

SQL

- [MySQL Resources](#)
- [PostgreSQL Resources](#)
- [Why doesn't mysqlshow work for databases or tables with underscores in their names?](#)
- [What is mysqlshow good for?](#)
- [How can I search/replace strings in MySQL?](#)
- [Microsoft Access, OpenOffice and MySQL](#)
- [SQL joins](#)
- [Get rid of default annoyances in MySQL Workbench](#)
- [Who uses PostgreSQL at UCLA?](#)
- [Why NoSQL Matters](#)

MySQL Resources

MySQL is a free, open-source relational database that has been used in thousands of websites, large and small. This is a contributed collection of resource links on MySQL. Feel free to add your favorites.

- Mike Franks – SSC
 - Collection of Mysql links and articles <http://www.sscnet.ucla.edu/mysql/>
 - Notes while learning Mysql <http://www.sscnet.ucla.edu/mysql/notes.htm>
 - Article on GRANTing access
http://www.devshed.com/Server_Side/MySQL/Access/print_html
 - Article I don't understand on storing hierarchical data in a relational database. //would this work with Mysql?\\
<http://www.oreillynet.com/pub/a/network/2002/11/27/bioconf.html>
 - Quick Tips to Optimize MySQL
 - 8 Reasons Not to Use MySQL
 - 5 Reasons To Use MySQL
- Ian Roessle November 9, 2002 at 13:23:
 - [phpMyAdmin](#) is a very useful tool for managing Mysql databases from the web. And this article tells how to [Turn on phpMyAdmin's hidden features](#)
- Client-side database tools
 - [Aqua Data Studio](#) is a database query and administration tool that allows developers to easily create, edit, and execute SQL scripts, as well as browse and visually modify database structures. My favorite features: auto-completing SQL editor, and Excel-like editing of SELECT resultsets – the software generates automatically the appropriate UPDATE commands. Written in Java and supported on Windows, OS X, Solaris and Linux. Comes with JDBC drivers for MySQL, MS SQL Server 7/2000, Oracle 8/9/10, PostgreSQL and Sybase, and you can add drivers for other DBs. Runs fast. Free for educational use.
 - [DBDesigner4](#) is a visual database design system that integrates database design, modeling, creation and maintenance into a single, seamless environment. I use it to visually layout and plan my database schemas. Free. Only on Windows and Linux

- [Point-and-Click MySQL Administration](#) – Linux Magazine – Feb. 1, 2007
- [The MySQL Query Browser](#) – Linux Magazine April 19, 2007
- Server-side Admin
 - [Performance Tuning Best Practices for MySQL](#) – Google Tech Talks – 43 min – Apr 28, 2006 *Recommended by Harry Mangalam, UCI on [UC-CSC Mailing List](#) “This is a bit dated, but since so many web services use MySQL, this may still be useful.. It’s 43 min long, but the guy has a lot of useful hints and if you let it run in the background, you’ll pick up some things that will help, especially if you’re confused about the diffs between myisam and innodb tables and how to tune each.”*
 - [Quick and Dirty MySQL Performance Troubleshooting](#) – Jeremy Zawodny, Monday, August 17th, 2009 [Linux Magazine](#)
- General documentation and training
 - [MySQL University](#) provides free download of talks on various topics.
 - [Be productive with the MySQL Command Line](#)
 - [Fun with the MySQL pager command](#)

This article was originally posted on the UCLA Programmers Wiki.

PostgreSQL Resources

See Also:

- [Who uses PostgreSQL at UCLA?](#)

PostgreSQL official website:

- <http://www.postgresql.org>

Modified versions:

- [Bizgres](#)
- [ExtenDB](#)

Tuning Guides:

- [PostgreSQL 8.0 Performance Checklist](#)
- [How the Planner Uses Statistics](#)
- [Server configuration options](#)

Graphical Admin Tools:

- [pgAdmin](#)
- [Aqua Data Studio](#)
- [EMS SQL Manager for PostgreSQL](#)

Web-based Admin Tools:

- [phpPgAdmin](#)

Why doesn't mysqlshow work for databases or tables with underscores in their names?

mysqlshow has a tricky feature that interprets SQL wildcard characters (*,?,%,_) as wildcards if they appear in the last argument you give it. While the first 3 characters don't get used much (if at all) in naming databases, the underscore is a common word separator and could cause problems if you try to use mysqlshow. Instead of showing the details of a database or table, if the object of interest has an underscore in its name, mysqlshow will treat it as a search and display only the name of the one object it finds.

The trick to getting mysqlshow to display details is to add another argument at the end. mysqlshow only interprets wildcards in its last argument, so adding another one will tell the command to interpret the underscores literally and not as expressions for a search.

To get the database `underscores_are_great` to show its tables, type:

```
mysqlshow underscores_are_great %
```

The '%' on the end is a wildcard which will show all tables, which is what you want it to do.

This works for displaying individual tables too, just add a '%' as a final argument.

[See what else mysqlshow is good for](#)

What is mysqlshow good for?

mysqlshow is a command-line tool included with standard MySQL distributions, similar to mysqladmin or other tools. mysqlshow is used to show summary information about databases and tables.

