State of MySQL (ver 5.0/5.1)

Michael "Monty" Widenius
CoFounder

MySQL AB, Creators of MySQL
Romania, 2007-05-19

Languages

- **Ruby**
- **C**
- **Java, MXJ (Native)**
- **C#/.Net**
- **ODBC**
- PHP
- Perl
- C++
- Python
- Delphi
- Objective C
- Visual Basic
- Smalltalk
- Pascal
- ADA
- APL
- Lasso
- Pike
- Rexx
- Dylan
- Common Lisp
- Scheme
- Gauche
- Guile
- Mathlab
- Eiffel
- Haskell
- Erlang
- Curl
- Forth
- Slang
- LUA
- OLEDB
- Active X
- TCL
- Fortran

And even Cobol!

The community are always adding more languages!

Platforms

- Linux
  - Ubuntu, RedHat, Suse, Debian, Feodora, Turbo WindRiver, MontaVista
- UNIX
  - OpenSolaris, HPUX, AIX
- Windows
  - Vista, NT, Win2k, XP
- MacOS X
- {Free,Open,Net}BSD
- IBM i5/OS

- Novell Netware
- OpenVMS
- QNX

- Intel [32 & 64]
- AMD [32 & 64]
- IBM PowerPC [32 & 64]
- Sun Sparc [32 & 64]

First 64bit MySQL in March 2000. Good code -> 64bit is a recompile!

All compiled from ONE source tree.
Code should be written with portability in mind from the beginning!

