



# Improvements in MySQL 5.5 and 5.6

Peter Zaitsev  
Percona Live NYC  
May 26, 2011

# State of MySQL 5.5 and 5.6

- MySQL 5.5
  - Released as GA December 2011
  - Percona Server 5.5 released in April 2011
  - Proven to be rather stable release
- MySQL 5.6
  - Development release in April 2011
  - Preview of features available at [labs.mysql.com](http://labs.mysql.com)
  - No release date promised so far
    - 18-24 month release cycle previously named

# MySQL 5.5 overview

- Better Scalability
- Performance\_Schema
- Improved Windows Support
- Semi-Synchronous Replication & Heartbeat
- SIGNAL/RESIGNAL in Stored Procedures
- Improved Partitioning Support

# LOCK\_Open

- “Table Cache” Operations did not scale well in MySQL
  - A lot of operations had to be done under global lock making it hot spot
  - Linear Search used for some operations
- MySQL 5.5 reduces contention introducing MDL (Meta Data Locking) Subsystem
  - Both optimizations for number of operations and reduced contention

# Smaller Optimizations

- LOCK\_alarm mutex
  - Replaced by SO\_SNDTIMEO/SO\_RCVTIMEO on platforms what support it
    - Most current do
- LOCK\_thread\_count mutex
  - Reduced usage to connect/disconnect from 1-2 times per query
- THR\_LOCK\_charset
  - Previously used for character set setup. Removed

# Improved Log\_sys\_mutex

- Served “double duty” protecting
  - Log records and LSN
  - Pages in buffer pool changed than MTR is committed
- Changed to 2 separate mutexes
- Contention can still be observed by heavy log write workloads

# Separate Flush\_list Mutex

- Buffer Pool operations and Flush List operations were bound on very hot buffer pool mutex
- Flush list has its own mutex reducing pressure on buffer pool mutex
  - Which remains rather hot

# Multiple Buffer Pool Instances

- “buffer pool mutex remains hot”
- But we can now create multiple buffer pools to relief this problem
  - **innodb\_buffer\_pool\_instances**
- Pages will be hashed across these buffer pools
  - No user level table mapping etc
- Some will contain hot pages and may get hot
- Many buffer pools can increase internal overhead

# Multiple Rollback Segments

- InnoDB used to have single rollback segment
  - Both contention spot and limit to 1023 concurrent transactions (or less)
- Now 128 Segments
  - Reducing contention and giving up to 128K transactions
- **innodb\_fast\_shutdown=0** needed for upgrade
- Old InnoDB Versions will benefit from this change
  - They can use more than one segment, just do not create more than one

# Asynchronous IO on Linux

- Before MySQL 5.5 InnoDB would use Native Async IO only on Windows
- On Linux it was “Simulated” by using helper threads
- Assumes `O_DIRECT`
  - Good to use anyway
- More outstanding IO requests
- Pushes burden of merging io requests to Kernel
- Can make crash recovery up to 1.5x faster

# Change Buffering

- Insert Buffer existed in Innodb forever
  - Speeding up Inserts for IO bound workloads
- Updates and Deletes were a lot slower, unable to use this feature
- MySQL 5.5 support “delete mark” and “purge” change buffering
  - It handles update case as well as delete because secondary indexes are never updated in place
- Up to 40x performance improvement for large deletes

# Separate Purge Thread

- Before MySQL 5.5 purging of old records is done by main thread
  - Making it stall on other activities while doing purge
  - Making purging lag more often than it should
  - Very large “History Length” with purge thread unable to keep up a big problem
- **innodb\_purge\_threads=1** to use separate purge thread

# InnoDB Adaptive Flushing


- Not a new feature in MySQL 5.5
- Was adjusted to deal with multiple buffer pools
- Getting constants small tunings
- MySQL 5.5 should be a lot “smoother” on many workloads
- This will likely continue to improve as MySQL 5.5 matures

# Windows Optimizations

- Using Native Windows calls instead of POSIX subsystem
- Use CriticalSection for mutexes and use native ConditionVariables in newer Windows OS
- Windows Optimizations graphs from Innodb Blog

