

Why we're excited about MySQL 8

Practical Look for Devs ~~and Ops~~

Peter Zaitsev, CEO, Percona

February 4nd, 2018

FOSDEM



In the Presentation

**Practical view
on MySQL 8**

~~**Exciting things
for Devs**~~

**Exciting things
for Ops**

Warning

This assessment is done for Pre-GA MySQL 8, based on documentation and limited testing. We're yet to see how they behave in production

Source Notes

Examples liberally borrowed
from Oracle team
presentations and Blog Posts

MySQL 8 for Ops

Ops care about

Stability

High Availability

Performance

Security

Observability

Manageability

Native Data Dictionary

About 10 years overdue


Atomic, Crash Save, DDLs

Much Faster Information Schema

No more MyISAM System Tables!

Fast Information Schema

```
1 SELECT t.table_schema, t.table_name, c.column_name
2 FROM information_schema.tables t,
3      information_schema.columns c
4 WHERE t.table_schema = c.table_schema
5        AND t.table_name = c.table_name
6        AND t.engine='InnoDB';
```






Scaling from 5.000 to 1.000.000 tables

Much Better and Faster UTF8

- utf8mb4 as Default Charset

MySQL 8.0 vs MySQL 5.7 utf8mb4



+300-350% in OLTP RO
+176-233% in OLTP RW
+1500-1800% in SELECT DISTINCT_RANGES

Security

ROLES

Breakdown of SUPER Privileges

Password History

Faster cached-SHA2 Authentication (Default)

Community Builds use OpenSSL

--skip-grants blocks remote connections

Redo and Undo Logs are now encrypted if Table Encryption is enabled

Persistent Auto Increment

Another feature 10 years overdue

Do not reset AUTO INCREMENT to the max value in the table on restart

Auto-Managed Undo Tablespace

Do not use system table space for undo space any more

Automatically reclaim space on disk from large transactions

Self Tuning (limited to Innodb)

Set
innodb_dedicated_server
to auto-tune

- innodb_buffer_pool_size
- innodb_log_file_size
- innodb_flush_method

Partial In-Place Update for JSON

Can update field in JSON object without full re-write

Great for counters, statuses, timestamps etc

Only update and removal of element is supported

Innodb Support Added as of MySQL 8.0.4

Invisible Indexes

- Test impact of dropping indexes before actually dropping them
- Can use use_invisible_indexes to use invisible indexes in a session

```
1  CREATE TABLE t1 (  
2      i INT,  
3      j INT,  
4      k INT,  
5      INDEX i_idx (i) INVISIBLE  
6  ) ENGINE = InnoDB;  
7  CREATE INDEX j_idx ON t1 (j) INVISIBLE;  
8  ALTER TABLE t1 ADD INDEX k_idx (k) INVISIBLE;
```

TmpTable Storage Engine

More efficient storage engine for Internal Temporary tables

Efficient storage for VARCHAR and VARBINARY columns

BLOB/TEXT Columns are not supported (yet?)

Backup Locks

Prevent operation which may result in inconsistent backups

LOCK INSTANCE FOR BACKUP

Optimizer Histograms

- Detailed Statistics on Columns, not just Indexes

```
1  {
2    "buckets": [
3      [
4        1,
5        0.3333333333333333
6      ],
7      [
8        2,
9        0.6666666666666666
10     ],
11     [
12       3,
13       1
14     ]
15   ],
16   "null-values": 0,
17   "last-updated": "2017-03-24 13:32:40.000000",
18   "sampling-rate": 1,
19   "histogram-type": "singleton",
20   "number-of-buckets-specified": 128,
21   "data-type": "int",
22   "collation-id": 8
23 }
```

Improved Optimizer Cost Model

**Keep in account how
much of data is cached vs
on disk**

<http://bit.ly/2CZu6z6>

More on MySQL 8 Optimizer

<http://www.unofficialmysqlguide.com/>

Performance Schema

(Fake) Indexes for Faster Access

Error Instrumentation

Response Time Histograms (Global and Per Query Digest)


Query Examples for Summary by Digest

Performance Schema Performance

- Now is Interactively Usable at Scale

Performance Comparison

Over 30x faster!



