

I don't like MySQL for various reasons, most important one being that I saw PostgreSQL early on and never looked back.

Here are few quick commands to ease my pain:

```
mysql> select count(*) from table_name
```

Contents: [Dobrica Pavlinu's random unstructured stuff]

- Dobrica Pavlinu's random unstructured stuff (Number of rows in table)
- Dobrica Pavlinu's random unstructured stuff (Create grepable dump)
- Dobrica Pavlinu's random unstructured stuff (Poor man's graphs)
- Dobrica Pavlinu's random unstructured stuff (MySQL Koha utf-8 encoding)
  - ◆ Dobrica Pavlinu's random unstructured stuff (connect)
- Dobrica Pavlinu's random unstructured stuff (my.cnf performance optimization)
- Dobrica Pavlinu's random unstructured stuff (Performance monitoring)
- Dobrica Pavlinu's random unstructured stuff (Linux tuning)

## Number of rows in table

`select count(*) from table_name` is terrably slow in MySQL, so use [information schema](#) to quickly have report...

```
mysql> select table_name,table_rows,avg_row_length,data_length
from information_schema.tables
where table_type = 'BASE TABLE'
order by table_rows desc
limit 10 ;
```

table_name	table_rows	avg_row_length	data_length
zebraqueue	2309022	87	200998912
sessions	1744788	675	1178402816
items	302540	126	38338560
biblio	268391	166	44646400
action_logs	225797	341	77168640
biblioitems	220004	8159	1795080192
items_print_log	205559	104	21528576
statistics	130166	76	9977856
old_issues	43852	84	3686400
import_biblios	19352	136	2637824

10 rows in set (1.90 sec)

## Create grepable dump

Put every value in it's own line

```
mysqldump -u user -p database --extended-insert=0 > dump.sql
```

# Poor man's graphs

```
mysql> select dateenrolled,count(*),rpad(' ',count(*) / 10,'**') from borrowers group by dateenroll
```

dateenrolled	count(*)	rpad(' ',count(*) / 10,'**')
NULL	18	**
2008-04-03	1	
2008-04-17	1	
2008-05-28	1	
2008-06-05	1	
2008-06-12	3	
2008-07-28	1	
2008-10-17	1	
2008-10-20	1	
2008-10-21	1	
2008-10-27	11	*
2008-10-28	2	
2008-10-29	4	
2008-10-30	2	
2008-11-03	1	
2008-11-04	9	*
2008-11-06	7	*
2008-11-10	16	**
2008-11-11	1	
2008-11-13	1	
2008-11-14	4	
2008-11-17	1	
2008-11-20	1	
2008-11-26	1	
2008-12-02	1	
2008-12-19	2	
2009-01-07	2	
2009-03-16	301	*****
2009-03-17	365	*****
2009-03-18	439	*****
2009-03-19	324	*****
2009-03-20	226	*****
2009-03-23	105	*****

```
33 rows in set (0.00 sec)
```

## MySQL Koha utf-8 encoding

### connect

```
use DBI();  
  
my $dbh = DBI->connect ('dbi:mysql:test_db', $username, $password,  
    {mysql_enable_utf8 => 1}  
);
```

## my.cnf performance optimization

- <https://launchpad.net/mysql-tuning-primer>

- <https://launchpad.net/mysqltuner/>
- <http://github.com/rackerhacker/MySQLTuner-perl>

## Performance monitoring

```
./mk-query-digest --processlist localhost --interval 0.01
```

## Linux tuning

```
# turn off readahead  
hdparm -a 0 /dev/sdb
```

```
# noop/deadline better than cfq  
echo noop > /sys/block/sdb/queue/scheduler
```