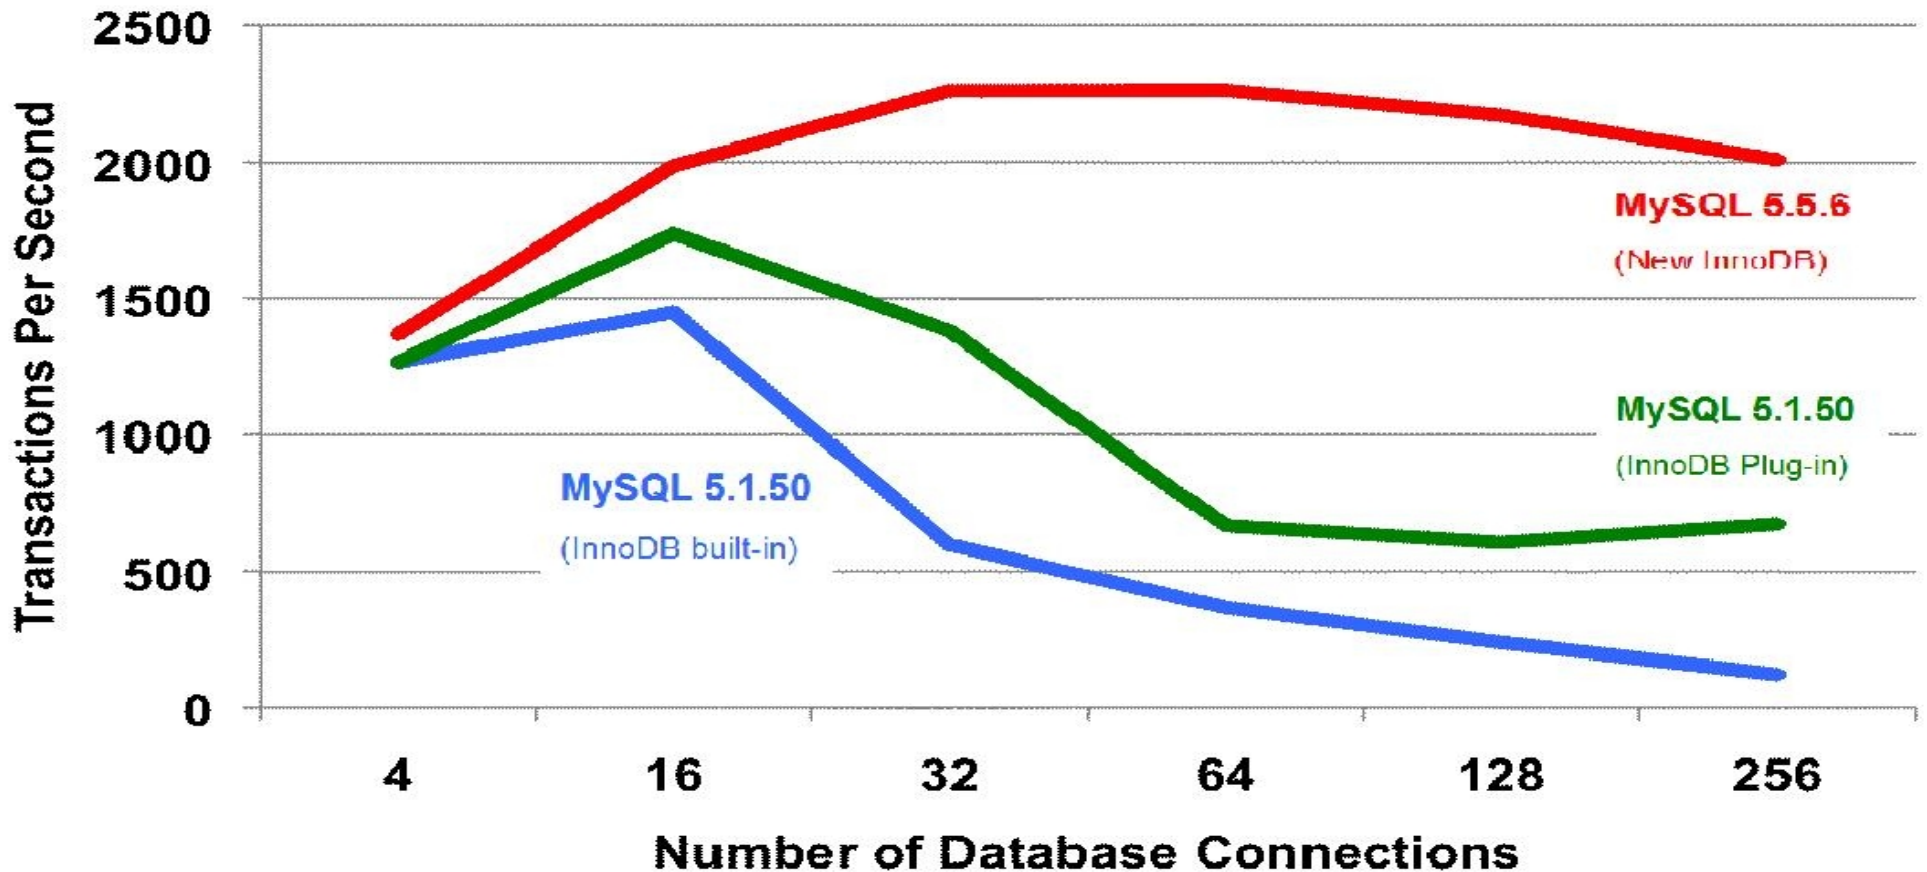
# Windows Sysbench Reads

## MySQL 5.5 vs. 5.1 - Read Only



# Windows Sysbench RW

## MySQL 5.5 vs. 5.1 - Read Write



# Is MySQL 5.5 Always Faster ?

- No. There are cases then MySQL 5.5 can be slower
  - Especially single thread workloads may be affected
  - This is natural for any new MySQL
- Mark Callaghan has reported some sysbench regressions
  - And there are some user complains
- There are no big specific regressions we can report now.

# Remaining Issues in MySQL 5.5

- Per Index Lock
  - Restricts performance for high update workloads
  - Use Partitioning as partial Workaround
- Kernel Mutex
  - Seems to be significant focus for Innodb optimization in MySQL 5.6
- Adaptive hash index
  - Single adaptive hash index often causes contention
- Compressed tables
  - Were not well optimized for scalability

# Remaining issues In MySQL 5.5


- Dropping Tablespace with Innodb File Per Table
  - Operation requires scan of LRU list
  - So the larger is buffer pool the slower it gets
- Log Mutex
  - Still causing contention
- Group Commit
  - Is “fixed” now only if you chose not to use **sync\_binlog=1** and so lose binary logs on crash

# What about Percona Server 5.5 ?


- Percona Server 5.5 builds on performance improvements in Percona Server 5.1
- No MySQL 5.5 features are removed
- Multiple adaptive hash indexes
- Fast index creation for OPTIMIZE TABLE, ALTER TABLE

# There are more problems to be solved

- Scalability on 24 way box is just 6x for some workloads



# MySQL 5.5 vs Percona Server 5.5



# MySQL 5.6 improvements

- MySQL 5.6 is Development release, so list is not final
- Optimizer Improvements
