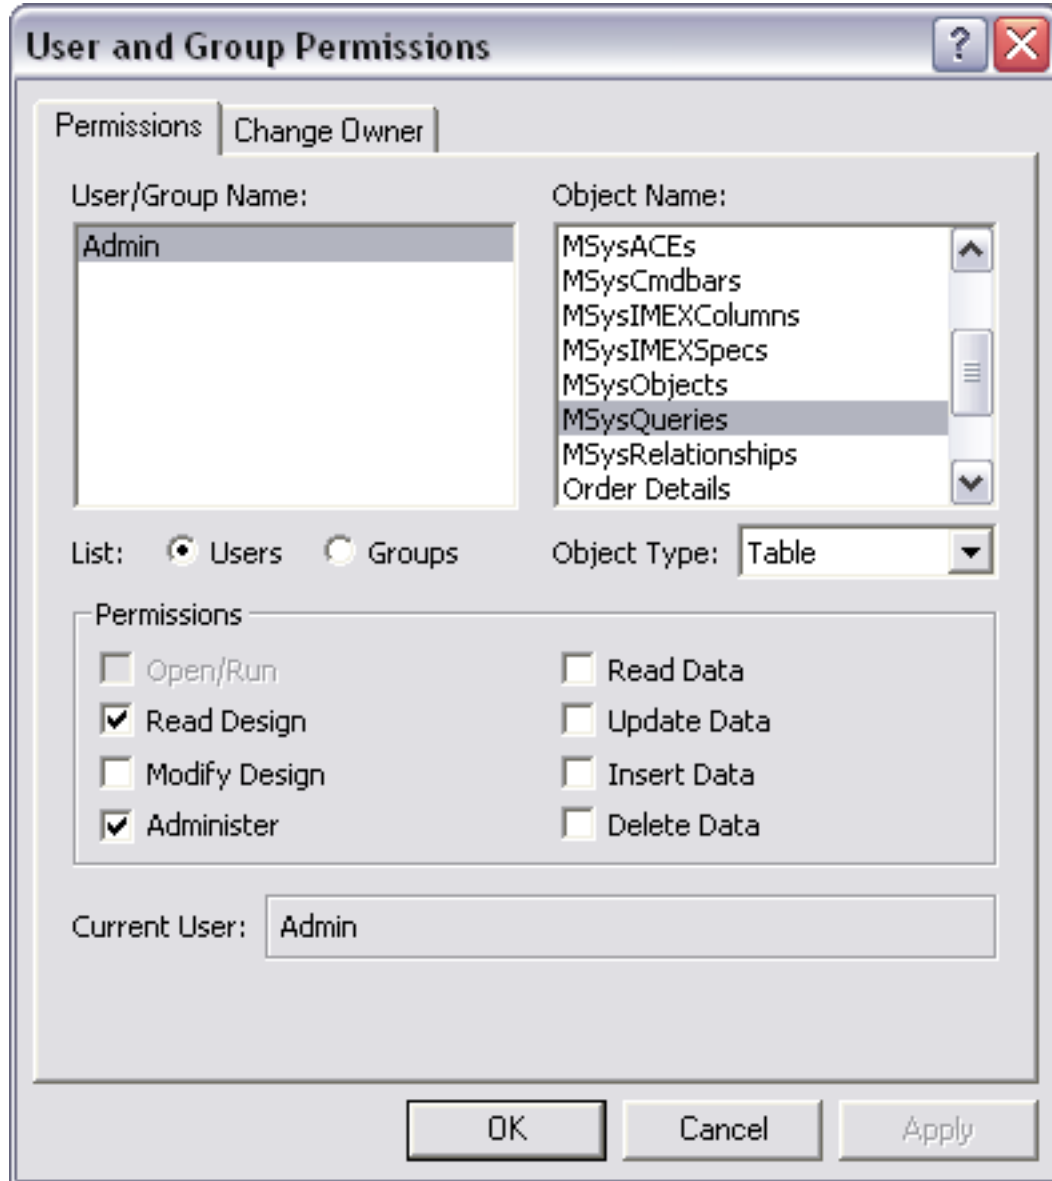


Note the presence of the various `MSys` tables.

After you expose the system objects to the MySQL Migration Toolkit, you must also grant permission to

access the objects. Choose the User and Group Permissions entry from the Security section of the Tools menu:

Figure 26.3. Granting access to the system objects



Enable the [Administer](#) permission for both the [MSysObjects](#), [MSysQueries](#), and [MSysRelationships](#) tables. You will need to click the [Apply](#) button after enabling each table and before moving to the next table.

After completing these steps, you can reverse engineer the schema in your Access database using the MySQL Migration Toolkit.

Part V. MySQL Workbench

Table of Contents

27. Introduction to MySQL Workbench	114
28. MySQL Workbench Overview	115
28.1. The Vertical Toolbar	117
28.1.1. The Standard Mouse Pointer	117
28.1.2. The Hand Tool	117
28.1.3. The Eraser Tool	117
28.1.4. Using the Layer Tool	117
29. Using MySQL Workbench	119
29.1. Creating a New Schema	119
29.2. Creating Tables	119
29.2.1. Adding Columns	120
29.2.2. Adding Indexes	121
29.2.3. Adding Foreign Keys	122
29.2.4. Adding Triggers	122
29.2.5. Adding Table Options	122
29.2.6. Standard Inserts	122
29.3. Creating Views	122
29.4. Creating Routines	123
29.5. Creating Routine Groups	123
29.6. Creating Synonyms	123
29.7. Creating Structured Data Types	123
29.8. Creating Sequences	123
29.9. Saving a MySQL Workbench File	123
30. MySQL Workbench Reverse Engineering	124

Chapter 27. Introduction to MySQL Workbench

MySQL Workbench is a schema design application that enables users to graphically design and modify tables and other database objects, then automatically implement their design in MySQL.

MySQL Workbench is designed to work with MySQL versions 4.0 and higher.

MySQL Workbench is beta software and should not be used on production systems. It is strongly recommended you backup all data before using MySQL Workbench. If you find it's lacking some feature important to you, or if you discover a bug, please use our [MySQL Bug System](http://bugs.mysql.com) [http://bugs.mysql.com] to request features or report problems.