MySQL Internal Architecture

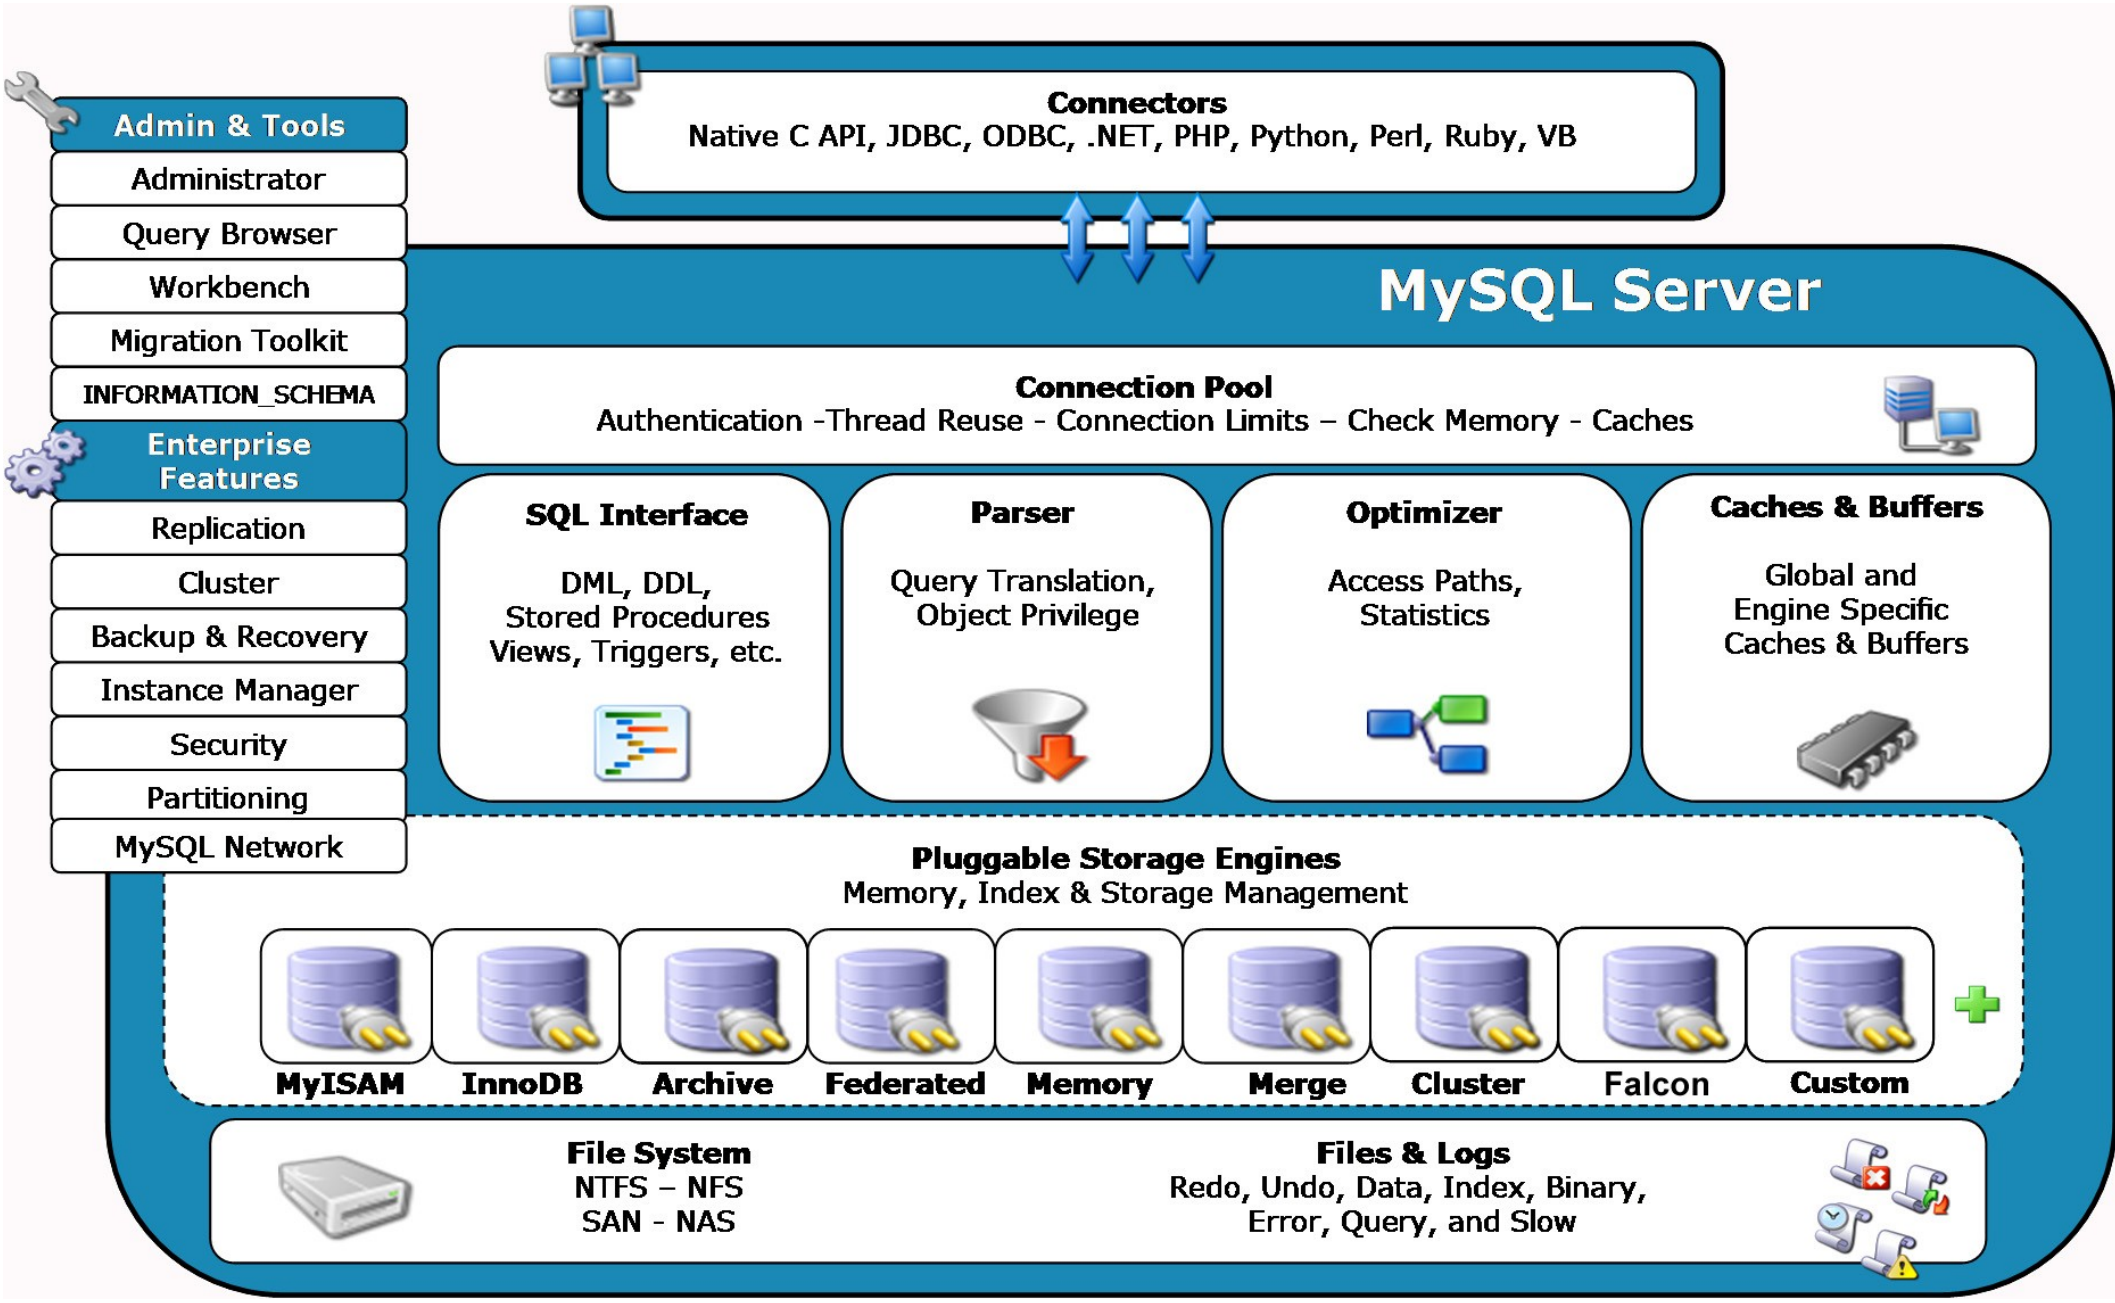

Storage Engines

- InnoDB
  - ACID, Transactions
- MyISAM
  - No Transact, Small footp
- Archive
  - Only Insert & Select
- Memory
  - In memory (temp table)
- NDB/Cluster
  - Failsafe distributed mostly in memory
- Falcon
  - ACID, opt for 64 bit
- Maria (MyISAM++)
  - Crash recovery, DW stuff
- PBXT (Streaming BLOB)
- Solid
- NitroSecurity
- InfoBright
- ScaleDB
- IBM DB2 (for i5 platform)
- Blackhole (/dev/null)

MySQL Cluster (NDB)

- Distributed in memory storage engine with
  - **Fault Tolerance:** shared nothing architecture
  - **High Availability:** fast auto failover (five 9's of availability)
  - **Scalability:** Scale by adding more commodity machines
  - **High Performance:** Really really fast (many 100000 ops per second) for primary key lookups (mixed read and write). Up to millions of queries per second using low level C API and high end hardware
  - **Simplified applications:** For the application MySQL Cluster is a just a table. In the case of failure you reconnect to another MySQL Server and immediately see the same data
  - Currently slow on joins so no replacement for MyISAM/InnoDB

MySQL 5.0 (Current GA version)

- Stored Procedures
- Triggers
- Views
- Strict Mode (classical DB error handling)
- Information_Schema (Data Dictionary)
- Precision Math (> 50 digits of precision)
- Many additions to our optimizer -> Faster complex queries