  - Index Condition Pushdown and Multi Range Reads
- ORDER BY LIMIT N for non indexed queries
- Persistent Innodb Statistics
  - Already available in Percona Server 5.1+

# Great stuff in MySQL 5.6

- Metrics Table
  - Lightweight way to get more information about innodb
- Expose internal data dictionary as system tables
  - Similar Available in Percona Server 5.1+
- Buffer Pool Information Table
  - Similar Available in Percona Server 5.1+

# Awesome features in MySQL 5.6

- Multiple Purge Threads
  - Available in Percona Server 5.1+
- Even better Innodb Scalability
  - Specially kernel\_mutex work
- NoSQL Access through memcached API
  - Percona Server has access through HandlerSocket
- Memcached access to MySQL Cluster (NDB)

# Amazing Enhancements MySQL 5.6

- Crash Safe Slaves when using InnoDB
  - Basic support exists in Percona Server 5.1
- Replication Events Checksums
- Improved Row Level Replication
  - Partial before/after images
- Time Delayed Replication
- Store Query Text in Row based binary log
- Remote Binlog backup

# Even more MySQL 5.6 features

- Server UUIDs (no needs to manually assign)
- Multi-Threaded Slave (parallel replication)
- Improved Performance Schema
  - Table level IO and logical operations
  - Table Locks
  - Statement Information
- Better Partitioning
  - Explicitly selecting partitions for queries

# Other MySQL 5.6 Changes

- Pruning InnoDB Internal Dictionary Cache
  - Available in Percona Server 5.1+
- Import/Export partitions from the table
  - Exchange partition and table

# DIAMOND SPONSORSHIPS





**baron@percona.com**

**We're Hiring! [www.percona.com/about-us/careers/](http://www.percona.com/about-us/careers/)**