Chapter 28. MySQL Workbench Overview

All functionality of MySQL Workbench is accessed from the design screen:

Figure 28.1. Workbench design screen – Windows

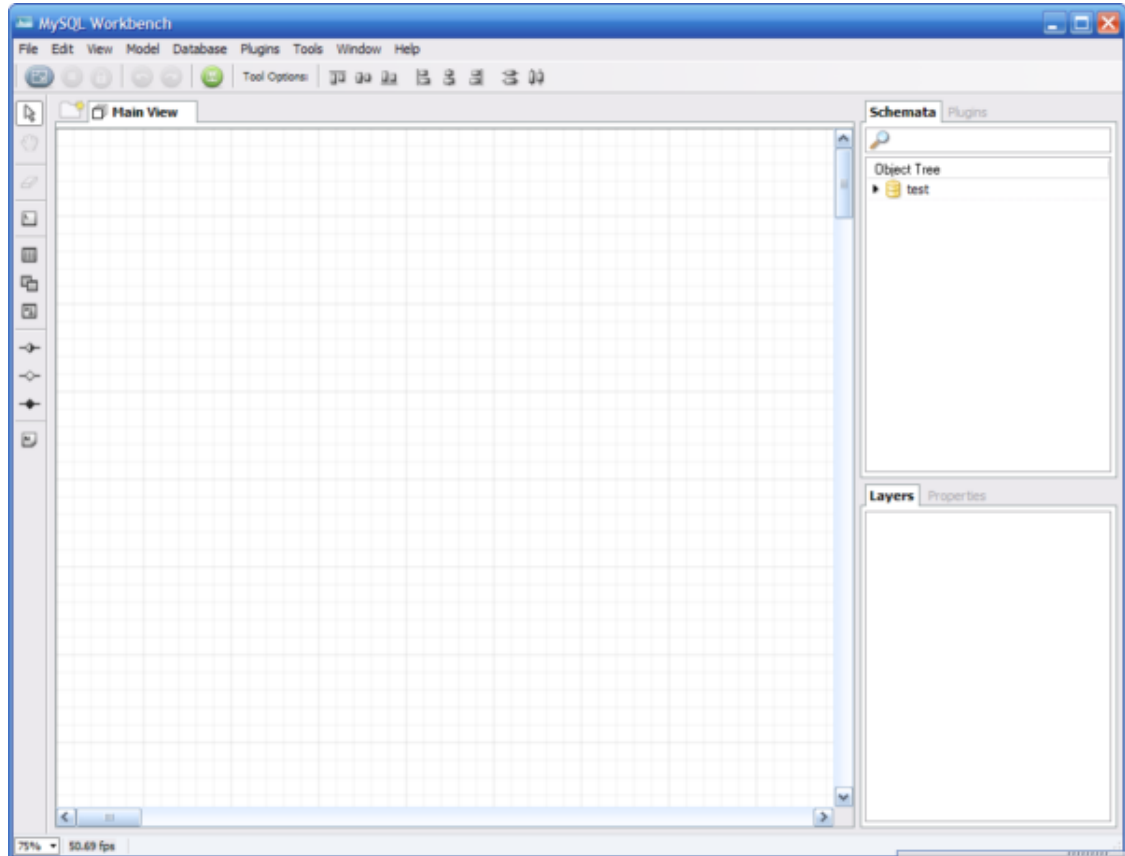


Figure 28.2. Workbench design screen – OS X

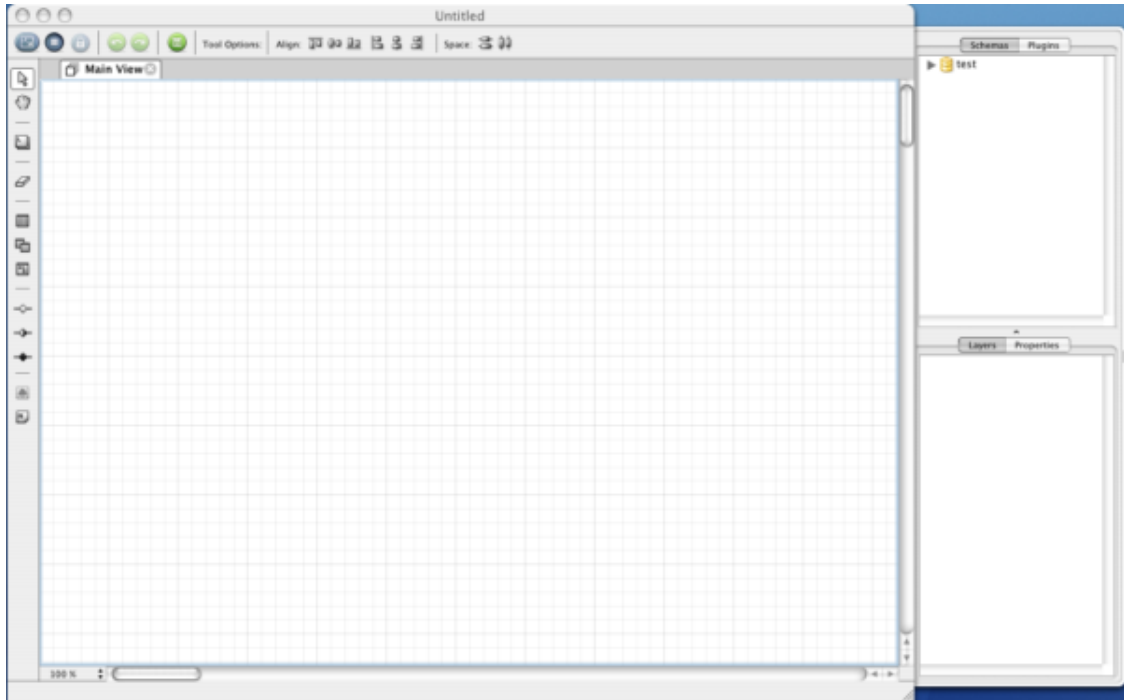
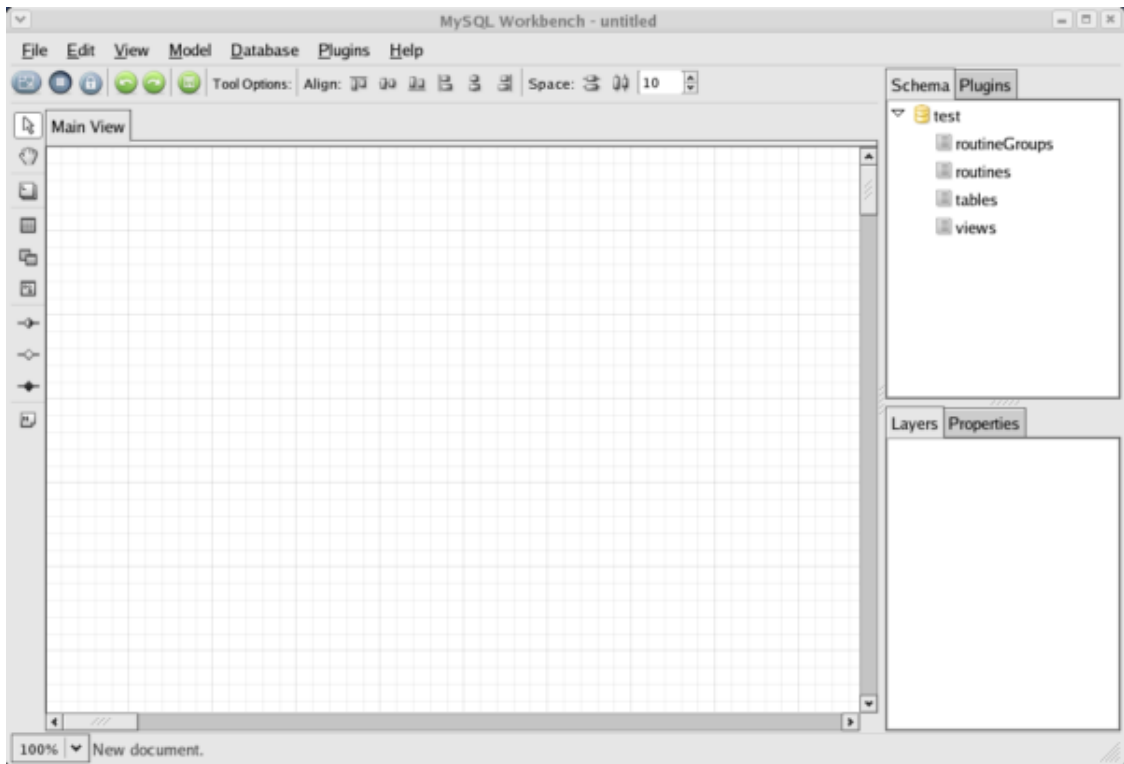


Figure 28.3. Workbench design screen – Linux



Most design operations are performed using the vertical and horizontal toolbars that border the center Model View area.

28.1. The Vertical Toolbar

The vertical toolbar contains the following icons:

- Standard Mouse Pointer
- Hand
- Eraser
- Layer Tool
- Table Tool
- View Tool
- Routine Group Tool
- The 1:n Tool
- The 1:1 Tool
- The n:m Tool
- The Note Tool

These tools are used to manipulate or create objects shown in the [Main View](#). To indicate that a tool is active, clicking a tool changes the mouse pointer.

28.1.1. The Standard Mouse Pointer

As its icon would indicate, the standard mouse pointer is the default mouse pointer for your operating system.

Use this tool to revert to the standard mouse pointer after using other tools.

28.1.2. The Hand Tool

This tool is not yet activated.

28.1.3. The Eraser Tool

This tool is not yet activated.

28.1.4. Using the Layer Tool

The layer tool is the rectangular icon with a capital **L** in the top left corner.

The layer tool is used to organize the objects presented in the [Main View](#). It is useful for grouping together similar objects. You may, for instance, want to group all your views together.