- Cursors (read only, forward scrolling)
- XA (distributed transactions)

MySQL 5.0: Stored Procedures

- MySQL's stored procedures:
  - Follow the SQL:2003 syntax
  - Can return multiple result sets in a single invocation
  - Supports `INOUT` parameters, partial exception handling, and flow control

Sample:

```
mysql> delimiter //
mysql> CREATE PROCEDURE simpleproc (OUT param1 INT)
    -> BEGIN
    ->   SELECT COUNT(*) INTO param1 FROM t;
    -> END
    -> //
mysql> delimiter ;
mysql> CALL simpleproc(@a);
mysql> SELECT @a;
+------+
| @a   |
+------+
| 3    |
+------+
```

MySQL 5.0: Triggers

- MySQL includes support for triggering statements and stored procedures before or after a table event

Sample (Keep the old value when updating):

```
CREATE TABLE account (acct_num INT, amount DECIMAL(10,2));
CREATE TRIGGER ins_sum BEFORE UPDATE ON account
   -> FOR EACH ROW NEW.previous_amount = OLD.amount;
```

MySQL 5.0: Views

- MySQL supports updatable and read-only views:
  - Useful for accessing the result of a query as if it were a table
  - Can restrict access to a set of rows in a table, database, or view

Sample:

```
CREATE TABLE t (qty INT, price INT);
INSERT INTO t VALUES(3, 50);
CREATE VIEW v AS SELECT qty, price, qty*price AS value FROM t;
SELECT * FROM v;
+------+-------+-------+
| qty  | price | value |
+------+-------+-------+
|    3 |    50 |   150 |
+------+-------+-------+
```

