What you need to know to upgrade from MySQL 5.0 to MySQL 5.1

Sheeri K. Cabral
Database Administrator
The Pythian Group, www.pythian.com

History of MySQL 5.1

- First Nov 29 2005
- RC Sept 24 2007
- GA Nov 14 2008

New in 5.1

- Partitioning
- Row-based replication
- Plugin API

New in 5.1 (part 2)

- Event Scheduler
- Log tables
- INFORMATION_SCHEMA
  - PROCESSLIST
  - ENGINES

New in 5.1 (part 3)

- TRIGGER privilege
- Prepared statements can use query cache
- XPath support
  - ExtractValue
  - UpdateXML

New Tools

- mysqlslap
- mysql_upgrade
  - mysqlcheck
    - --fix-db-names, --fix-table-names
    - --check-upgrade
  - CHECK TABLE FOR UPGRADE

New Tool Functionality

- mysqldump
  - --replace
  - --single-transaction does not use BEGIN
  - --master-data changes
- mysqlbinlog
  - --server-id
  - --base64-output
  - --dump-date, --skip-dump-date
  - --verbose, -v

New Tool Functionality (part II)

- mysqlimport
  - --use-threads=N
- mysql.server
  - --service-startup-timout
- mysqladmin
  - --no-beep

Storage Engine Changes

- BDB
- FEDERATED
- have_isam

Incompatible Changes

- mysqld_safe only checks for mysqld
  - not msyqld-max
- system variables and binary logs
- Change in handling prepared statements; may need to upgrade client library

Stored routines, triggers, views

- New fields in mysql.proc table, created NULL
  - character_set_client
  - collation_connection
  - db_collation
  - body_utf8

mysqldump --default-character-set=[definition charset]

Configuration File Changes

- table_cache is now table_open_cache
- tmp_table_size default changed
- max_connections default changed

Configuration File Changes (part II)

- have_raid is removed
- --master-% deprecated
- --one-thread gone, --thread-handling added

Configuration File Changes (part III)

- hostname
- Error logging
  - --syslog, --skip-syslog
  - --syslog-tag=tag
- --innodb_log_arch_dir removed

New Parameters

- table_definition_cache
- myisam_use_mmap
- binlog_format, binlog_row_event_max_size

New Parameters (part II)

- max_prepared_stmt_count
- Prepared_stmt_count (global status variable)
- innodb_stats_on_metadata

New Parameters (part III)

- old
- --skip-write-binlog
- plugins
  - --loose-x

New Parameters (part IV)

- --slave-exec-mode
  - STRICT, IDEMPOTENT
- Uptime_since_flush_status
- --report-%

New Features

- Meaningful stack traces

- charset on mysql commandline

New Features

- SHOW AUTHORS
- SHOW PROFILE, SHOW PROFILES
  - INFORMATION_SCHEMA.PROFILING
- SHOW CONTRIBUTORS

Important Bugfixes

- BLACKHOLE transactions
- FULLTEXT search and apostrophes
- INFORMATION_SCHEMA uses less memory
- --replicate-%-table options

New Functionalities

- Faster ALTER TABLE for some commands
- RAND() can take non-constant initializers
- ARCHIVE supports AUTO_INCREMENT

New Functionalities (part II)

- MySQL now escapes identifiers
- FULLTEXT search
  - IN NATURAL LANGUAGE MODE
  - .... WITH QUERY EXPANSION
- Stored routines allow DEFINER

New Functionalities (part III)

- CHANGE MASTER TO
  - MASTER_SSL_VERIFY_SERVER_CERT = {0|1}
- PAD_CHAR_TO_FULL_LENGTH
- long_query_time
  - microseconds
  - 0

New Functionalities (part IV)

- --log-slow-slave-statements
- --min_examined_row_limit
- --innodb_autoinc_lock_mode

Deprecated Features

- SHOW INNODB STATUS
- SHOW MUTEX STATUS
- SHOW [BDB] LOGS

Deprecated Features (part II)

- TYPE as a synonym for ENGINE
  - @@table_type
  - SHOW TABLE TYPES
- --skip-thread-priority
- --log, --log-slow-queries

Package Changes

- mysqld-debug included
- safe_mysqld excluded
- mysqldumpslow part of server, not client
- MySQL Cluster is a separate product

New Collations

- cp1250_polish_ci
- utf8_hungarian_ci
- ucs2_hungarian_ci

Sheeri Cabral

cabral@pythian.com

www.sheeri.com

What you need to know to upgrade from MySQL 5.0 to MySQL 5.1

Sheeri K. Cabral
Database Administrator
The Pythian Group, www.pythian.com

History of MySQL 5.1

- First Nov 29 2005
- RC Sept 24 2007
- GA Nov 14 2008

2 years of development, 1 year from RC to GA.

