

# MySQL: The Commercial OpenSource Database

*David Axmark*

*CoFounder & VP of Community Building*

*MySQL AB*

*MySQL AB Creators of MySQL*

*Tokyo 2002-02-14*

<http://www.mysql.com>

- Why use MySQL
- Mini History of MySQL
- The Company and how we make money
- Technical features & Storage Engines
- MySQL version 4.0 and beyond
- NuSphere
- Questions

# Why use MySQL

- „ Fast and reliable
- „ Designed for Mission Critical / Heavy load systems
- „ Easy to install and use
- „ Well documented (in Japanese & English)
- „ Extensively used worldwide (more than **10000** server downloads per **day**)
- „ Transactional/Non transactional storage engines
- „ Low cost of ownership
  - Dual licensed under GPL and Commercial licenses

# Some References

- " **Aizawa Securities**
- " **Japanese Environmental Information Center**
- " **NASA**
- " **Yahoo Finance**
- " **Slashdot.org, Freshmeat.net & Linux.com**
- " **Cisco**
- " **Nortel(InSight) / Ericsson / Alcatel / Telia / Nokia**
- " **Motorola / Compaq**
- " **Handles mission critical application under heavy load**

## " # installations

- Difficult to judge
- Included in Linux distributions
- For example 1 million copies has been shipped with MacOS X

## " >10000 server downloads per day

## " Goto.com bids (nr/max)


- MySQL 22/0.53
- Oracle 55/2.60
- Postgres 1/0.05

## " Links according to Google.com

- MySQL 24300
- Oracle 16700
- Postgres 6170

## " Mentions according to Google.com

- MySQL 4400000
- Oracle 5990000
- Postgres 532000



- " LAMP is a model for developing web applications
  - **Linux**
    - " Operating system that is gaining marketshare for servers
  - **Apache**
    - " Webserver with more than 50% of the word market
  - **MySQL**
    - " Database developed with the web in mind
  - **PHP**
    - " Scripting language with over 6 million installed domains
- " **LAMP** was coined by Michael Kunze



# Aizawa Securities Co., Ltd.

- " A Mission-critical web based customer system for viewing transactions
- " 100000 new transactions are loaded per day
- " "We get much more performance out of MySQL Server than we did with Oracle"
- " Technologies used: MySQL, PHP, Ruby, Apache on IBM AIX servers
- " Aizawa is one of the leading Japanese Securities Houses, with a market capitalization of 3200000000 Yen







# Environmental Information Center

- " Web site about Japanese Environmental issues
- " **LAMP** (Linux/Apache/MySQL/PHP) Solution
- " "Because, to provide and search much information about environment intelligibly and quickly, we required fast engine"
- " PHP needed a patch for Japanese characters but MySQL supported Hiragana, Katakana and Kanji character sets by default



- " MySQL has very good documentation as a online reference
- " About 600 pages, updated as new code is written
- " A printed version in English will appear very soon
- " MySQL AB now has a new full time technical writer
- " A partial Japanese translation exists
- " A lot of third party books are also available



# Some MySQL Books in Japanese



" I will not try to write any titles our authors here :-)

MySQL™

# *Some MySQL Books in English*



*MySQL by Paul DuBois*

*MySQL by Michael Kofler*

*MySQL & PHP from scratch*

*MySQL in 21 Days*

*MySQL/PHP Database Applications*

*PHP and MySQL Web Development*

*JSP, Servlets, and MySQL*

## *Some Awards and notable events*

- " 1998 WEBTechniques*
- " 1998 CNET Builder.com: Best affordable database*
- " 1998, 1999, 2000 & 2001 Linux Journal: Readers choice*
- " 2000 & 2001 Linux Magazine: Editors choice*
- " 2001 MySQL AB is formed and takes investment*
- " 2001 Oracle introduces MySQL migration kit*
- " 2001 Linux New Media (Germany): Best commercial software*