Click the layer tool to use it. Draw a rectangle on the [Main View](#). Change to the standard mouse pointer tool and pick up any objects you would like to place in the newly created layer.

To change the size of a layer, mouse over the lower left corner of a layer. When the mouse changes to a

double-headed arrow, hold down the left button and adjust the size of the layer.

To change the name of a layer, click the [Properties](#) tab beside the [Layers](#) tab on the lower right side of the screen. You may also adjust its position and size.

Chapter 29. Using MySQL Workbench

When MySQL Workbench is first opened, an empty [Main View](#) tab occupies most of the screen and [Schemata](#) and [Layers](#) tabs appear on the right.

If you are not opening an existing file, creating a schema is the natural place to begin.

29.1. Creating a New Schema

You may begin by using the [test](#) schema shown in the [schemata](#) panel on the right.

However, to create a new schema right-click in the [Schemata](#) tab. Choose the [Create MySQL Schema](#) option. This will create a schema called [New Schema](#) in the [Object Tree](#).

Click on the [arrow](#) icon to the right of your new schema in order to see all the elements in the object tree. Those elements are:

- routineGroups
- routines
- sequences
- structuredTypes
- synonyms
- tables
- views

Note

MySQL Workbench is alpha software. Not all features are implemented and some are only implemented on specific platforms.

29.2. Creating Tables

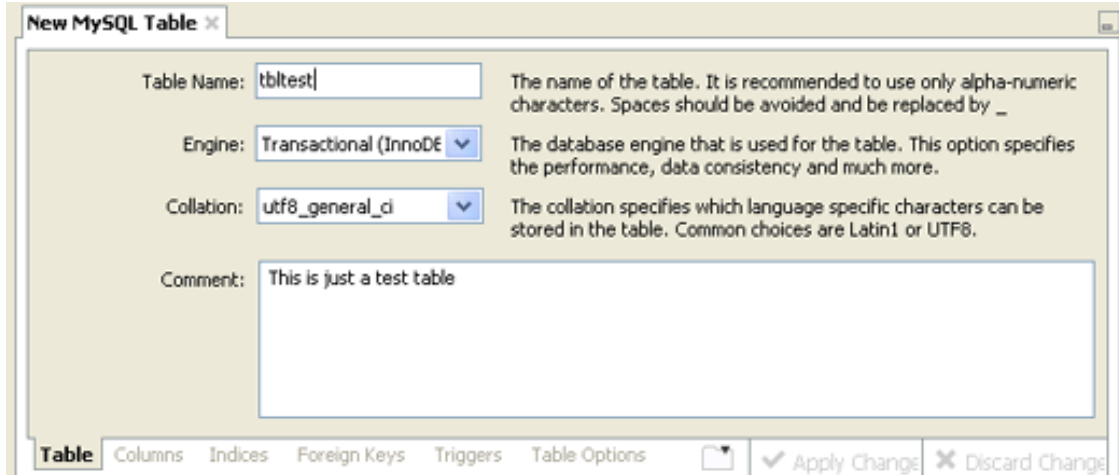
The first thing you will want to do is create tables for your schema.