MySQL 5.0: Data Dictionary

- Supports for a new information_schema database with meta information
  - Useful for finding any meta information about your data, in addition to the present SHOW command

Sample:

```
use information_schema
select TABLE_SCHEMA, COLUMN_NAME, CHARACTER_SET_NAME from COLUMNS
      where TABLE_NAME = "tables" limit 1;
+--------------------+--------------+--------------------+
| Table_schema       | COLUMN_NAME  | CHARACTER_SET_NAME |
+--------------------+--------------+--------------------+
| information_schema | TABLE_CATALOG | utf8               |
+--------------------+--------------+--------------------+
1 row in set (0.05 sec)
```

MySQL 5.0: Precision Math

- Exact calculations with
  - Well defined rounding
  - At least 56 digits precision
  - Very fast with static memory allocation

Sample:

```
create table d2 (n decimal(64,3));
insert into d2 values (233221213212312312321321321321),
                      (3454354332432543543 5),
                      (32432432432454374435234543456);
Query OK, 3 rows affected (0.00 sec)
select sum(n) from d2;
+-----------------------------------+
| sum(n)                            |
+-----------------------------------+
| 265653680188310011081991300212.000 |
+-----------------------------------+
1 row in set (0.00 sec)
```

MySQL 5.0: Strict Mode

- Get rollback/errors instead of closest value/warning
  - Was one of the most asked for features for 10 years!

```
Sample:
CREATE TABLE d1 (d date);
Query OK, 0 rows affected (0.23 sec)

mysql> INSERT INTO d1 SET d = "2005-04-31";
Query OK, 1 row affected, 1 warning (0.00 sec)

mysql> SET sql_mode='STRICT_ALL_TABLES';
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO d1 SET d = "2004-04-31";
ERROR 1292 (22007): Incorrect date value: '2004-04-31' for column 'd' at row 1
```

MySQL 5.0: Server Side Cursors

- Read only, forward scrollable cursors

Sample:

```
CREATE PROCEDURE cursor_demo()
BEGIN
  DECLARE a, b CHAR(16);
  DECLARE cur1 CURSOR FOR SELECT id,data FROM test.t1;
  OPEN cur1;
  REPEAT
    FETCH cur1 INTO a, b;
    ...

  CLOSE cur1;
END
```

MySQL 5.0: XA

- MySQL can Participate as a Resource Manager (RM)
  - Commit grouping to improve performance
  - Java support via XADatasource implementation
- MySQL can be a Transaction Manager
  - Coordinates transactions across multiple storage engines
  - XA is part of storage engine interface

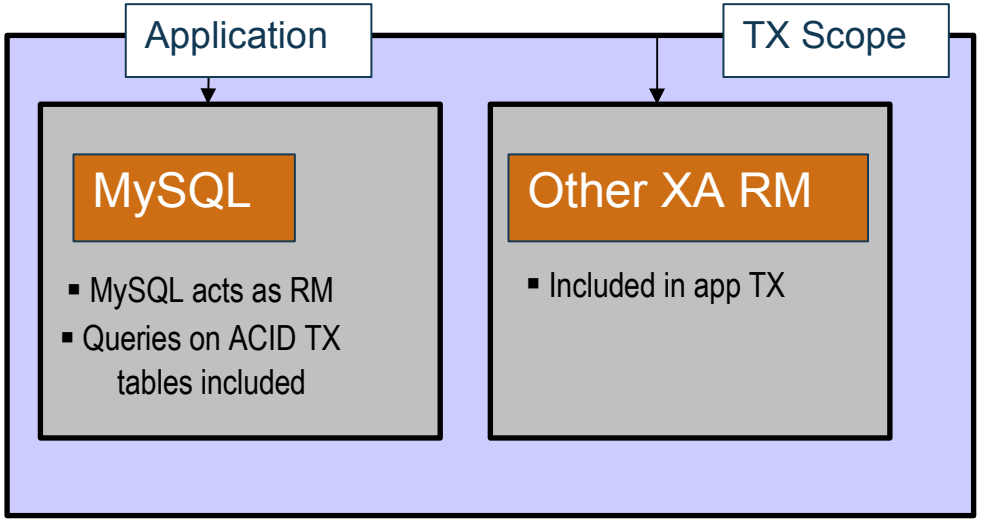

 Creators of MySQL

5.0: Even Better Optimizer

- Greedy Optimizer (Fast handling of **many** tables in a join)
- Index Merge (for `WHERE a=1 or b=2`)
- Equality propagation (WHERE a=b AND a=c -> implies b=c)
- Conversion of outer joins into inner joins. (Allows us to consider different join orders)
- Correct execution of LEFT and RIGHT outer joins
- Loose index scan for queries of the type
  - SELECT MAX(b) FROM t GROUP BY a
  - SELECT DISTINCT a, b, c FROM t

5.0: 'Small' things also gets added!