Remote Management Features

RESTART (Command)

SET PERSIST innodb buffer pool size =
1024 * 1024 * 1024;

Assumes storage is SSD by Default

Start of the long journey

Binary Log On by Default

bin_log is enabled by default

log_slave_updates is enabled by default

Expire logs after 30 days by default

Query Cache Removed

It's design caused more problems than it fixed

Use ProxySQL (or other) external query cache instead

Native Partitioning Only

Only “Native” Partitioning supported, not Generic One

Remove partitions from MyISAM partitioned tables or convert them

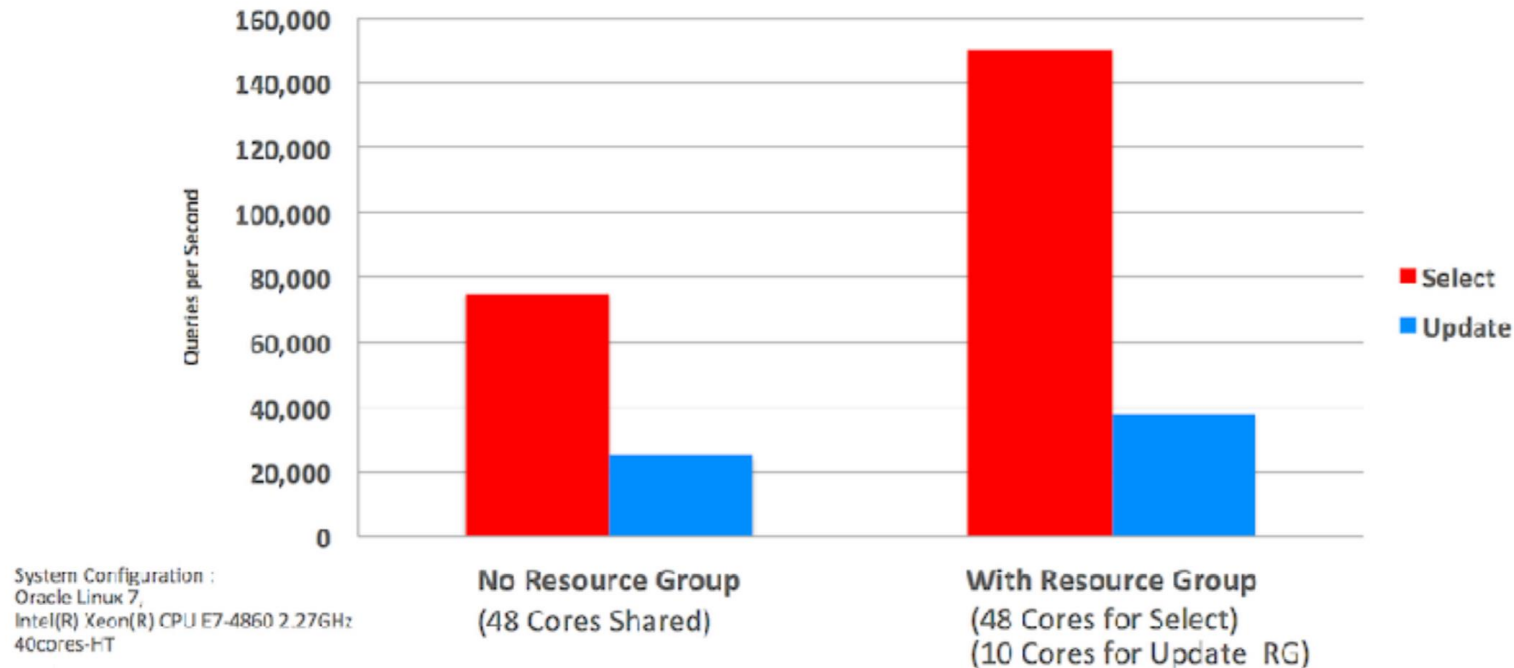
ALTER TABLE ... REMOVE PARTITIONING

ALTER TABLE ... ENGINE=INNODB

Resource Groups

- Isolation and Better Performance

MySQL 8.0 Resource Groups - 100% Faster




Plain Better Performance at Scale

Sysbench: OLTP_RO Point-Selects

2.1x Faster than MySQL 5.7

2.8x Faster than MySQL 5.6



MySQL 8 for PMM

MySQL 8 will allow us to make Percona Monitoring and Management even Better!