# Background

- " *Both Monty and I used a lot of OpenSource software when we started working on UNIX in the early 80s*
- " *I heard Richard Stallman talk in Stockholm 1996*
- " *Needed a SQL Database for our web pages*
  - *Technical reasons: Our tools was not suited for web use*
  - *Political reasons: Our customers wanted standards and accessibility from Windows (ODBC)*
- " *So we decided to develop our own commercial database inspired by OpenSource projects*
- " *MySQL development started 1995 based on experience and code from our previous database*

# *How MySQL started*

- " *1996: MySQL 3.11 publicly released*
- " *1997: First commercial licenses and support agreement sold*
- " ***2000: We change the license to the GPL***
  - *We wanted maximum spread of MySQL and needed a standard license*
- " *2001: MySQL AB is being sued by NuSphere and is suing NuSphere for violating the GPL*
- " *2001: A single MySQL AB is formed*
  - *Lead investor ABN Amro one of europe's largest investors*

# MySQL AB



- " *Distributed company*
- " *Head office in Uppsala, Sweden*
- " *Currently people in over 10 countries*

MySQL™




# MySQL AB Revenue

- " *By **Commercial Licences** (such as embedded)*
- " *By **Support Agreements***
- " *By **Partnership Agreements***
- " *By **Training Courses***
- " *By **Professional Services (Consulting)***
- " *By **minimal advertising** on our website*
- " *We get this question all the time!*
- " *Place orders online at [order.mysql.com](http://order.mysql.com)*



# Training



- " *Advanced MySQL training is now provided by MySQL AB (In English, German, French, Swedish and Finnish)*
- " *Scheduled courses (Feb-Mar) in Hamburg, San Francisco, Toronto, Amsterdam, Paris and Brisbane*
- " *[mysql.com/training](http://mysql.com/training)*

- " Modular design*
- " Never pay in performance for unused features*
- " Easy to install and use*
- " Write documentation while writing code*
- " Occams razor: No complexity beyond necessity*

A large, stylized version of the MySQL logo, with "MY" in blue and "SQL" in yellow, followed by a trademark symbol (TM). A faint blue outline of a fish is visible in the background behind the text.

# MySQL Benchmarks

- " *Most benchmarks give a single number*
- " *Intention is to show speed of single operations*
  - *This shows tradeoffs in each database!*
- " *Tests lots of things like connect, simple inserts, join of two tables using a key and sorting*
- " *Tests only common functionality between databases*
- " *Need more tests: multi-user, transaction etc*
- " *Written in Perl/DBI, so a good DBI/DBD driver is required for the benchmark to work.*

# Crash-me

- *Crash-me tests what a SQL server can do by actually trying to do it*
  - *It finds out about things like*
    - *Maximum size of strings, query's, expressions ...*
    - *What ANSI/ODBC functions and types are supported*
  - *Because of extreme query's some databases crash*
  - *Very useful for writing portable applications*
  - *Crash-me is **not** a feature list!*
- Both actively worked on again by MySQL AB*

# MySQL Features

- " An extended subset of ANSI SQL89.
- " Numerous extensions like SELECT ... LIMIT 10
- " Very few limits: disk, memory and OS set the limits
- " Localized to many languages (German, Japanese, Chinese)
- " Powerful security system
- " Master-Slave Replication for High Availability
- " Per table choice of backend (with/without transactions)
- " Can be extended with User Defined Functions (UDF)
- " MyISAM has very compact tables
- " Embeddable engine

