

# **MySQL: New features in Version 4.1, 5.0, 5.1**

**Peter Zaitsev**  
**High Performance Group, Manager**

**Open Source Forum 2005, Moscow, Russia**  
**April 27-29, 2005**

**MySQL AB**

## About Speaker

- Peter Zaitsev, MySQL Inc.
  - Senior Performance Engineer
  - MySQL Performance Group Manager
  - MySQL Performance consulting and partner relationships
- Previously
  - Co-Founder, CTO of SpyLOG.RU
  - Graduated from Computer Science Department, Moscow State University

## The World's Most Popular Open Source Database

- Founded in 1995; operations in 27 countries
- Over 6,000,000 installations; 40,000 downloads / day
- Part of a rapidly growing open source stack
- Dramatically reduces Total Cost of Ownership (TCO)
- Used by leading IT organizations and ISVs

The SIEMENS logo, consisting of the word "SIEMENS" in white capital letters on a yellow rectangular background.The Travelocity logo, featuring a stylized line graph with a red arrow pointing up and the word "Travelocity" in black.The GAP logo, consisting of the word "GAP" in white capital letters on a dark blue square background.The CISCO SYSTEMS logo, featuring the words "CISCO SYSTEMS" in red above a green bar chart.The NOKIA logo, consisting of the word "NOKIA" in blue capital letters above the tagline "CONNECTING PEOPLE" in smaller blue capital letters.The Google logo, with the word "Google" in its multi-colored font.The YAHOO! logo, with the word "YAHOO!" in red capital letters.The Novell logo, with the word "Novell" in red capital letters.The Sun microsystems logo, featuring the word "Sun" in white on a purple background with "microsystems" in smaller white text below.The SAP logo, consisting of the word "SAP" in white capital letters on a dark blue background.The hp invent logo, featuring the "hp" logo in white on a blue square background with the word "invent" in lowercase below.

# MySQL Features

- Very easy to use
  - 15 minutes install rule
- Good quality track record
- High Performance for variety of workloads
- Multiple storage engines can be used at the same time
  - You do not need transactions for your log, but for your billing
- Small disk and memory foot print
  - Same code scales from hand-held devices to highend servers
- Great documentation (available in Russian)
- Connectors for C, ODBC, Java, PHP, Perl, .NET, ...
- Wide range of supported platforms
- Great community and Commercial Support

## State of MySQL 4.0

- MySQL 4.0 - older production release
- Subset of SQL 92
- Wide set of supported data types
- Multiple Storage Engines
- InnoDB storage Engine
  - Transactions, Foreign keys, row level locks, versioning
- Full text search support
- Connectors for most of programming languages
- Some character set support
  - No Unicode support, single charset for server
- Lacking
  - Subselects, Prepared Statements, Stored Procedures, Views..

## MySQL 4.1

- Current production release, since October 2004
- More Secure authentication
- Non Blocking Key Cache
  - With Index Preloading, Multiple Key Cache support
- Improved SSL support, can be used with Replication
- Warnings
- Built in HELP function
- GIS (Geographical Data) support for MyISAM tables
- Optimizer improvements
  - faster sorts, group by
- Innodb:
  - Multiple Tablespace support – each table in its own file.

## MySQL 4.1 SubQueries

- **SELECT \* FROM t1 WHERE column1 = (SELECT column1 FROM t2);**
- Dependent subqueries supported
- Can be used in any SQL Statement (UPDATE etc)
- **IN,ANY/SOME, EXISTS** operators
- Inline Views/Derived Tables
  - **SELECT AVG(sum\_column1) FROM (SELECT SUM(column1) AS sum\_column1 FROM t1 GROUP BY column1) AS t1;**
- Support for ROW subqueries
  - **SELECT \* FROM t1 WHERE ROW(1,2) = (SELECT column1, column2 FROM t2);**
- Limits:
  - Some cases are not optimized yet, faster if rewritten as joins
  - No LIMIT support in subqueries, supported in derived tables

# Prepared Statements

- “Standard” prepare/execute type of interface
- Save on parsing
- Significantly reduce SQL Injection problem
- Comes with binary protocol – less data conversion
- Connectors: C API, JDBC, .NET, PHP 5.0 (mysqli)
  - Perl DBI support available in CVS version
- Limits
  - Not all statements can currently be prepared
  - For single execution standard statements can be faster
    - Single trip to the server
  - Placeholders are not allowed in some places, e.g. LIMIT