When creating or editing a table you can choose to dock the table editor or not. To toggle the table editor view, select the View, Advanced, Open Editors Dock option. The description of the table editor given here applies to the docked view.

To create a table, right click the schema in the [Object Tree](#) and choose the [Create MySQL Table](#) option.

This will open a panel at the bottom of the screen, similar to the figure shown below:

Figure 29.1. Creating a table



This panel has seven tabs along the bottom. The first tab, the table information view, is the default view.

In this view you may set the table name, engine type, collation, and also add a comment.

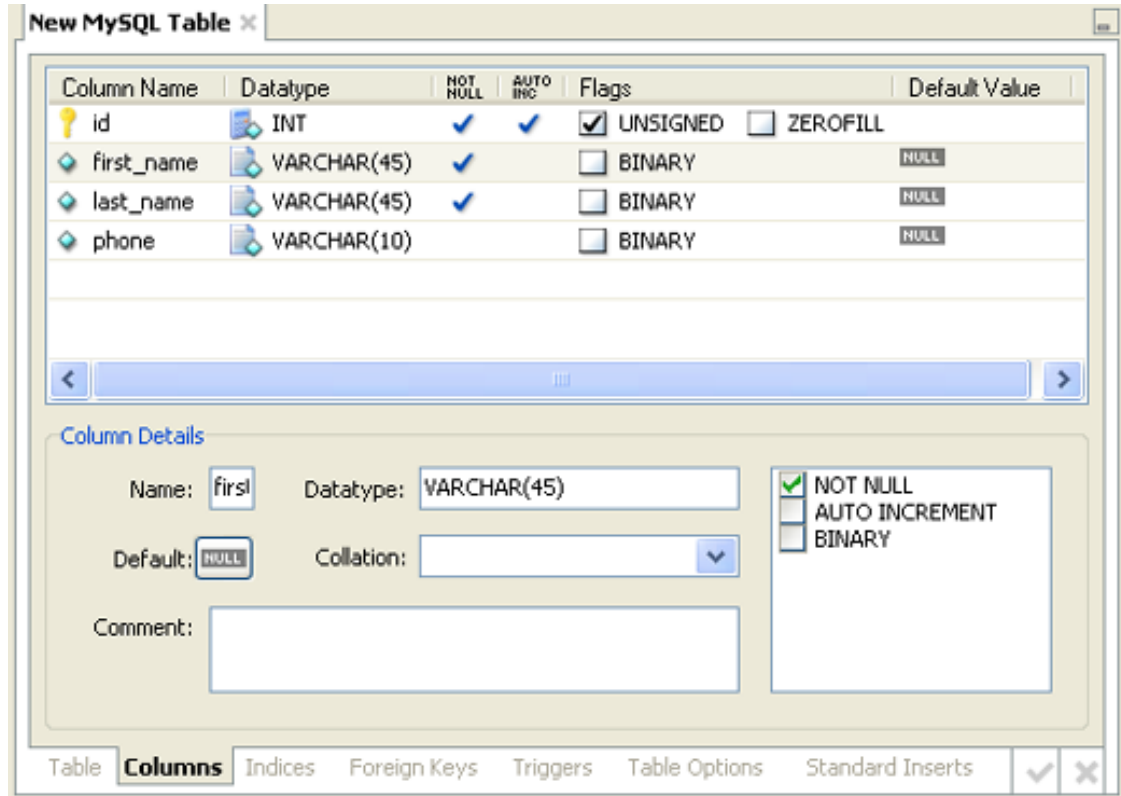
To apply your changes, click the [Apply Changes](#) button shown on the lower left in [Figure 29.1, “Creating a table”](#). This button will only be enabled if there are changes that have not been applied. To back out of any changes, choose the [Discard Changes](#) button. Again, this button will only be enabled if there are changes to discard.

29.2.1. Adding Columns

Add columns to your table by choosing the `columns` tab.

You may choose the name, data type, default value, and collation for each column. New columns may be added using the column details frame in the lower part of the screen or by using the tabular layout—both are shown in [Figure 29.2, “Creating columns”](#).

Figure 29.2. Creating columns



To edit a column select the column you wish to edit and choose the property you wish to change.

Depending upon circumstances one or the other of these views may be more convenient.

However, changing the data type in the tabular view is especially easy because a drop-down list shows all available data types.

29.2.2. Adding Indexes

Once you have created a table, it will appear in the [Main View](#) panel and its properties will be displayed in a separate window below the [Main View](#) panel. You may close this panel by clicking the **x** beside the table name.

To add an index, reopen the properties window by right clicking the table and choosing the [Edit](#) option. This will re-open the table edit window below the [Main View](#) tab. Choose the [indices](#) tab.

This will open a horizontally split screen, showing any existing indexes in the top half and the index properties in the lower half.

To create a new index, find an empty row in the upper half of the table edit window and double click the [Index Name](#) column. Choose an appropriate name for the index.

Double click the [Type](#) column and choose an index type from the drop down box. You may choose from:

- INDEX
- PRIMARY

- UNIQUE
- FULL TEXT
- SPACIAL

Once you have named the index and chosen its type, you may set the properties in the lower half of the split screen.

29.2.3. Adding Foreign Keys

To add a foreign key, choose the [Foreign Keys](#) tab.

You create a foreign key in much the same way that you create an index. Assign a name to the foreign key, identify the referred table, and then choose the action to perform on update and on delete.