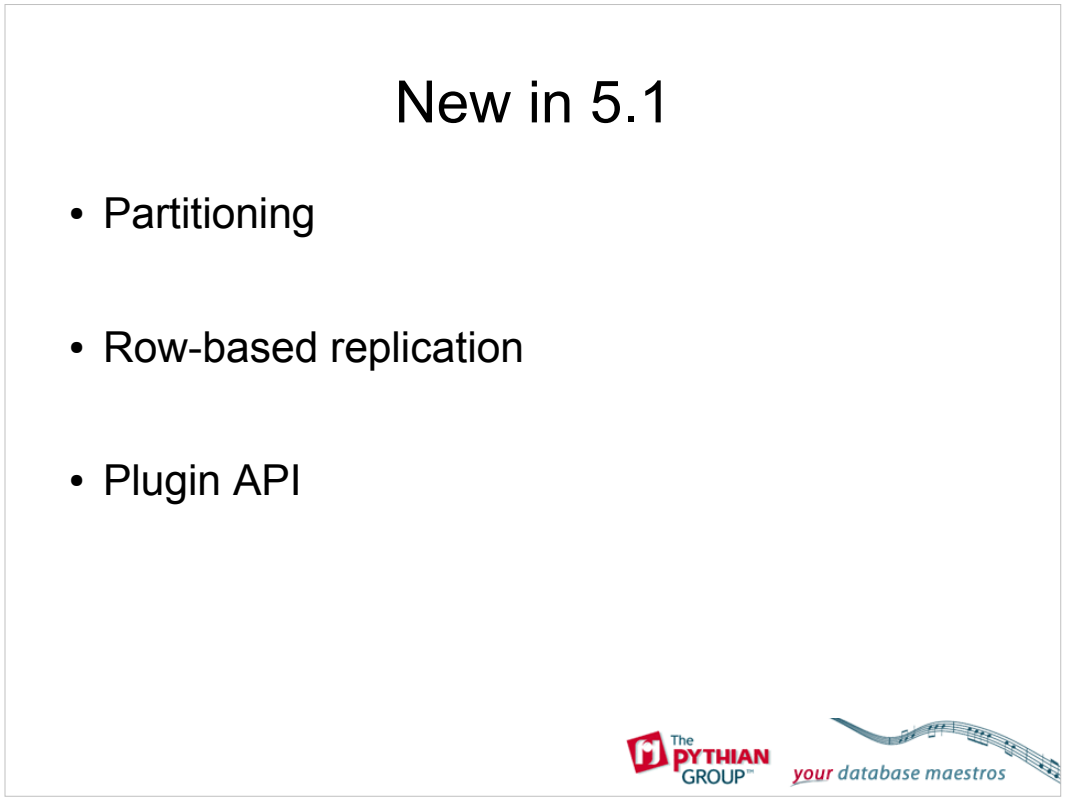

Partitioning – automatic merge tables. Different parts of the table are stored in different file locations, but are queried like a single table.

Row based replication – currently, replication happens like this: MySQL writes statement to binlogs, slave reads binlogs, applies statement.

Plugin API – load and unload runtime components. Ie, plugin fulltext parser, though it's not finished yet – basically you can implement your own input filter on indexed text, enabling fulltext search on data like PDF or other docs. Pre-parser fulltext plugin you write, then hands it over to MySQL fulltext search, so still needs to be MyISAM, etc.

New in 5.1 (part 2)

- Event Scheduler
- Log tables
- INFORMATION_SCHEMA
  - PROCESSLIST
  - ENGINES

Similar to cron, the event scheduler has scheduled events. Puts notes in the MySQL error log. Now you can use this instead of running an SQL query every so often, saving passwords in a cron script.
allows one to schedule statements for execution at predetermined times. transient (one-time-only) or recurrent at regular intervals, may execute queries & stmts. same as allowed in stored routines, including compound statements. Events can be altered after creation, and dropped when no longer needed.
SHOW EVENTS and SHOW CREATE EVENT, INFORMATION_SCHEMA.EVENTS table.
EVENT privilege