# Localisation

- Specify character set on database, table, column
  - **CREATE TABLE t (c VARCHAR(255)) CHARSET=utf8;**
- List available charsets by **SHOW CHARACTER SET**
- Multiple collations (sort orders)
  - **c VARCHAR(255) COLLATE utf8\_icelandic\_ci**
- Metadata stored in utf8
  - Watch on upgrade if you have Russian table names
  - Some size and performance overhead
- Multiple timezone support
  - Different connections may use different timezones
- Things to watch
  - utf8 takes more space than latin1/cp1251 can be slower
  - Different collations can't be compared

## MySQL 4.1: NDB Cluster

- Available in MySQL 4.1-Max
- In memory clustered storage engine
- Shared nothing architecture
- No single point of failure
- Automatic hash based data partitioning
- Database has to fit in memory of the cluster
- Usable for special applications mostly at this point
  - Telecom application, metadata storage, user database
  - slow joins, timeout based deadlock detection, tricky extension
- Support for on disk data is on a way

# MySQL 5.0 Stored Procedures

- MySQL's stored procedures:
  - Follow the SQL:2003 syntax
  - Can return multiple result sets in a single invocation
  - Support `INOUT` parameters, 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 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 - Triggers

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

Sample (sum the values of newly inserted rows):

```
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: Information Schema

- Supports for a new information\_schema database with meta data information
  - Useful for finding any meta information about your data, in addition to the present SHOW command
  - Inefficient in case of many tables

### 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 Strict Mode

- Ability to get rollback/errors instead of closest value / warning message
- ANSI SQL 2003 style behavior

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 Precision Math

- Exact calculations with
  - Well defined rounding
  - At least 56 digits precision
  - Very fast with static memory allocation
- Must have for financial applications

### Sample:

```
create table d2 (n decimal(64,3));
insert into d2 values (233221213212312312321321321321),
                    (34543543324325435435435),
                    (32432432432454374435234543456);
```

Query OK, 3 rows affected (0.00 sec)

```
select sum(n) from d2;
```

```
+-----+
| sum(n) |
+-----+
| 531307360376620022163982600424.000 |
+-----+
```

1 row in set (0.01 sec)

## MySQL 5.0 More features

- Cursors (Read-only, forward only)
- Improved MySQL Cluster performance
- True Varchar
  - End spaces no more truncated
- XA – Distributed transactions
  - Both Resource Manager and Transaction Manager roles
  - ACID multiple storage engine transactions
- Improved optimizer
  - Index merge – using several indexes per table instance
  - “Greedy optimizer” - more efficient plan search
  - Equality propagation, nested joins etc

## MySQL 5.0 More features

- Federated storage engine
  - See remote table (MySQL) as local table
  - Very simple implementation at first
  - Support for other DBMS to be added later
- Instance Manager
  - replaces mysqld\_safe script. Monitors and controls the instance
- BIT Type (Compact storage)
- Innodb improvements
  - Compact storage format (some 20% smaller tables)
  - Background transaction rollback on crash recovery

## MySQL 5.1 Preview

- Early in development cycle features may change
- Partitioning and parallel query
  - Storage engine independent
- Improved replication
  - Row level replication (safer, works with all features)
  - Multi Source replication
- Online Backup and Backup API
- Hash and Merge Joins
  - Data Warehouse applications with large data sets
- XML type and XPath expressions
- Performance improvements of existing features

## Yes we're hiring

- MySQL Company Is rapidly growing
  - 3 years ago I was #45, now some 180 employees
- About 20 employees in Russia and Ukraine
- Developers typically working in their home offices
  - Hiring good talent everywhere
- MySQL mostly hires engineers in Russia
  - Server kernel developers
  - GUI Developers
  - Quality Assurance Engineers
  - GUI Developers
  - Support Engineers
- Contribution – Code, Bug reports, Ideas, Articles, Community support is best way to be hired
  - Even if you're not hired, being visible in OpenSource community is rewarding

## Resources

- MySQL Web site <http://www.mysql.com>
- MySQL Discussion forums <http://forums.mysql.com>
- MySQL Mailing lists: <http://lists.mysql.com>
- MySQL Manual:  
<http://dev.mysql.com/doc/mysql/en/index.html>
- Contact me: [peter@mysql.com](mailto:peter@mysql.com)