

New Features in MySQL 5.0, 5.1, and Beyond

Jim Winstead
`jimw@mysql.com`

**Southern California Linux Expo
February 2006**

MySQL AB

5.0: GA on 19 October 2005

- Expanded SQL standard support:
 - Stored procedures and functions
 - Triggers
 - Views
 - INFORMATION_SCHEMA
 - True VARCHAR
 - True DECIMAL (precision math)
 - SQL Mode
- XA (Distributed transactions)
- Federated and archive storage engines

Stored procedures and functions

```
mysql> DELIMITER //
```

```
mysql> CREATE PROCEDURE simpleproc (OUT param1 INT)
-> BEGIN
->   SELECT COUNT(*) INTO param1 FROM t;
-> END//
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> CALL simpleproc(@a) //
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> SELECT @a//
```

```
+-----+
| @a    |
+-----+
| 3     |
+-----+
```

1 row in set (0.00 sec)

Triggers

```
mysql> delimiter //
mysql> CREATE TRIGGER upd_check
-> BEFORE UPDATE ON account
-> FOR EACH ROW
-> BEGIN
->     IF NEW.amount < 0 THEN
->         SET NEW.amount = 0;
->     ELSEIF NEW.amount > 100 THEN
->         SET NEW.amount = 100;
->     END IF;
-> END; //
```

Views

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

qty	price	value
3	50	150

INFORMATION_SCHEMA

```
mysql> SELECT table_name, table_type, engine
-> FROM information_schema.tables
-> WHERE table_schema = 'db5'
-> ORDER BY table_name DESC;
```

table_name	table_type	engine
v	VIEW	NULL
t2	BASE TABLE	MyISAM
t	BASE TABLE	MyISAM
pk	BASE TABLE	InnoDB
fk	BASE TABLE	InnoDB

5 rows in set (0.01 sec)

True VARCHAR and True DECIMAL

- VARCHAR
 - Preserves trailing spaces
 - Length can be up to 65,535 characters (but only 65,532 effectively)
- DECIMAL
 - Calculations done using new fixed-point precision math library, with a precision of up to 65 digits

SQL mode

```
mysql> SET sql_mode = 'traditional';  
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> INSERT INTO t1  
-> SET date_field = '2006-2-29';  
ERROR 1292 (22007): Incorrect date value:  
'2006-2-29' for column 'date_field' at row 1
```


New storage engines

- **FEDERATED**
 - Allows seamless access to a table on another MySQL server
 - Transaction support added in 5.1
 - Will be extended for heterogeneous federation in a later release
- **ARCHIVE**
 - No indexes (except auto_increment in 5.1)
 - Does zlib compression of data
 - Great for data warehousing
 - Fast BLOB skipping in 5.1
- **BLACKHOLE**
 - Useful in some replication setups
 - Supports transactions in 5.1

5.1: GA in Q2 2006

- Partitioning
- Row-based replication
- Cluster disk data
- Cluster replication
- Task scheduler
- Logs as tables
- Fast ALTER TABLE
- Plugin API
- XML (XPath)

Partitioning

- Allows for splitting data across different tables (which could be located on different disks)
- Supported by all storage engines (some natively, like Cluster)
- Partition pruning done by optimizer
- No parallelism (yet)
- Can partition by:
 - Range
 - Hash
 - Key
 - List
- Subpartitioning supported

Partition example


```
mysql> CREATE TABLE list_by_area (  
->   store_number    INT           NOT NULL,  
->   location        INT           NOT NULL,  
->   rollup_date     DATE          NOT NULL,  
->   store_receipts DECIMAL(10,2) NOT NULL)  
-> PARTITION BY LIST(location) (  
->   PARTITION p1 VALUES IN (1,2) ,  
->   PARTITION p2 VALUES IN (3) ,  
->   PARTITION p3 VALUES IN (4,5)  
-> );
```

Row-based replication

- Start server with `--binlog-format=row`
- Instead of logging statements, the binary log will instead contain images of the rows that are added, changed, or deleted
- DDL statements (like **ALTER TABLE**) are still logged as statements

New cluster features

- Disk-based data
 - Indexes must still fit in memory
- Variable-sized records (**VARCHAR** only takes up as much space/memory as necessary)
- Cluster replication



Task scheduler

```
mysql> delimiter //  
mysql> CREATE EVENT optimize_tables  
-> ON SCHEDULE EVERY 1 WEEK  
-> STARTS '2006-03-05 1:00:00' /*Sunday*/  
-> ON COMPLETION PRESERVE  
-> DO  
-> BEGIN  
->     OPTIMIZE TABLE test.table1;  
->     OPTIMIZE TABLE test.table2;  
-> END//
```

Plugin API

- New framework for extending server functionality
- First plugins are for FULLTEXT indexing
 - Work sponsored by CNET
- Future areas for plugins:
 - Storage engines
 - Authentication and authorization
 - ?

XML (XPath)

```
mysql> SELECT
->   ExtractValue ('<a>ccc<b>ddd</b></a>', '/a') AS v1,
->   ExtractValue ('<a>ccc<b>ddd</b></a>', '/a/b') AS v2;
```

```
+-----+-----+
| v1    | v2    |
+-----+-----+
| ccc   | ddd   |
+-----+-----+
```

1 rows in set (0.01 sec)

```
mysql> SELECT
->   UpdateXML ('<a><b>ccc</b><d></d></a>', '/a',
->             '<e>fff</e>') AS v1,
->   UpdateXML ('<a><b>ccc</b><d></d></a>', '/b',
->             '<e>fff</e>') AS v2 \G
```

```
***** 1. row *****
v1: <e>fff</e>
v2: <a><b>ccc</b><d></d></a>
1 rows in set (0.01 sec)
```

...and Beyond

- 5.2 will (probably) include ROLES and pluggable authentication/authorization support
 - I'm (one of the people) writing it
- And a lot more that I can't talk about.

Questions?

Contact:

Jim Winstead
jimw@mysql.com