- Log to file or table or both with log_output.
- INFORMATION SCHEMA has new tables including FILES, EVENTS, PARTITIONS, PROCESSLIST, ENGINES, PLUGINS

New in 5.1 (part 3)

- TRIGGER privilege
- Prepared statements can use query cache
- XPath support
  - ExtractValue
  - UpdateXML

Prep stmt and query cache http://bugs.mysql.com/735

. ExtractValue() returns the content of a fragment of XML matching a given XPath expression.
UpdateXML() replaces the element selected from a fragment of XML by an XPath expression supplied by the user with a second XML fragment (also user-supplied), and returns the modified XML.
This release introduces the TRIGGER privilege. Previously, the SUPER privilege was needed to create or drop triggers. However, the requirement that the account named in a trigger's DEFINER clause must have the SUPER privilege has changed to a requirement for the TRIGGER privilege. After upgrading, be sure to update your grant tables with “mysql_fix_privilege_tables This process assigns the TRIGGER privilege to all accounts that had the SUPER privilege.

New Tools

- mysqlslap
- mysql_upgrade
  - mysqlcheck
    - --fix-db-names, --fix-table-names
    - --check-upgrade
  - CHECK TABLE FOR UPGRADE

myqsl_upgrade checks all existing tables for incompatibilities and repairs if necessary (for instance, table version in information_schema).

Added the --check-upgrade to mysqlcheck that invokes CHECK TABLE with the FOR UPGRADE option. Added the --fix-db-names and --fix-table-names options to mysqlcheck.

mysqlslap emulates client load

New Tool Functionality

- mysqldump
  - --replace
  - --single-transaction does not use BEGIN
  - --master-data changes
- mysqlbinlog
  - --server-id
  - --base64-output
  - --dump-date, --skip-dump-date
  - --verbose, -v

mysqldump can now dump NDB tablespaces. It can also use –replace, to do REPLACE INTO instead of INSERT INTO. W00t!

Uses START TRANSACTION /*!40100 WITH CONSISTENT SNAPSHOT */

mysqldump adds the LOCAL qualifier to the FLUSH TABLES statement that is sent to the server when the --master-data option is enabled. This prevents FLUSH TABLES stmt from replicating to slaves, which is disadvantageous because causes slaves to block while statement executes. (Bug#35157)

.....produces a -- Dump completed on DATE comment at the end of the dump if --comments is given. The date causes dump files for identical data take at different times to appear to be different. --dump-date and --skip-dump-date control whether the date is added to the comment. The default is --dump-date (include date). (Bug#31077)

mysqlbinlog --server-id enable only those events created by the server having the given server ID to be extracted. **base64-output** prints the entries using base64 encoding, for debugging only, do not apply these logs to production systems.

--verbose = will print out statement too in comment, see http://dev.mysql.com/doc/refman/5.1/en/mysqlbinlog-row-events.html

New Tool Functionality (part II)

- mysqlimport
  - --use-threads=N
- mysql.server
  - --service-startup-timout
- mysqladmin
  - --no-beep

mysqlimport now has a --use-threads=N option for loading data files in parallel using N threads.

mysqladmin -- no warning beep. :)

mysql.server – how long to wait for the server to start.
http://bugs.mysql.com/26952

Storage Engine Changes

- BDB
- FEDERATED
- have_isam

BDB is unsupported in 5.1 -- if you have skip-bdb in your my.cnf, you will need to change that.

Federated is not enabled by default in 5.1 -- if you want it, put federated in your my.cnf

have_isam is removed

Incompatible Changes

- mysqld_safe only checks for mysqld
  - not msyqld-max
- system variables and binary logs
- Change in handling prepared statements; may need to upgrade client library

mysqld_safe no longer checks for a mysqld-max binary. Instead, mysqld_safe nows checks only for the standard mysqld server unless another server binary is specified explicitly via --mysqld or --mysqld-version. If you previously relied on the implicit invocation of mysqld-max, you should use an appropriate option now. (Bug#17861)

statements that refer to a system variable are marked as “unsafe” -- just a warning in STATEMENT mode, uses RBR in MIXED mode

http://dev.mysql.com/doc/refman/5.1/en/news-5-1-25.html

Stored routines, triggers, views

- New fields in mysql.proc table, created NULL
  - character_set_client
  - collation_connection
  - db_collation
  - body_utf8

mysqldump --default-character-set=[definition charset]

Configuration File Changes

- table_cache is now table_open_cache
- tmp_table_size default changed
- max_connections default changed