The options are:

- NO ACTION
- RESTRICT
- CASCADE
- SET NULL

Choose the appropriate action for each event. For instance, on delete, you may wish to perform a cascading delete of related records.

29.2.4. Adding Triggers

A trigger is a database object that is associated with a table and is activated when a particular event occurs for the table.

Use the [Trigger](#) tab to add a trigger to a table.

Use the textbox provided to enter the SQL statement for creating a trigger.

Be sure to click [Apply Changes](#) to save your work.

29.2.5. Adding Table Options

Clicking the [Table Options](#) tab

Use this option to pack keys, add a table password or set the initial AUTO_INCREMENT value.

Row options, storage options and merge options can also be set from this screen.

29.2.6. Standard Inserts

Use the [Standard Inserts](#) tab to ...

29.3. Creating Views

To create a view right click the schema in the [Object Tree](#) and choose the Create MySQL View option.

Enter the view name and the SQL statement to create the view.

Choose [Apply Changes](#) to save the view.

29.4. Creating Routines

This feature is not yet implemented.

29.5. Creating Routine Groups

To create a routine group right click the schema in the [Object Tree](#) and choose the Create MySQL Routine Group option.

29.6. Creating Synonyms

This feature is not yet implemented.

29.7. Creating Structured Data Types

This feature is not yet implemented

29.8. Creating Sequences

This feature is not yet implemented

29.9. Saving a MySQL Workbench File

Before exiting be sure to save your MySQL Workbench file.

Chapter 30. MySQL Workbench Reverse Engineering

This chapter explains how to use the reverse engineering capabilities of MySQL Workbench.

Choose the Reverse Engineering ... menu item from the Database menu or alternately, press `Shift+Ctrl+R` to open the `Reverse Engineering` window.

Enter appropriate connection information. Once you have connected to the MySQL server the schemata will be displayed. Choose a schema to reverse engineer.

You have the option of importing some or all of the objects contained within the chosen schema. In the `Object Selection` window, click the `Detailed Selection > >` button to choose which objects to import.

Subsequent screens explain the process and, upon finishing, all the objects selected will appear on the canvas in the `Main View` tab.

Part VI. Appendices

Table of Contents

A. Installing From Source	127
A.1. Introduction	127
A.2. Downloading The Source Code	127
A.3. Building from Source under Windows	127
A.3.1. Prerequisites	127
A.3.2. Compiling and Building	127
A.4. Building from Source under Linux	129
A.4.1. Prerequisites	129
A.5. Compiling and Building	130
A.5.1. Building Static Binaries	131
A.5.2. Building RPMs	132
B. Troubleshooting Application Errors	133
B.1. Troubleshooting Connection Errors	133
B.2. Troubleshooting Display Problems	133
C. How Connections Are Stored	134
D. XML Files Common to the MySQL GUI Applications	135
E. XML Files Used by MySQL Query Browser	136
F. XML Files Used by MySQL Administrator	137
G. The MySQL System Tray Monitor	139
H. Third-party Software Used by the MySQL GUI Tools	141
H.1. PCRE Library	141
H.2. PNG Support	141
I. Notes for Translators	142

Appendix A. Installing From Source

A.1. Introduction

Normally, you will want to install the MySQL GUI tools from binaries. We put a lot of effort into making sure that our binaries are built with the best possible options. If, for whatever reason, you wish to compile the tools yourself, follow these instructions.

A.2. Downloading The Source Code

The source code for the MySQL GUI tools is available from our public Subversion trees at <http://dev.mysql.com/tech-resources/sources.html>.

You will need the `mysql-gui-common` tree regardless of which GUI tools you intend to build. The `mysql-gui-common` tree contains source and graphics files used for all MySQL GUI tools.

To be able to access these trees, you must have Subversion installed. Subversion is freely available from <http://subversion.tigris.org/>.

A.3. Building from Source under Windows

Installing from source under Windows requires multiple commercial compilers and because of this it is preferable to install from binaries. If you meet the prerequisites needed you can build the MySQL GUI tools from source under Windows.

A.3.1. Prerequisites

To compile the MySQL GUI tools from source under Windows, you need Borland Delphi 7, Microsoft Visual Studio.NET 2003, and some libraries which are available on the internet as open source software. Make sure the programs `DCC32.exe` and `DEVENV.exe` / ? can be launched from the command line.

You will need the following libraries to build from source:

- glib-2.0
- libxml-2
- mysql client libraries (4.0 or newer, 5.0 preferred)
- pcre-1.4
- MS Visual C runtime libraries 7.1
- Java include files 1.4.2_04
- Lua 5.0.2

A.3.2. Compiling and Building

Make sure you have cloned the `mysql-administrator`, `mysql-gui-common`, and `mysql-query-browser` trees, and that their local directories are located in the same directory, for example in `Personal Files\MySQLGUI\`.

The source files and libraries should be placed into the following tree structure:

```
Working Directory
|
- mysql-administrator
- mysql-gui-common
- mysql-query-browser
- mysql-gui-win-res
|
-include
|
  -glib-2.0
  -java
  |
  -win32
|
  -libxml
  -lua
  -mysql
  -pcre
  -windows
|
-lib
|
  -glib-2.0
  -java
  -libxml
  -lua
  -mysql
  -pcre
  -windows
```

A.3.2.1. Building from Command Line

To build from the command line open a windows command line prompt and change into the main repository of the tool (i.e. for MySQL Administrator this is `work\mysql-administrator`). There are three batch-files to start the build:

- `build_debug.cmd`
- `build_release.cmd`
- `build_commercial.cmd`

Execute one of those to generate a runtime image of the application. It will be generated in a `bin\windows` directory.

Another method is to doubleclick on one of the `build_buildtype.cmd` files, where `buildtype` is one of `debug`, `release`, and `commercial`. The `commercial` build script is for users with commercial licenses, and makes use of the commercial version of the MySQL client library.

A.3.2.2. Building from the IDE

Before you can open any projects, you have to install these components. Otherwise some of the forms will not open and you will receive error messages.