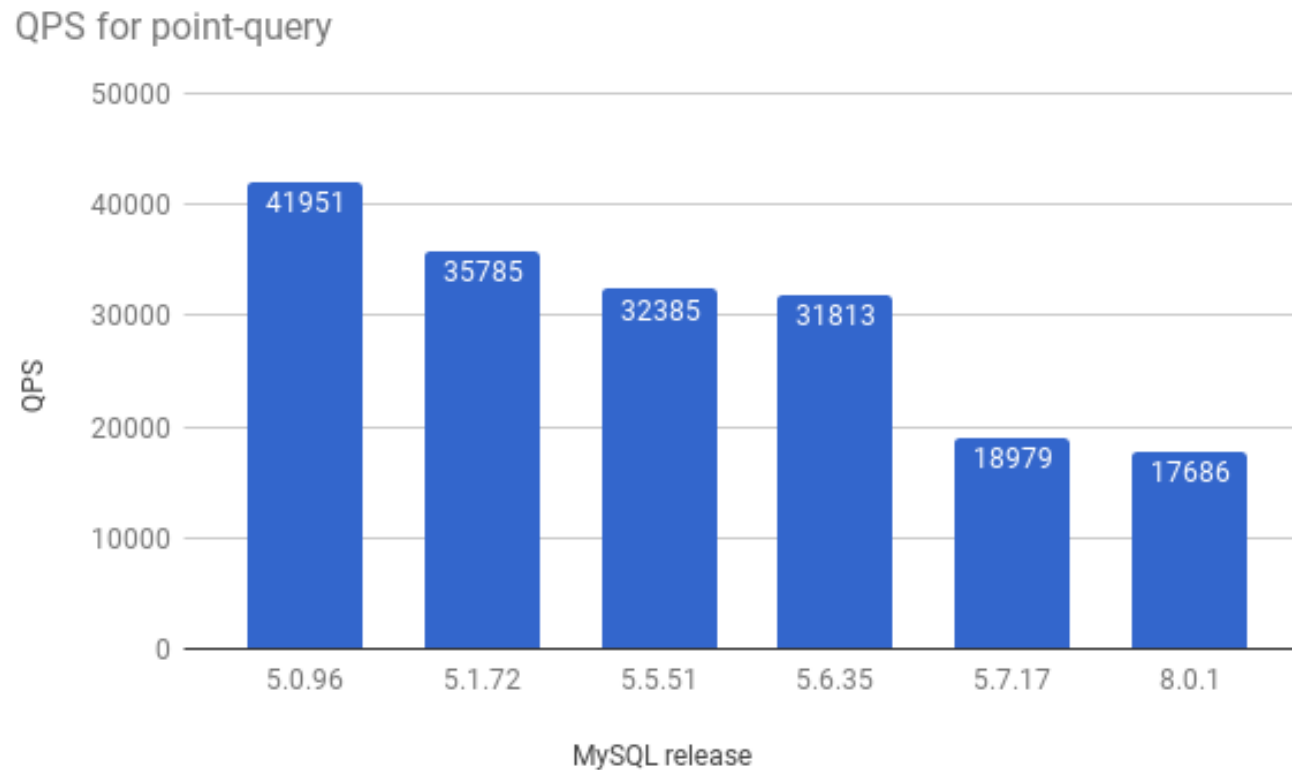
Feature Requests

Better Single Thread Performance

**Parallel Single Query Processing
Please**

Sorry state of MySQL single thread performance

<http://bit.ly/2oMvu2a>



SAVE THE DATE!

April 23-25, 2018
Santa Clara Convention Center



PERCONA

LIVE

www.perconalive.com

Thank You!