The default value of tmp_table_size (32M) is meaningless, as it is less than max_heap_table_size (16M), which then becomes the real maximum for an in-memory temporary table.

now default is 16M

The default value of the max_connections variable has been increased to 151 in order that Websites running on Apache and using MySQL will not have more processes trying to access MySQL than the default number of connections available.
(The maximum number of Apache processes is determined by the Apache MaxClient, which defaults to 256, but is usually set to 150 in the httpd.conf (Bug#23883)

Configuration File Changes (part II)

- have_raid is removed
- --master-% deprecated
- --one-thread gone, --thread-handling added

he following options for controlling replication master configuration on a slave are now deprecated. Use CHANGE MASTER TO instead

--master-host --master-user
--master-password --master-port
--master-connect-retry
--master-ssl --master-ssl-ca
--master-ssl-capath --master-ssl-cert
--master-ssl-cipher --master-ssl-key

Added the thread_handling system variable to control whether the server use a single thread or one thread per connection. The --one-thread option now is deprecated; use --thread_handling=one-thread instead.

Configuration File Changes (part III)

- hostname
- Error logging
  - --syslog, --skip-syslog
  - --syslog-tag=tag
- --innodb_log_arch_dir removed

Added the hostname system variable, the server sets this variable to the server hostname at startup.

mysqld_safe now supports error logging to syslog on systems that support the logger command. The new --syslog and --skip-syslog options can be used instead of the --log-error option to control logging behavior, as described in Section 4.3.2, “mysqld_safe — MySQL Server Startup Script”. The default is to use –skip-syslog

--syslog-tag=tag modifies the default tags written by mysqld_safe and mysqld to syslog to be mysqld_safe-tag and mysqld-tag rather than the default tags of mysqld_safe and mysqld.

innodb_log_arch_dir deprecated since 5.0.24

New Parameters

- table_definition_cache
- myisam_use_mmap
- binlog_format, binlog_row_event_max_size

table_definition_cache: The number of table definitions that can be stored in the definition cache. If you use a large number of tables, you can create a large table definition cache to speed up opening of tables. The table definition cache takes less space and does not use file descriptors, unlike the normal table cache....The minimum and default are both 256, max is 524288.

myisam_use_mmap – use memory mapping for reading and writing myisam table, false by default.

binlog_format = statement, row, mixed. (1,2,3)

New Parameters (part II)

- max_prepared_stmt_count
- Prepared_stmt_count (global status variable)
- innodb_stats_on_metadata

global max_prepared_stmt_count system variable = limit the total # of prepared statements in the server. Limits the potential for denial-of-service attacks based on running the server out of memory by preparing huge numbers of statements. current # of prepared statements is available through the prepared_stmt_count system variable. (Bug#16365)

The prepared_stmt_count system variable has been converted to the Prepared_stmt_count global status variable (viewable with the SHOW GLOBAL STATUS statement). (Bug#23159)

When this variable is enabled (which is the default, as before the variable was created), InnoDB updates stats during metadata statements such as SHOW TABLE STATUS or SHOW INDEX. If disabled, InnoDB does not. Disabling this variable can improve access speed for large number of tables or indexes. It can improve stability of execution plans on InnoDB

New Parameters (part III)

- old
- --skip-write-binlog
- plugins
  - --loose-x

Added the old system variable to cause the server to revert to certain behaviors present in older versions. Currently, this variable affects handling of index hints. Section 12.2.8.2, “Index Hint Syntax”.

--skip-write-binlog: Added --write-binlog option for mysqlcheck. enabled by default; --skip-write-binlog to ANALYZE, OPTIMIZE, and REPAIR TABLE stmts from mysqlcheck not to be binlogged (Bug#26262)

New command-line options: ambiguities in variable names, all vars related to plugins can be specified using a plugin part in the name. For example: innodb in options, plugin-innodb: --skip-plugin-innodb --plugin-innodb-buffer-pool-size=# this is preferred.

--loose-skip-innodb The --loose prefix modifier should be used for all command-line options when uncertain whether plugin exists and when you want to proceed even if the option is necessarily ignored due to the absence of the plugin. Section 4.2.3.1

New Parameters (part IV)

- --slave-exec-mode
  - STRICT, IDEMPOTENT
- Uptime_since_flush_status
- --report-%

Controls whether IDEMPOTENT or STRICT mode is used in replication conflict resolution and error checking. IDEMPOTENT mode causes suppression of some errors, including duplicate-key and no-key-found errors. Beginning withMySQL 5.1.24, this mode should be employed in multi-master replication, circular replication, and some other special replication scenarios. STRICT mode is the default, and is suitable for most other cases.