To install the components, goto the `mysql-gui-common\source\windows\` directory. Copy the following directories to your Delphi source directory (`C:\Program Files\Borland\Delphi7\Source\`).

- `png`

- SynEdit
- TNT
- UniCodeEditor
- VirtualTreeView

Then doubleclick each `*.dpk` file outlined below. It will open Delphi. Press compile and install to install the components. Repeat this for each `*.dpk` file.

- SynEdit\Packages\SynEdit_D7.dpk
- TNT\Packages\TntUnicodeVcl_D70.dpk
- VirtualTreeView\Packages\VirtualTreesD7D.dpk
- UniCodeEditor\Packages\UniCodeEditorD7.dpk

These components have to be installed via the Install option of the Component menu. Click on the [Browse](#) button and select the files below. Press [OK](#) after each file.

- png\PNGImage.pas

After installing the components you can open the projects as outlined below.

To build the application from the Delphi 7 IDE simply doubleclick the `edit_XXXXX.cmd`. This will launch Delphi 7 and open the file. To generate the complete runtime image select Build All Projects option from the Project menu.

Like for building from command line there are three different files.

- edit_debug.cmd
- edit_release.cmd
- edit_commercial.cmd

The runtime image will be generated in a `bin\windows` directory.

Opening the Delphi project not using the `edit_.cmd` script will lead to incorrect project settings and problems while building the projects.*

A.4. Building from Source under Linux

A.4.1. Prerequisites

The easiest way to install required packages is by using an automatic dependency resolving package manager, such as `apt-get` or `yum`. You can install these dependencies by hand, but it will be a very tedious process since some of them require other packages that require other packages and so on.

Common requirements are listed below. For some distributions the package name is also mentioned.

Table A.1. Required Components

Software	Version	Ubuntu 6	SuSE 10.x
gcc/g++	4	g++	gcc-c++
autoconf, automake lib-tool			
x11-libs		libx11-dev	xorg-x11-devel
glibc		libc6-dev	xorg-x11-devel
libstdc++		libstdc++6-4.0.dev	libstdc++devel
glib	2.6+	libglib2.0-dev	glib2-devel
gtk+	2.6+	libgtk2.0-dev	gtk2-devel
libglade2	2.5+	libglade2-dev	libglade2-devel
libsigc++	2.0.x	libsigc++-2.0-dev	libsigc++2-devel
glibmm-2.4	2.4	libglibmm-2.4-dev	glibmm24-devel
gtkmm-2.4	2.4	libgtkmm-2.4-dev	gtkmm24-devel
pcre	4.x or newer	libpcre3-dev	pcre-devel
pkgconfig	0.15+	pkg-config	pkgconfig
libxml2	2.6.x	libxml2-dev	libxml2-devel
MySQL 5.0 or 5.1 client libs and headers		libmysqlclient15-dev	

Additional requirements for building MySQL Query Browser are:

Table A.2. Requirements for MySQL Query Browser

Software	Version	Ubuntu 6	SuSE 10.x
libgtkhtml	2.11 or 3.x	libgtkhtml2-dev	
libgnomeprint	2.2	libgnomeprint2.2-dev	

A.5. Compiling and Building

To build any of the tools, you need to have the source tree for `mysql-gui-common` and the tree for the desired app (i.e.: `mysql-administrator`, `mysql-query-browser`) in the same top-level directory.

First build `mysql-gui-common`:

1. `cd mysql-gui-common`
2. `sh ./autogen.sh`
3. `./configure --prefix=/opt/mysql-administrator`

4. `make`
5. `make install`

The `--prefix` option can have whatever value you want, normally one would specify `/usr/local` (the default) or `/usr`, but for making precompiled binaries, something in `/opt` might be more appropriate.

After you have successfully built `mysql-gui-common`, you can build the individual tools (in this example we will build `mysql-administrator`):

1. `cd mysql-administrator`
2. `sh ./autogen.sh`
3. `./configure --prefix=/opt/mysql-administrator`
4. `make`
5. `make install`

If everything went well, things should be properly built and installed.

A.5.1. Building Static Binaries

Building static binaries is non-trivial under Linux, due to the enormous number of library dependencies that come from GNOME and associated libraries. Additionally:

- Many libraries do not come with their static (`lib*.a`) versions shipped, so we need to build them ourselves.
- Some libraries can't be linked statically, such as `glibc`, because of `libnss*.so`.
- Some libraries depend on data files and modules from the original package, esp. in case of `gtk/gnome`. Problems don't only appear when the user is missing these data files; things may not work because different distributions put files in different paths.

The current solution is to build a partially static binary, with the most common libraries left dynamically linked (`glibc` and `gtk`). `gtkmm`, `gtksourceview`, `gtkhtml` and dependencies (like `gnome`) are linked statically. The ideal solution would be to build custom binaries for each major/chosen distribution, but we don't have time/resources for that at the moment.

To build these binaries:

1. Look at the `source/linux/static_make` script for all the files that should be removed/moved out of `/usr/lib`.
2. Run the script.

That's needed to force the linker to look for the static version of each library and also to explicitly list the dependency libs that were otherwise linked to the dynamic libs.

A.5.2. Building RPMs

For building RPMs, a spec file is already supplied and will be made automatically after `./configure` is ran. The spec file expects a source tarball with the following structure:

```
mysql-administrator/ mysql-administrator/mysql-gui-common/* »  
mysql-administrator/mysql-administrator/*
```

The contents of each subdirectory is the entire source tree for each tool (Windows specific files and source metadata are optional, of course). You must make a `tar.gz` file and put it in `/usr/src/redhat/SOURCES` (or whatever is your rpm SOURCES directory).

Then, execute: `rpmbuild -ba mysql-administrator.spec`

That should build the srpm and rpm files for the tool.

Appendix B. Troubleshooting Application Errors

B.1. Troubleshooting Connection Errors

If an error occurs when you select a connection profile and attempt to establish a connection to the MySQL server named in the profile, a window containing error information will be displayed (see figure).

Figure B.1. Error message dialog



Normally, this means either that the credentials specified in the profile are wrong (wrong username, wrong password), or that you do not have privileges to connect to the MySQL server from the client machine. See [Causes of Access denied Errors](#) [<http://dev.mysql.com/doc/refman/5.1/en/access-denied.html>], for more information.

