



Performance Improvements and New Features in MySQL 5.6

Vadim Tkachenko,
CTO, Percona
Percona MySQL University
Portland, OR
Jun 17, 2013

About Presentation

Brief Overview

Birds eye view of features coming in 5.6

Mainly documentation based review

- MySQL 5.6 has not been long used in production yet

MySQL 5.6

Looks to be great release!

- Much better than initial 5.1 and 5.5 releases

Have been GA since 5th of February 2013

A bit over 2 years since MySQL 5.5 GA

- Work on some 5.6 features started 4+ years ago

A lot of focus on conventional MySQL usage

When to Upgrade to MySQL 5.6

What specific benefits are you looking from upgrade ?

Start looking at it for development purposes & testing

Upgrade in 6 months unless you critically need something

- Let “baby bugs” to be worked out
- Production usage experience gathered

Percona and MySQL 5.6

Percona is ready to help you to be successful with MySQL 5.6

MySQL 5.6 is fully supported by our Support, Consulting, RemoteDBA teams

Percona Software support for 5.6 is current or on a way

Learn more

- <http://www.percona.com/mysql-5.6-solutions-from-percona>
- <http://www.mysqlperformanceblog.com/2013/02/25/percona-welcomes-mysql-5-6/>

Percona Server 5.6

Currently RC

GA Release : approximately mid Jul 2013

Includes All MySQL 5.6 features plus Percona Improvements



PERCONA
SERVER

Percona XtraDB Cluster 5.6

- Work has started on PXC 5.6
- Great potential of joining MySQL 5.6 replication improvements
 - With Galera Replication technology



PERCONA
XTRADB CLUSTER

Feature Categories

Scalability

Optimizer and Execution

Replication

Transparency

Database Operations

Functionality for Developers

Security

Scalability

InnoDB Storage engine in Focus

Scalable Read Only Transactions

Concurrent InnoDB Data File Extension

Non Recursive Deadlock Detection

Faster Locking Primitives

Improved InnoDB Thread Concurrency

Multiple Background Purge Threads

Improved Purge Lag Control

Improved Adaptive Flushing

Support for Large (over 4GB) redo logs

Scalability

Split of “Kernel
Mutex”

Data Dictionary
Cache

Page Cleaner /
Separate Flush
Thread

Group Commit for
Binary Log

Fight Cache
Coherence and
False Sharing
Issues

Reduced Innodb
Memory
Fragmentation

Reduced Locking
for Partitioned
Tables


Reduced
Contention for
LOCK_open

Multiple
table_open_cache
instances

MySQL 5.6 performance

Multiple CPU cores at focus

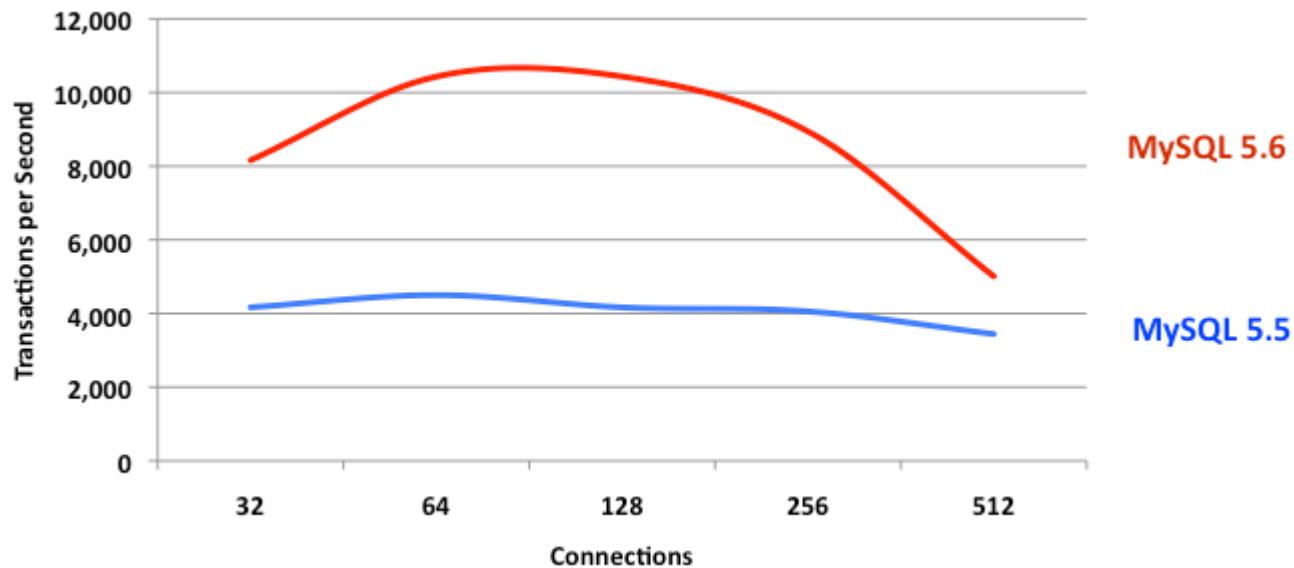
- Results from Oracle



MySQL 5.6 Performance

Number of connections at Focus


- Results from Oracle



Binary log Group Commit

Gains with sync_binlog=1

- Results by Oracle



Optimizer and Execution

Index Condition Pushdown (ICP)

Batched Key Access Joins (BKA)

Multi Range Read (MRR)

Faster **ORDER BY** `nidxcol` **LIMIT N**

Subquery Optimizations

Optimizer and Execution

Fast Page Checksums (CRC32)


4K and 8K Page sizes for Innodb

Persistent Statistics for Innodb

Improvements to Innodb Compression

BKA+MRR May rock your world

- Reporting Query Benchmark
 - Results from Oracle



Replication

Multi-Threaded Slave

Global Transactions Identifiers

Crash Safe Slave and Binlog

Replication Event Checksums

Time Delayed Replication

Server UUID

Improved Logging for ROW based Repl.