Uptime_since_flush_status seconds since FLUSH STATUS, provided by Jeremy Cole, #24822

The new read-only global system variables report_host, report_password, report_port, and report_user system variables provide runtime access to the values of corresponding --report-host, --report-password, --report-port, and --report-user options.

New Features

- Meaningful stack traces
- charset on mysql commandline

New charset command added to mysql command-line client. By typing charset name or \C name (such as \C UTF8), the client character set can be changed without reconnecting. (Bug#16217)

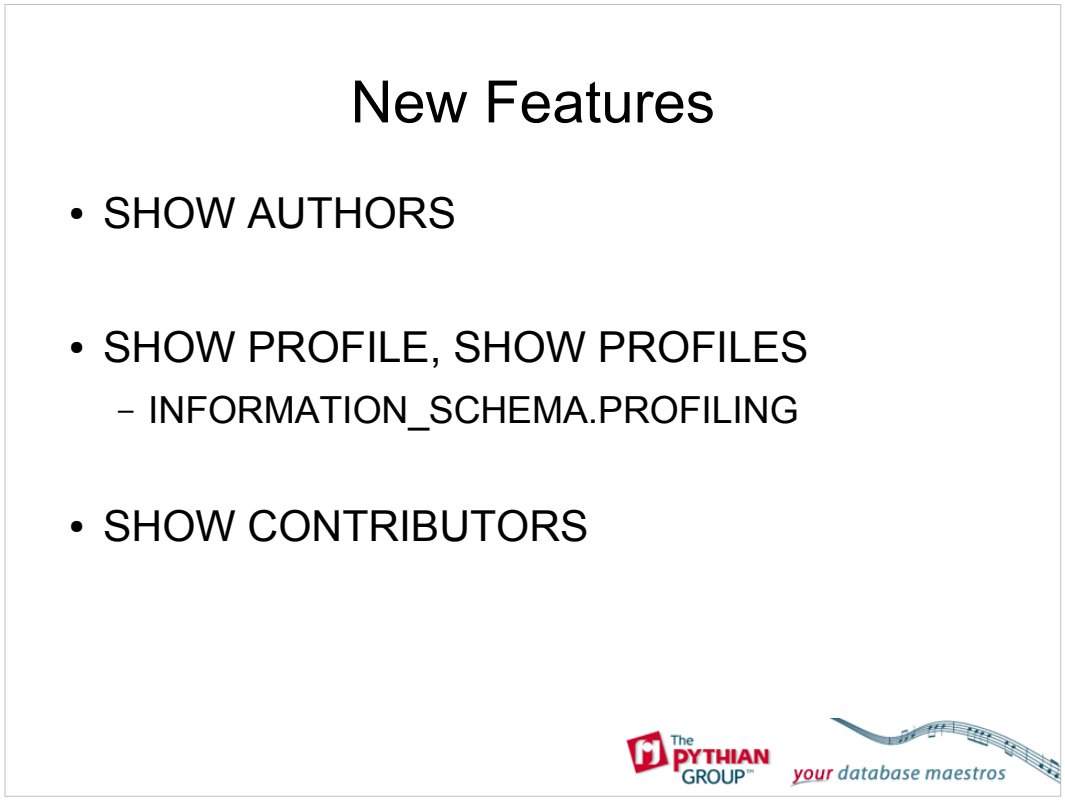

Added the SHOW PROFILES and SHOW PROFILE statements to display statement profile data, and the accompanying INFORMATION_SCHEMA.PROFILING table. Profiling is controlled via the profiling and profiling_history_size session variables. see Section 12.5.5.33, “SHOW PROFILES Syntax”, and Section 20.26, “The INFORMATION_SCHEMA PROFILING Table”. (Community contribution by Jeremy Cole)

The profiling feature is enabled via the --enable-community-features and --enable-profiling options to configure. These options are enabled by default; to disable them, use --disable-community-features and --disable-profiling. (Bug#24795)