However, another cause of connection failure is that there might be a network problem that prevents you from reaching the machine where the MySQL server is running. For this reason, the error dialog provides a **PING** button. Clicking this button will send an ICMP PING request to the MySQL server host machine. If the host machine is available over the network, you will see something like this:

```
Reply from 127.0.0.1: Time=0ms TTL=128
Reply from 127.0.0.1: Time=1ms TTL=128
```

This would indicate that the machine is reachable, and that the cause of connection failure is not a network problem.

B.2. Troubleshooting Display Problems

If, under Windows, the controls in the windows look somehow displaced, select Start, Control Panel, Display. Select the **Settings** tab and click the **Advanced...** button. Change your DPI settings to **Normal (96 DPI)**.

If, under Linux, messages and captions appear very big, your display resolution is probably set too high. Use [gnome-font-properties](#) to change your default font to a smaller value or go to **Details** and set your resolution to **72** or some other more appropriate value.

Appendix C. How Connections Are Stored

On Windows, connection profiles are stored in `C:\Documents and Settings\UserName\Application Data\MySQL` directory, where `UserName` is the name of the current Windows user. On Linux, the files are stored in the `~/ .mysqlgui` directory. The settings are stored in a file called `mysqlx_user_connections.xml`. The contents of that file look like this:

```
<?xml version="1.0"?>
<user_connections>
<last_connection>1</last_connection>
<password_storage_type>2</password_storage_type>
<user_connection>
  <connection_name></connection_name>
  <username>root</username>
  <hostname>localhost</hostname>
  <port>3306</port>
  <schema></schema>
  <advanced_options/>
  <storage_path></storage_path>
  <notes></notes>
  <connection_type>0</connection_type>
  <storage_type>2</storage_type>
  <password></password>
</user_connection>
...
```

You can edit the file manually, but take care not to invalidate the XML. When applying changes by editing and saving the file, those changes will show up the next time you open the [Connections](#) section of the [Options Dialog](#). You do not need to restart your application for the changes to take effect.

Database administrators may want to edit the file and then copy it to another machine running the MySQL GUI tools. This makes it easy to have identical connection profiles on a number of machines, without having to set up those profiles individually. All XML configuration files for the MySQL GUI Tools are cross-platform compatible and can be transferred between Windows and Linux machines.

Appendix D. XML Files Common to the MySQL GUI Applications

On Windows, XML files common to all MySQL GUI applications are stored in the `C:\Documents and Settings\UserName\Application Data\MySQL` directory, where `UserName` is the name of the current Windows user. On Linux, the files are stored in the `~/mysqlgui` directory.

- `mysqlx_common_options.xml`: Stores options selected in the [Options dialog](#).
- `mysqlx_user_connections.xml`: This file is described in [How Connections are Stored](#).

Appendix E. XML Files Used by MySQL Query Browser

Besides the [common files](#), MySQL Query Browser uses a number of XML files for internal purposes.

- `mysqlqb_functions.xml`: Stores a list of functions and operators for use in the [Function Browser](#).
- `mysqlx_dbm_data_types.xml`: Stores a list of the available MySQL data types for use with auto-completion and with the [MySQL Table Editor](#)
- `mysqlx_dbm_charsets.xml`: Stores the available MySQL character sets.
- `mysqlqb_history.xml` : Stores the queries that make up the [History Browser](#)
- `mysqlqb_bookmark.xml`: Stores the queries that form the [Bookmark Browser](#)

Appendix F. XML Files Used by MySQL Administrator

Besides the documentation you are currently reading (which is written in DocBook XML), MySQL Administrator uses a number of XML files for internal purposes.

- `mysqlx-options.xml`: Stores options selected in the [Options dialog](#).
- `mysqlx-status-variables.xml`: Stores the status variables known to MySQL Administrator, so that they can be assigned descriptions in various languages (see below for the XML file containing those descriptions). Those descriptions will show up in the [Status Variables](#) tab of the [Health](#) section. You can select the language of the descriptions in the [General Options](#) section of the [Options dialog](#). That file also determines in which categories or subcategories a variable will appear, and the order in which categories and subcategories appear. However, it does not determine the order of variables within a category or subcategory, as these are sorted lexically by MySQL Administrator.
- `mysqlx-system-variables.xml`: Stores the system variables known to MySQL Administrator, so that they can be assigned descriptions in various languages (see below for the XML file containing those descriptions). Those descriptions will show up in the [System Variables](#) tab of the [Health](#) section. You can select the language of the descriptions in the [General Options](#) section of the [Options dialog](#). That file also determines in which categories or subcategories a variable will appear, and the order in which categories and subcategories appear. However, it does not determine the order of variables within a category or subcategory, as these are sorted lexically by MySQL Administrator.
- `mysqlx-translations-administrator.xml` : Stores descriptions specific to MySQL Administrator, like descriptions for status and server variables, or category names, or labels used in dialogs. Here is an excerpt from that file, which contains English and German texts:

```
<textgroup id="MySQLVariables">
<text id="status-var-group-General">
  <trans xml:lang="en">General</trans>
  <trans xml:lang="de">Allgemein</trans>
</text>
<text id="status-var-Uptime">
  <trans xml:lang="en">How many seconds the server has been
up.</trans>
</text>
...
<textgroup id="Form-InstanceForm">
  <text id="FileMI">
    <trans xml:lang="en">&File</trans>
    <trans xml:lang="de">&Datei</trans>
  </text>
  <text id="NewConnectionMI">
    <trans xml:lang="en">&New Connection
...</trans>
    <trans xml:lang="de">&Neue Verbindung
...</trans>
  </text>
  <text id="SaveConnectionMI">
    <trans xml:lang="en">&Save Connection
...</trans>
    <trans xml:lang="de">Verbindung &speichern
...</trans>
  </text>
  <text id="CloseMI">
    <trans xml:lang="en">Clos&e</trans>
    <trans xml:lang="de">Be&enden</trans>
  </text>
</textgroup>
...
```

This excerpt demonstrates a number of things:

- A description (called *text*) does not necessarily have to be translated (see `status-var-Uptime`). If there is no translation, English will be used instead.
- Adding another translation is as simple as adding another `<trans xml:lang="language">` tag.

