MySQL

New Features 5.6

FOSDEM MySQL and Friends Devroom
February 5, 2012, ULB Brussels

Oli Sennhauser

Senior MySQL Consultant, FromDual

oli.sennhauser@fromdual.com

FromDual

- FromDual provides neutral and independent:
  - Consulting
  - Remote-DBA
  - Support for MySQL, Galera, Percona Server
  - Training

- Oracle Silver Partner (OPN)

- More information about us: http://www.fromdual.com

Contents

- Milestone Release Model
- Partitioning
- InnoDB
- Optimizer
- Performance Schema (P_S)
- Replication
- Various

Milestone Release Model

- Between 5.4 and 5.5 MySQL introduced the new "Milestone Release Model"
- Dynamic Model for development in theory:
  - Starts and always is at least in beta quality
  - Milestone releases, with RC quality, every 3 – 6 months
  - Between Milestones new features allowed
  - GA releases every 12 – 18 months (5.5: Oct. 2010)
  - No more than 2 releases in active support!
- MySQL Lab releases

Milestone Release Model

- Look at the schedule:
  - 5.6.0+1 Milestone 4+5, not released
  - 5.6.2 no Milestone number, released April 2011
  - 5.6.3 + 5.6.4 (M6 + M7) both October 2011
  - 5.6.5 (M8) not released yet (February 2012)
- In practice:

  → I cannot see much practical differences
- My guess: 5.6 GA at Collaborate in April 2012

  → As a consequence: 5.0 and 5.1 EOL!
  - Who is still at 5.0 and 5.1?

NF 5.6 / Partitioning

- Explicit Partition Selection

```
SELECT *
  FROM sales PARTITION (p2010, p2011)
 WHERE sum < 100;
```

  - → Caution: implicit `WHERE` clause!

- Exchanging Partitions

```
ALTER TABLE sales_hist
EXCHANGE PARTITION p2011
WITH TABLE sales;
```

  - ETL, DWH

NF 5.6 / InnoDB

- InnoDB `INFORMATION_SCHEMA` tables
  - `INNODB_BUFFER_*`
  - `INNODB_FT_*`
  - `INNODB_METRICS`
  - `INNODB_SYS_*`

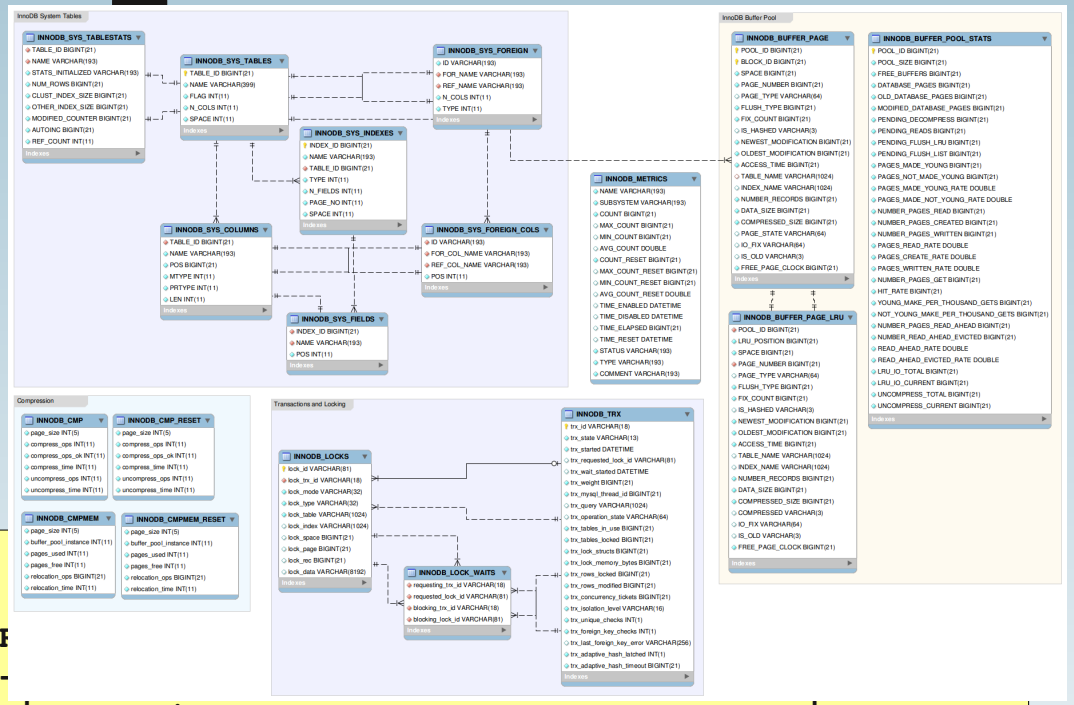


```
SELECT name, subsystem, count, comment
  FROM INFORMATION_SCHEMA.innodb_metrics WHER
+----------------------+-------------+-------
| name                 | subsystem   | count | comment                                    |
+----------------------+-------------+-------+--------------------------------------------+
| trx_rseg_history_len | transaction |   654 | Length of the TRX_RSEG_HISTORY list        |
| trx_rseg_curent_size | transaction |     0 | Current rollback segment size in pages     |
+----------------------+-------------+-------+--------------------------------------------+
```

- Most of InnoDB Monitor is now obsolete!