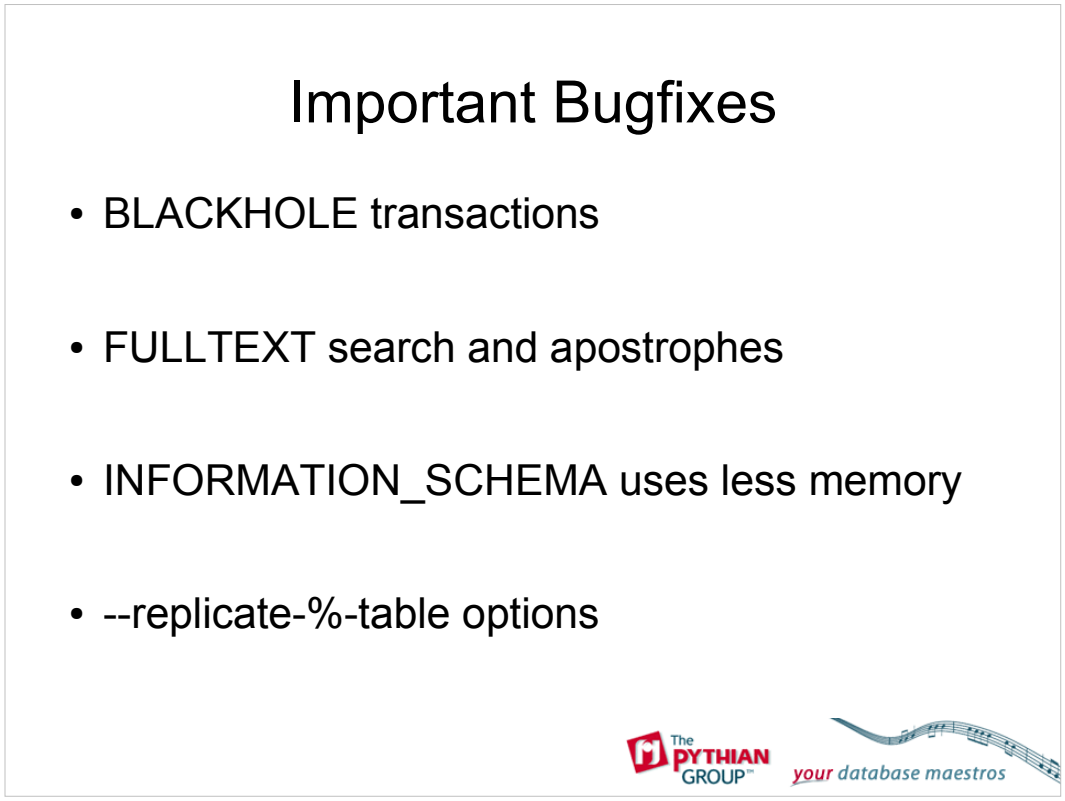

A rolled-back transaction in a BLACKHOLE table was written to the binary log, now they aren't.
http://bugs.mysql.com/15406

Jerry will now match Jerry's

http://bugs.mysql.com/bug.php?id=14194
Users upgrading to this version must issue REPAIR TABLE ... QUICK statements for tables containing FULLTEXT indexes.

There was some data structure cleanup in 5.1

The --replicate-*-table options were not evaluated correctly when replicating multi-table updates.
As a result of this fix, replication of multi-table updates no longer fails when an update references missing table but doesn't update its columns. (Bug#37051)

New Functionalities

- Faster ALTER TABLE for some commands
- RAND() can take non-constant initializers
- ARCHIVE supports AUTO_INCREMENT

Rand() -- constant initializers – initialized once when the query is compiled.

Before 5.0.13, non-constant initializers were undefined.
After 5.0.13, only non-constant initializers were accepted.
In 5.1, non-constant initializers are accepted, but the seed is initialized for each RAND() invocation.

Note: RAND() is replication-safe!

Fast ALTER TABLE: Operations that change only table metadata and not table data do not require a temporary table to be used. For example, renaming a column changes only the .frm file and no longer uses a temporary table. Add column, ADD INDEX and DROP INDEX operations are performed online when the indexes are on variable-width columns only.

New Functionalities (part II)

- MySQL now escapes identifiers
- FULLTEXT search
  - IN NATURAL LANGUAGE MODE
  - .... WITH QUERY EXPANSION
- Stored routines allow DEFINER

http://dev.mysql.com/doc/refman/5.1/en/identifier-mapping.htm

-- anything except ascii null is now allowed in a table or db name. Note that "." is now allowed, as is #. though it is escaped.

IN NATURAL LANGUAGE MODE and IN NATURAL LANGUAGE MODE WITH QUERY EXPANSION -- default is natural language mode, but now there's a new parameter for it.