Consequently, translating the software is a relatively painless task; however, please look at the [Notes for Translators](#) below.

- `mysqlx-translations-general.xml`: Stores terms and descriptions that appear not only in MySQL Administrator, but also in other graphical MySQL programs. Its format, however, is identical with the format of the `mysqlx-translations-administrator.xml` file. Here is an excerpt from that file:

```
<textgroup id="ResourceStrings">
<text id="OKBtn">
  <trans xml:lang="en">OK</trans>
</text>
<text id="CancelBtn">
  <trans xml:lang="en">Cancel</trans>
  <trans xml:lang="de">Abbruch</trans>
</text>
...
```

- `mysqlx-user-connections.xml`: This file is described in [How MySQL Administrator Stores Connections](#).

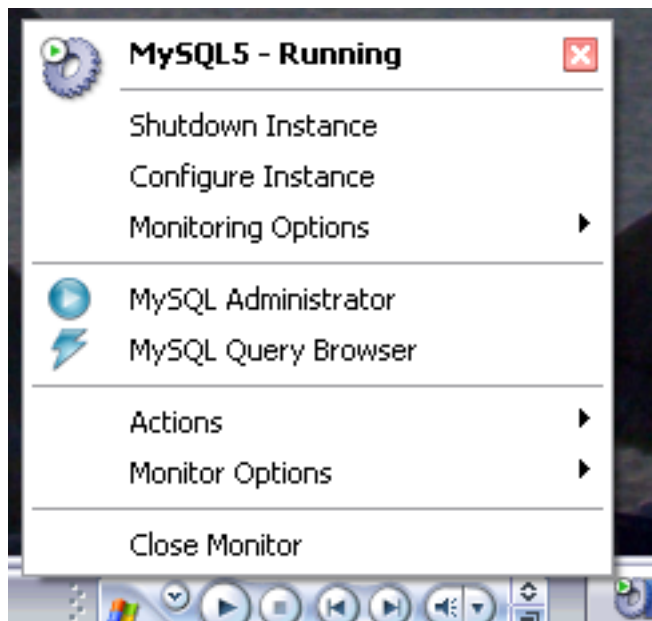
Appendix G. The MySQL System Tray Monitor

The MySQL System Tray Monitor is a tool that allows you to monitor the status of your local mysql server through an indicator that resides in the system tray. The MySQL System Tray Monitor also gives quick access to the various MySQL GUI tools through its context menu.

To start the MySQL System Tray Monitor, you can either execute the `MySQLSystemTrayMonitor.exe` application that resides in your MySQL Administrator installation directory, or choose the MySQL System Tray Monitor option of the Tools menu.

Once the MySQL System Tray Monitor is running, it resides in the system tray and provides visual status information for your MySQL server. A green icon is displayed at the top left corner of the MySQL System Tray Monitor if the MySQL server is running, and a red icon is displayed if the service is stopped. In addition, the CPU load may optionally be displayed along the right side of the MySQL System Tray Monitor.

Figure G.1. The MySQL System Tray Monitor



If you right-click on the MySQL System Tray Monitor, a pop-up menu will appear, listing all installed MySQL services and their individual statuses. For each service, the following options are available:

- [Shutdown Instance](#): this will begin the shutdown operation on the selected running service.
- [Start Instance](#): this will start a currently stopped server instance.
- [Configure Instance](#): this will launch the MySQL Administrator in configure service mode for the currently selected instance. See [Configure Service Mode](#) for more details.
- [Monitoring Options](#): currently the [Monitoring Options](#) submenu has one available option: [Ignore Instance State](#). [Ignore Instance State](#) will cause the MySQL System Tray Monitor to stop monitoring the selected service.

In addition to the service controls, you can also start the MySQL Query Browser and MySQL Administrator from the pop-up menu.

The [Manage MySQL Instances](#) option of the [Actions](#) menu can be used to start MySQL Administrator in configure service mode so you can manage your server configurations. See [Configure Service Mode](#) for more details.

The [Monitor Options](#) menu has two options, which affect the general behavior of the MySQL System Tray Monitor. [Display CPU Load](#) determines whether the CPU usage is displayed on the right of the MySQL System Tray Monitor. [Launch Monitor after Login](#) sets the MySQL System Tray Monitor to launch when you log into Windows.

Appendix H. Third-party Software Used by the MySQL GUI Tools

The MySQL GUI tools incorporate PCRE and PNG support through the use of third-party libraries.

H.1. PCRE Library

Regular expression support is provided by the PCRE library package, which is open source software, written by Philip Hazel, and copyright by the University of Cambridge, England. The source for the PCRE library can be found at: <ftp://ftp.csx.cam.ac.uk/pub/software/programming/pcre/> [<ftp://ftp.csx.cam.ac.uk/pub/software/programming/pcre/>].

H.2. PNG Support

PNG support for the Windows version is provided by the TPNGImage component which is open source software, written and copyright by Gustavo Daud. The source for the TPNGImage can be found at: <http://pngdelphi.sourceforge.net> [<http://pngdelphi.sourceforge.net>].

Appendix I. Notes for Translators

You may find that the MySQL GUI tools and/or their documentation are not available in your preferred language. If you would like to translate the software, or its documentation, or participate in that translation, *please contact the MySQL documentation team before starting your translation.*

The MySQL documentation team's address is <docs@mysql.com>. In your email, please state what you would like to translate (the tools, their documentation, or even the whole MySQL reference manual if you dare), and give some information regarding your background:

- What is your name?
- Where do you live?
- How long have you used MySQL?
- Have you done other translations?
- Will you work alone or with a group of translators?
- What's your timetable regarding the translation?
- What is your motivation for translating?

We will get back to you as soon as possible.

The format used for both software and documentation is **DocBook XML**. From that base format, all other available formats (HTML, CHM, PDF, etc.) are being generated. It would be a pity if you started translating, say, the HTML version of this documentation, because we will not be able to use it in that format. Also, you might find that someone else has already done (or is in the process of doing) that translation. So, once again, please contact the MySQL documentation team first.