Replication Utilities for Failover and Adm.

Multi Threaded Slave

- Peak gains with Multiple Schema
 - Results by Oracle



Transparency

Many new **INFORMATION_SCHEMA** tables

- **INNODB_METRICS**
- Table Meta Data Information Tables
- Buffer Pool Contents Information Tables

Improved **EXPLAIN**

- Handling **INSERT/UPDATE/DELETE**
- **JSON** output with additional information

Optimizer Tracing

Deadlock Logging

Innodb Metrics Example

```
mysql [localhost] {msandbox} (information_schema) > select * from innodb_metrics where
status!="disabled" and count!=max_count limit 1 \G
***** 1. row *****
      NAME: buffer_pool_pages_free
      SUBSYSTEM: buffer
      COUNT: 2359
      MAX_COUNT: 7957
      MIN_COUNT: 2359
      AVG_COUNT: NULL
      COUNT_RESET: 2359
      MAX_COUNT_RESET: 7957
      MIN_COUNT_RESET: 2359
      AVG_COUNT_RESET: NULL
      TIME_ENABLED: 2013-03-06 09:43:43
      TIME_DISABLED: NULL
      TIME_ELAPSED: 471
      TIME_RESET: NULL
      STATUS: enabled
      TYPE: value
      COMMENT: Buffer pages currently free (innodb_buffer_pool_pages_free)
1 row in set (0.01 sec)
```

JSON Explain Example

```
mysql [localhost] {msandbox} (sptest) > explain format=json delete from sptest where k>2 \G
***** 1. row *****
EXPLAIN: {
  "query_block": {
    "select_id": 1,
    "table": {
      "delete": true,
      "table_name": "sptest",
      "access_type": "range",
      "possible_keys": [
        "PRIMARY",
        "k"
      ],
      "key": "k",
      "used_key_parts": [
        "k"
      ],
      "key_length": "4",
      "rows": 1,
      "filtered": 100,
      "attached_condition": "(`sptest`.`sptest`.`k` > 2)"
    }
  }
}
1 row in set (0.00 sec)
```

Transparency

Improved PERFORMANCE_SCHEMA

- Reduced Overhead
- Simplified Configuration
- Table Access Instrumentation
- Statement Instrumentation
- Stages Instrumentation
- Aggregation by User, Host etc
- Network IO Instrumentation
- Host Cache contents
- Improved File I/O Instrumentation

PERFORMANCE_SCHEMA Example

```
mysql [localhost] {msandbox} (performance_schema) > select * from users;
```

```
+-----+-----+-----+
| USER   | CURRENT_CONNECTIONS | TOTAL_CONNECTIONS |
+-----+-----+-----+
| NULL   |          18         |          21        |
| msandbox |           1         |           2        |
+-----+-----+-----+
```

```
2 rows in set (0.00 sec)
```

.. Not replacement for User Statistics

```
mysql> select * from information_schema.user_statistics limit 1\G
***** 1. row *****
      USER: root
    TOTAL_CONNECTIONS: 2
  CONCURRENT_CONNECTIONS: 0
    CONNECTED_TIME: 4
      BUSY_TIME: 0
      CPU_TIME: 0
    BYTES_RECEIVED: 67
      BYTES_SENT: 0
  BINLOG_BYTES_WRITTEN: 0
    ROWS_FETCHED: 2
    ROWS_UPDATED: 0
  TABLE_ROWS_READ: 0
  SELECT_COMMANDS: 2
  UPDATE_COMMANDS: 0
    OTHER_COMMANDS: 0
  COMMIT_TRANSACTIONS: 0
  ROLLBACK_TRANSACTIONS: 0
  DENIED_CONNECTIONS: 0
  LOST_CONNECTIONS: 0
    ACCESS_DENIED: 0
    EMPTY_QUERIES: 0
  TOTAL_SSL_CONNECTIONS: 0
1 row in set (0.00 sec)
```


Operational Improvements

Separate Tablespaces for UNDO Logs

Fast Restart – Innodb BP preloading

Online DDL

Import/Export for Partitioned Tables

Remote Binlog Backup

Innodb Transportable Tablespaces

New configuration variables defaults

User defined **DATA DICTIONARY** for Innodb tables

New Functionality for Developers

Memcached API for Innodb

Explicit Partition Selection in Queries

Full Text Search indexes for Innodb

Microsecond TIME precision

Precise spatial operations for GIS

Security

Passwords hashes in Query Logs

SHA256 hashing with salt for Authentication

Support obfuscated password storage for command line tools

Policy Based password validation

Plugin based Authentication support in Replication

Where to Learn More ?

MySQL
5.6 Manual
is great

- <http://dev.mysql.com/doc/refman/5.6/en/mysql-nutshell.html>

Blogs

- <http://blogs.innodb.com>
- <http://www.planetmysql.org>
- <http://www.mysqlperformanceblog.com>

Santa Clara, April 22-25, 2013



www.percona.com/live

Thank You!

Vadim Tkachenko
vadim@percona.com
@VadimTK