Now you can specify who the definer should be, for use with SQL SECURITY DEFINER/INVOKER

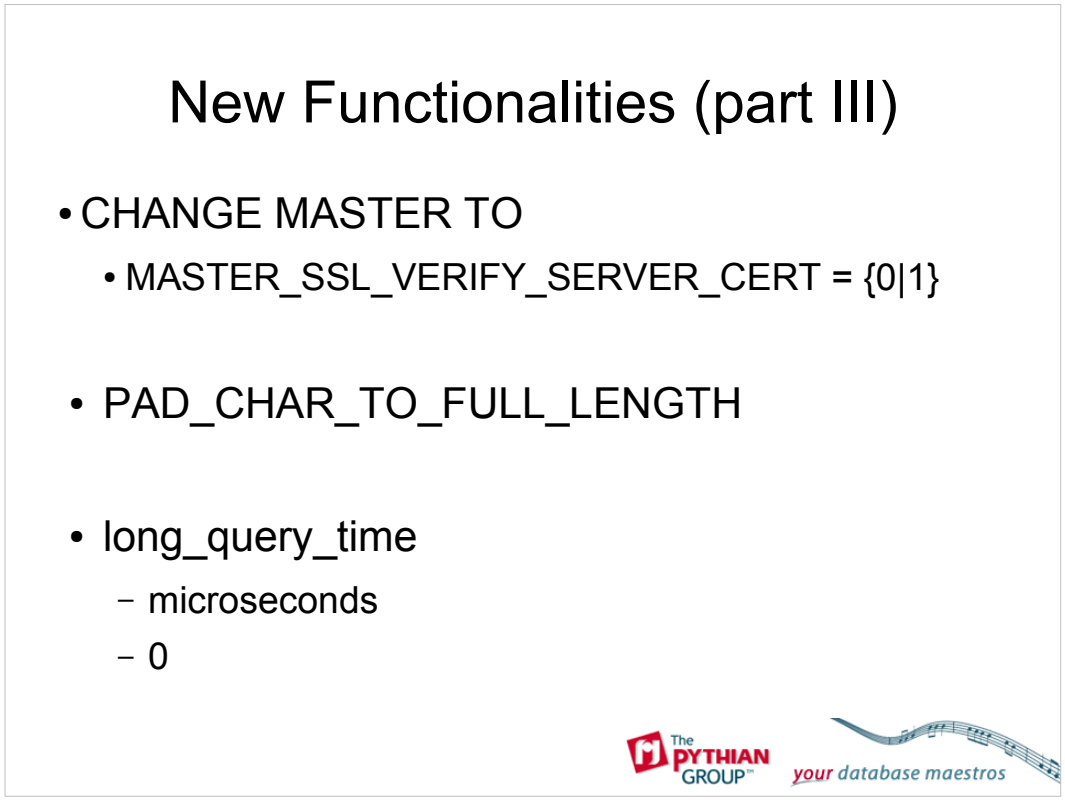

Add the option ssl-verify-server-cert to CHANGE MASTER so that the replication between master and slave can prevent MITM attack.

CHANGE MASTER TO master_def [, master_def] ...
| MASTER_SSL_VERIFY_SERVER_CERT = {0|1}

PAD_CHAR_TO_FULL_LENGTH

By default, trailing spaces are trimmed from CHAR column values on retrieval. If PAD_CHAR_TO_FULL_LENGTH is enabled, trimming does not occur and retrieved CHAR values are padded to their full length. This mode does not apply to VARCHAR columns, for which trailing spaces are retained on retrieval. This mode was added in MySQL 5.1.20.

It is now possible to set long_query_time in microseconds or to 0. Setting this value to 0 causes all queries to be recorded in the slow query log.

New Functionalities (part IV)

- --log-slow-slave-statements
- --min_examined_row_limit
- --innodb_autoinc_lock_mode

--log-slow-slave-statements causes slow statements executed by a replication slave to be written to the slow query log;

min_examined_row_limit can be used to cause queries which examine fewer than the stated number of rows not to be logged.

configure the locking behavior that InnoDB uses for generating auto-increment values. The default behavior now is slightly different from before, which involves a minor incompatibility for multiple-row inserts that specify an explicit value for the auto-increment column in some but not all rows. See Section 13.5.5.3, “AUTO_INCREMENT Handling in InnoDB”.