Here is its basic usage:

```
mysqlshow [ OPTIONS ] [ database [ table [ field ] ] ]
```

There are four basic modes:

- **No database given** – all databases on the server are shown
- **Database given only** – all tables in that database are shown
- **Database and table given** – all fields in the given table of the database are shown
- **Database, table, field given** – only that field of the table is shown

There is also a search mode, which is activated if the last argument contains an SQL wildcard (*,?,%,_). In this mode, all objects at the level of the last argument are searched. For example,

```
mysqlshow ucla c%
```

 matches all tables in ucla starting with the letter 'c'.

Numerous options are available:

- -i shows extra information on the table view mode (when you specify only the database name), such as the storage engine, number of rows, creation and modification date.
- -v gives more info in all modes. In database listing mode, it shows the number of tables in each database. In table listing mode, it lists the number of columns. This option can be invoked multiple times for even more info!
- Many more (do `mysqlshow --help`) for a complete list for your version.

Gotchas:

- mysqlshow may not behave as you want if your database or table name has an underscore

in it. See the fix [here](#).

Also see the [mysqlshow entry](#) in the MySQL manual.

How can I search/replace strings in MySQL?

MySQL lets you replace all occurrences of a character or string of characters in a table with some other character or string.

```
UPDATE table SET field1 = REPLACE(field1, 'replace_that', 'with_this'), field2 = REPLACE(field2, 'and_that', 'with_this')
```

As an example, let's say that you have a pets table and you want to change every cat into a dog:

```
UPDATE pets SET species = REPLACE(species, 'cat', 'dog')
```

This feature is also useful for transposing characters from one encoding to another. Here's how to change Windows-style CRLF line endings into Unix-style LF line endings:

```
UPDATE users SET bio = REPLACE(bio, CONCAT(CHAR(13), CHAR(10)), CHAR(10))
```

Of course, you don't have to replace **all** occurrences in a table if you don't want to. Just supply a WHERE clause and only those rows will be affected.

Microsoft Access, OpenOffice and MySQL

Using ODBC it is possible to connect Microsoft Access (or OpenOffice) to a MySQL database.
Access or OpenOffice can even be used as “front end” to MySQL.

The “MySQL Connector/ODBC” drivers are available at: <
<http://dev.mysql.com/doc/refman/5.0/en/connector-odbc.html>>

Installation information and documentation is available for Windows, Mac OS X and Unix platforms.

WindowsXP: <http://dev.mysql.com/doc/refman/5.0/en/connector-odbc-installation-binary-windows.html>

MacOS X: <http://dev.mysql.com/doc/refman/5.0/en/connector-odbc-installation-binary-macosx.html>

SQL joins

Please add other helpful links:

SQL joins: a visual explanation

<http://www.codinghorror.com/blog/2007/10/a-visual-explanation-of-sql-joins.html>

Get rid of default annoyances in MySQL Workbench

By default, if you make any changes to table rows, there is an annoying 2-step confirmation dialog when you click Apply. It also prevents you from doing many mass UPDATES and DELETES (this is called 'Safe Updates').

Safe Updates is a good idea in principle, since it stops you from running a DELETE without a WHERE clause, for example. But at the same time it prevents you from running an UPDATE/DELETE with a WHERE based on a column other than an ID, which I find myself doing sometimes.

The confirmation dialog seems utterly pointless. You already have to click the 'Apply' button in the lower right corner to get your query to run, anyway.

To turn off these options:

1. In the menu, go to Edit > Preferences...
2. Click on the 'SQL Queries' tab
3. In the 'General' section, uncheck 'Safe Updates'
4. In the 'Query Results' section, uncheck 'Confirm Data Changes'

These instructions are for version 5.2.47.

Who uses PostgreSQL at UCLA?

The purpose of this page is to create a directory of people who are using PostgreSQL and what advanced features they are using. Please try to mention how you have found certain features useful so that you can give other people ideas on how to better utilize the database.

Jason Fong

Database's purpose: collecting large sets of data, and querying data (usually big long queries) to discover trends or other interesting results

Total database size: approximately 500 GB

Largest table: approximately 500 million rows

Advanced features in use:

- Procedural languages:
 - PL/pgSQL
 - PL/perl: This is very useful— it lets one run bits of Perl code inside of a query. This can also be used to create a Perl function index on a table so that the Perl code doesn't even need to be run for future queries.
- Triggers: I use these for a number of things:
 - Automated auditing: a table can have a trigger that inserts a row into a log table whenever any insert/update/delete occurs
 - Duplicate row detection: A table can have a trigger on inserts that checks for a duplicate primary key and then compares the other column values to the matching row. If all the columns are the same, the insert can be dropped. If any column is different, then the rows can be merged (or the new row could be inserted into a separate table to keep track of these duplicate primary keys with different column values).

Why NoSQL Matters

Useful survey of NoSQL (and SQL) technologies and a summary of the philosophy...

<http://blog.heroku.com/archives/2010/7/20/nosql>

Taken from email by Jose Hales-Garcia to [Campus Web Publishers List](#)