NF 5.6 / InnoDB performance

- Page cleaner thread (before master thread)
- `innodb_purge_threads` can be set > 1
- Kernel mutex split → improved concurrency
- Concurrent read while creating secondary index
- Improved warm-up:
  - `innodb_buffer_pool_dump_at_shutdown`
- InnoDB REDO log size up to 512 Gbyte
- InnoDB threads scheduling better > 16 threads
- UNDO log → separate TS (random I/O → SSD!)
- Improved concurrency extending TS files `innodb_file_per_table`

NF 5.6 / InnoDB

- InnoDB and Optimizer:
  - Persistent Optimizer Statistics
  - Control of Statistics sampling (random dives)

```
SET GLOBAL innodb_analyze_is_persistent = 1;
ANALYZE TABLE valuemaps;
SELECT * FROM mysql.innodb_table_stats;
+---------------+------------+--------+----------------------+--------------------------+
| database_name | table_name | n_rows | clustered_index_size | sum_of_other_index_sizes |
+---------------+------------+--------+----------------------+--------------------------+
| zabbix        | valuemaps  |      7 |                    1 |                        1 |
+---------------+------------+--------+----------------------+--------------------------+
```

- Deadlocks to the error log
  - `innodb_print_all_deadlocks`
- InnoDB page size can be: 16k, 8k, 4k
  - `innodb_page_size`

NF 5.6 / Optimizer

- ORDER BY on non indexed columns

  → `sort_buffer` avoid sort merge passes
- Multi-Range Read (MRR)

  → optimize Range Scan on secondary indexes
- Index Condition Pushdown (ICP)

  → `WHERE` is evaluated in the Storage Engine
- Query Execution Plan for DML Statements!
- Optimization of derived tables (`FROM` clause)

  → Postponed materialization and index on derived table
- Batched Key Access (BKA) → improved `JOIN` performance
- Optimizer trace!

NF 5.6 / Performance Schema

- MySQL introduced with 5.5 the P_S
  - The idea measure everything!
  - Has some performance impact!
- New instrumentation for:
  - Table read and write (row-level accesses)
  - Stages and statements (stage = state)
  - Connections, Sockets
  - Table lock wait events
  - Table and index I/O wait events
- Filter by
  - thread
  - object

NF 5.6 / Replication

Sergei Petrunia at 15:00

Replication features of 2011

NF 5.6 / Replication

- Globally Unique Server ID

```
cat $datadir/auto.cnf
[auto]
server-uuid=db731167-2b4c-11e1-928c-bcaec586ca65
```

- Delayed Replication
  - Before: `mk-slave-delay` (Maatkit)

```
CHANGE MASTER TO MASTER_DELAY = 42;
```

- Timestamp added to `SHOW SLAVE STATUS`:

```
SHOW SLAVE STATUS\G
...
     Last_IO_Error_Timestamp: 120130 16:59:12
     Last_SQL_Error_Timestamp:
```

NF 5.6 / Replication

- Row Image Control
  - RBR
  - Save: disk space, network resources, memory

```
binlog_row_image = {full | minimal | noblob}
```

- Crash safe Binary Logs
  - Event length + CRC32 checksum
  - Reading and writing on Master and Slave
  - Complete Events/Trx from/to binary log

NF 5.6 / Replication

- Slave Log Tables
  - `master.info` and `relay-log.info` into tables:
  - `mysql.slave_*_info`
  - MyISAM :(→ convert to InnoDB?

```
master-info-repository          = TABLE
relay-log-info-repository       = TABLE
```

Parallel Event Execution (multi-threaded slave)

- Per schema

```
slave_parallel_workers = <n>
```

NF 5.6 / Various

- Fractional seconds
  - Up to microseconds (0.000001 s)
- `GET DIAGNOSTICS`
  - For Stored Programs
- Fulltext indexes on InnoDB tables!
- Pluggable authentication (Socket)
- Memcached Plug-in (still in Labs :( )
- `mysqlbinlog --read-from-remote-server --raw`

Summary

- Bugs fixed: +400 bugs
- Clean-up!
- And many many more smaller features...
- Incompatible changes!

  → Upgrade

Q & A

Questions ?

Discussion?

We have some time for face-to-face talks...