Deprecated Features

- SHOW INNODB STATUS
- SHOW MUTEX STATUS
- SHOW [BDB] LOGS

These will generate WARNINGS:

SHOW ENGINE INNODB STATUS
SHOW ENGINE INNODB MUTEX

REMOVED: SHOW BDB LOGS, SHOW LOGS

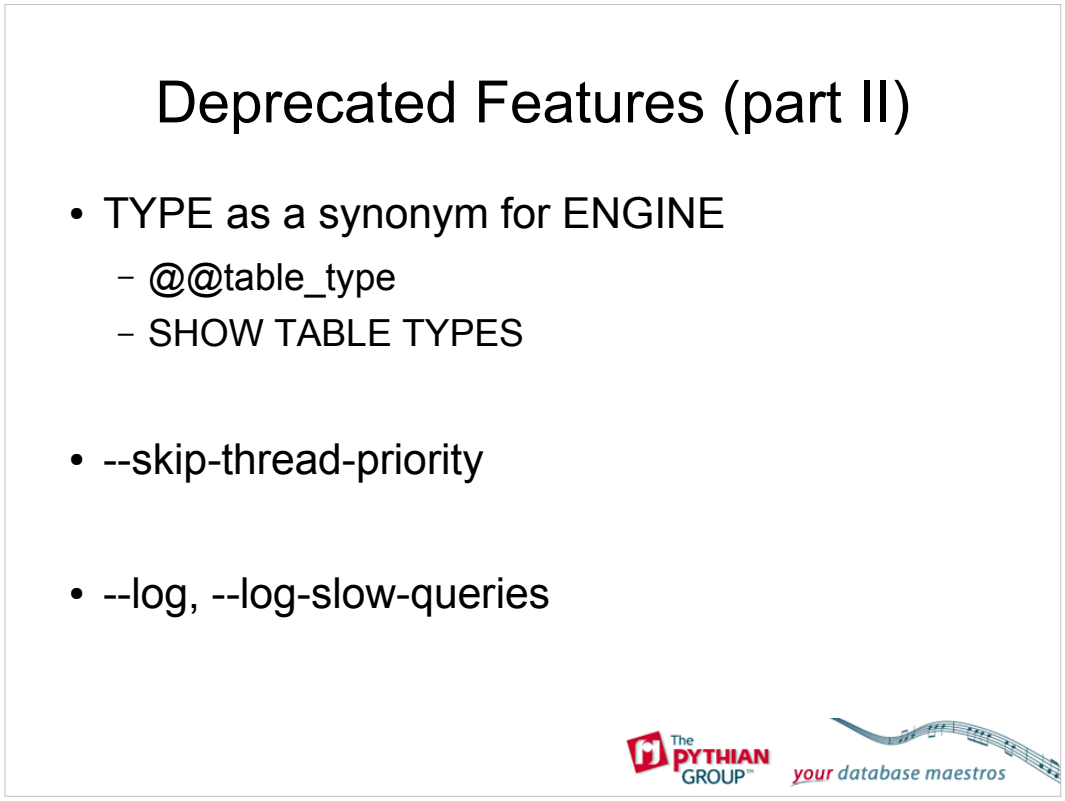

These will generate WARNINGS:

TYPE = engine_name has been deprecated. Use ENGINE=engine_name. (@@storage_engine for @@table_type), SHOW [STORAGE] ENGINES

deprecated: The --skip-thread-priority option is now deprecated in MySQL 5.1 and is removed in MySQL 6.0 such that the server won't change the thread priorities by default. Giving threads different priorities might yield marginal improvements in some platforms (where it actually works), but it might instead cause significant degradation depending on the thread count and number of processors. Meddling with the thread priorities is a not a safe bet as it is very dependent on the behavior of the CPU scheduler and system where MySQL is being run. (Bug#35164, Bug#37536)

--log is deprecated, use general_log, general_log_file, --slow_query_log, --slow_query_log_file

Package Changes

- mysqld-debug included
- safe_mysqld excluded
- mysqldumpslow part of server, not client
- MySQL Cluster is a separate product

mysqld-debug used to be a separate release.

safe_mysqld has been an alias to mysqld_safe, deprecated since 4.0. It's now out.

Mysqldumpslow relies on my_print_defaults

New Collations

- cp1250_polish_ci
- utf8_hungarian_ci
- ucs2_hungarian_ci

cp1250_polish_ci collation for the cp1250 character set.

These support the correct sort order for Hungarian vowels. However, they do not support the correct order for sorting Hungarian consonant contractions; we expect to fix this issue in a future release.

Sheeri Cabral

cabral@pythian.com

www.sheeri.com