# Language Support

***PHP***

***C***

***Java (JDBC)***

***Ruby***

***Guile***

***ADA***

***Dylan***

***Common Lisp***

***Perl***

***Python***

***Pike***

***C++***

***TCL***

***Delphi***

***ODBC***

***And lots more ...***

# *Supported platforms*

## *Native Threads*

*Sun Solaris 2.5+*

*Linux (x86, Alpha, PowerPC,  
Sparc)*

*HP UX 11+*

*IBM AIX 4+*

*Apple MacOS X*

*DEC (Tru64) UNIX*

*BSDI 3+*

*FreeBSD 3+*

*IBM OS/2*

*SCO Unixware 7*

*SGI Irix 6.x*

*Windows 95, 98, NT and  
2000*

## *Emulated threads (1 CPU)*

*BSDI 2.x*

*FreeBSD 2.x*

*HP HPUX 10.20*

*NetBSD Intel/Alpha*

*OpenBSD 2.x*

*Sun SunOS 4*

*SCO Openserver*

*AmigaOS*

- " A storage engine is a low level data storage / retrieval module (disk or memory)
- " This allows you to choose locking and speed tradeoffs per table (instead of when choosing db!)
- " MySQL supported multiple storage engines from the very beginning
- " **CREATE TABLE TYPE=HEAP** (key int, value char(10), PRIMARY INDEX key);
- " **ALTER TABLE table\_name TYPE=InnoDB;**



# *Storage Engine: MyISAM*

- „ Developed by MySQL AB (replaced original ISAM)*
- „ Static, dynamic and compressed (read-only) row formats but no transactions*
- „ Text and compressed indexes*
- „ Data and indexes in separate files*
- „ Fast read/write performance but low r/w concurrency*
- „ Extremely good concurrency in the select and insert at end case (logs)*
- „ External check and repair program (myisamchk)*
- „ Especially useful for websites & logging*

# *Storage Engine: InnoDB*

- " Actively developed code from InnoDB Oy*
- " Full transactions (ACID) with **row** level locking*
- " Better concurrency than MyISAM for read/write on the same table*
- " Consistent reads (Oracle style MVCC)*
- " Uses tablespaces instead of individual files*
- " MySQL AB provides support for InnoDB through a contract with the creators/developers*
- " Is included in MySQL 4 & the MySQL Max binaries*
- " Has now been in active use under heavy load*

# Storage Engine: Merge

- " Developed by MySQL AB
- " Merge tables is a collection of identical MyISAM tables that can be used as one table
- " Some of the tables can be compressed
- " ALTER TABLE is used to change the set of tables
- " Insert/Delete/Update/Select works
- " VERY useful for logging systems since MyISAM has perfect locking for this case



# Storage Engine: Hash/InMemory

- „ Developed by MySQL AB
- „ Completely in Memory with very fast **hash** based indexing
- „ Insert/Delete/Update/Select works
- „ Useful for
  - Temporary tables
  - Lookup tables

MySQL™

# *Storage Engine: BerkeleyDB*

- " Developed and supported by Sleepycat Software Inc*
- " Full transactions (ACID) with **page** level locking*
- " Better concurrency than MyISAM for read/write on the same table*
- " Primary key lookups can be faster than MyISAM*
- " Data and indexes in one file per table*
- " Is included in the MySQL Max binary*
- " Can also be used at a lower level independent of MySQL*

# High Availability by Replication

- MySQL supports many sites that need **high** reliability
- This is done by replicating the system to many machines
- Examples of users are
  - Yahoo Finance (And other parts of Yahoo)
  - Mobile.de (sells used cars in Germany, Over 100 Million page views per month)
  - Slashdot.org & Slashdot.jp (The very popular tech news site)

# Replication Technology

- " The single **Master** keeps binary log of SQL commands that update data
- " **Slaves** connect to the master or another slave to read, and rerun the updates
- " In MySQL 4.0 the slaves use two threads
  - One to read the all queries
  - One to actually do the updates
- " This makes sure that every slave has all data even if the master goes down while the slave is working on a slow query

# Case Study - A Sweepstake Portal

- „ Case study: A sweep-stake portal
- „ 6 web+6 web servers for dynamic/static content
- „ Each dynamic hit involves a database query
- „ In June 2001, the site served 91 million pages and handled 342 million hits (max 1 mill hits/hour)
- „ Platform: **LAMP** (Linux/Apache/MySQL/PHP)
- „ Hardware power range (totally about 15 Machines):
  - Dual Pentium 600 Mhz with 1GB RAM
  - Quad Xeon 700 Mhz with 2 GB RAM