Extension to LOAD DATA for doing transformation/calculation at the time you load the data.

```
LOAD DATA INFILE 'file.txt' INTO TABLE t1
   (col1, @var1, @var2)
  SET col2 = @var1-@var2, col3 = @var2;
```


MySQL 5.1 (now in late Beta)

- A better 5.0 (Thousands of small fixes)
- Table/Index Partitioning - for data warehousing etc
- Row-based replication – Transfers data instead of commands
- Full-text indexing parser plugins – Flexible full text search
- Disk-based Data Support for MySQL Cluster
- Replication Support for MySQL Cluster – Running with multiple clusters in different locations
- XPath Support - helps any customer wanting to better navigate and search XML documents stored in MySQL
- Internal Task Scheduler (Events)- new utility that helps you to schedule, run, and monitor database-driven tasks and jobs
- And much more...

Upcoming Features (6.0, 6.1)

- Partitioning
  - Parallel query
- Replication
  - Multi source replication
- Global Backup API
- Hash & Merge joins
- Falcon and Maria (MyISAM++) storage engine
- Federated tables over ODBC
- Foreign key support for all engines

Development Road Map

- More plugins (Storage engines are already a plugin)
- More online features (ALTER TABLE)
- Features to run Enterprise Software (SAP etc)
- Extending current features (XPATH, Subquery optimizer, geometrical data etc...)
- Profiling, stored procedure debugging...
- SQL standards
- Useful extensions
  - Data warehousing (New storage engines, new index types)
  - Store 20 Peta (20,000 Terabytes) bytes in MySQL

MySQL Graphical (GUI) Tools

- **Available Now!**
  - MySQL Administrator (Win, Linux, MacOS)
  - MySQL Query Browser (Win, Linux, MacOS)
  - MySQL Migration Toolkit (Win, Linux, MacOS)
    - Plug-in Architecture for sources (Java/JDBC)
- **In Development**
  - MySQL Workbench
- All freely available under GPL license


MySQL Administrator


MySQL Query Browser


MySQL Migration Tool


MySQL Workbench

Contributions

- The most common contribution is repeatable bug reports & ideas!
- MySQL has always accepted contributions
  - Windows port, JDBC, .Net, PHP, Perl connector...
- Not only the server: GUI Tools, Connectors, Articles, Documentation...
- HOWTO:
  - Subscribe to the internals@lists.mysql.com mailing list
  - Check out forge.mysql.com/contribute
  - Work from the latest development version
    - There is a read only BitKeeper/SVN tree available with the latest
  - Shared copyright needed (Just like the FSF, OpenOffice)

Free Databases get Better all the time!

- Good bug reports since bugs gets fixed
  - **Repeatable bug reports are as valuable as code!**
- Much faster user/developer feedback than closed source
- Lots of testing of **all** code. All features available for all!
- Freedom & Independence!
  - You have the ultimate documentation, the source!
- Security is not by obscurity, No hidden hooks in the code
- Lots of Eco system code gets written by the community
- We can hire people who has shown they already know the code/product
- Result: **Low Total Cost of Ownership**

All software discussed here is OSS

- So you can go and download the source for Cluster, MySQL server, GUI tools and connectors directly!
- The piece that is not freely downloadable is MySQL Enterprise that is a pay per year per machine offering (a so called subscription).
  - Very useful for companies who want to get their DBAs more efficient.

No Software Patents!

- Software Patents are a threat against Free Software, Software innovation and developing local software industries countries
- In Europe our side was successful and the SW Pat proposal was thrown out (a real thriller!)
- MySQL has been spending lots of cash and lots of Management time (CEO,VPs & Founders) fighting Software Patents
  - Other backers included RedHat
- Please help make sure that Romania does not follow the US into this minefield.
- See recent article “Microsoft takes on the free world”

We are looking for developers?

- MySQL AB are looking for good C/C++ developers
  - We strongly prefer people with OpenSource experience
  - People who send in a good CV has < 1% chance
  - People who send in good patches and bug fixes has close to 100% chance historically
  - BTW: we prefer CVs that does not list every computer buzz world for the last 30 years :-(
- C/C++ Developers with system programming experience. Filesystem work, Kernel internals (VM System), Database internals etc

Questions?