# MySQL Features

- " An extended subset of ANSI SQL89.
- " Numerous extensions like SELECT ... LIMIT 10,10
- " Few limits: disk, memory and OS set the limits
- " Handles large datasets (hundreds of GB)
- " Powerful security system
- " Master-Slave Replication
- " Per table choice of backend (with/without transactions)
- " Can be extended with User Defined Functions(UDF) in C
- " Embeddable engine
- " Full text search

# New MySQL 4.x Features

- " **Reconstruction of file formats and table management**
  - **Better scalability, Faster and easier to implement new features**
- " **Embedded server library**
  - **Use mysql internal to you application without a separate process**
- " **UNION & TRUNCATE table SQL commands**
- " **DELETE & UPDATE with multiple tables (MS Access Syntax)**
- " **More features for Replication**
  - **Better safety and more features**
- " **Handler interface: HANDLER name READ INDEX FIRST™**
  - **Usefull for porting applications from older ISAM like systems**



# New MySQL 4.x Features

## Query cache

- Improves the speed of queries on read mostly tables. Most (all?) websites has some query that will run many times a second

## New fulltext search

- Much faster bulk updates



- " Foreign keys
- " Change options on the fly (per thread and globally)
- " Secure connections with SSL
- " New Key cache with better concurrency
- " Boolean operations for fulltext search
- " Online backup of MyISAM tables
- " Sub Select (Yes we know this is very late)
- " Fail-safe replication

- " MySQL is actively developed and tested on Windows
  - But windows is still not recommended as a high load platform (UNIX performs much better)
- " MyODBC active developed again by a MySQL AB
  - Cursor support
  - Transaction support
  - Lots of small annoying long lived bugs fixed
- " A ODBC 3.5 compliant driver was released in January

## Temporary tables

- Stays around until you drop them or your thread dies
- Can be in memory (TYPE=heap), on disk (TYPE=MyISAM or TYPE=InnoDB)
- `CREATE TEMPORARY TABLE type=MyISAM ... SELECT a+b as isum, ...`
- **CREATE TEMPORARY TABLE** sums  
(PRIMARY KEY (id), INDEX (total), comment TEXT)  
`SELECT t1.id, SUM(t2.number) AS total FROM t1,t2 WHERE t1.id = t2.id GROUP BY t1.id;`

# Why use Free Software

## " Independence

- Access to all technical information and not only a vendor chosen subset
- Any good technical person has the same chance to become an expert as company employees

## " Competing commercial services available

## " Expanded use because of lower investment in time and money before

- No extra money for advanced features like replication/free text search

## " Result: Low Total Cost of Ownership

# *OpenSource databases are better*

*People invest more time in their bug reports when the possibility of a quick fix exits*



*Good testing since you get more users*

*You can hire people who already know the code*

*Lots of code gets written by others (for us especially in the client area)*

*A repeatable bug report is worth its weight in gold!*



# MySQL AB vs NuSphere

- " NuSphere sued MySQL AB, and David & Monty personally in June 2001 for unfair competition, terminating of our interim agreement and interference with third party contracts
- " MySQL AB answered with suing NuSphere for **violating** our Trademark and the **GNU General Public License (GPL)**
- " We have tried to solve this fast but NuSphere has deployed a delaying and misinformation tactic. Examples:
  - ❑ Saying that Gemini was released under the GPL in January when the source was released in July/August after we sued them (the worlds first **GPL release without source code?**)
  - ❑ That we violate the GPL (even if the FSF say we do not!)
  - ❑ That we wanted them to assign the whole of Gemini when we only wanted the glue code

# Questions!

- " *Now it is your turn to ask questions!*
- " *For more information about MySQL in English go to:*

**[www.mysql.com](http://www.mysql.com)**

- " *MySQL Nippon Assiation*

**[www.mysql.gr.jp](http://www.mysql.gr.jp)**

- " *And for our Japanese partner:*

**[www.softagency.co.jp](http://www.softagency.co.jp)**