

Why we're excited about MySQL 8

Practical Look for Devs and Ops

Peter Zaitsev, CEO

October 2, 2017



About Percona

Solutions for your success with MySQL, MariaDB and MongoDB

Support, Managed Services, Consulting, Training, Software

Our Software is 100% Free and Open Source

Support Broad Ecosystem – MySQL, MariaDB, Amazon RDS, Google CloudSQL

In Business for 11 years

More than 3000 customers, including top Internet companies and enterprises

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

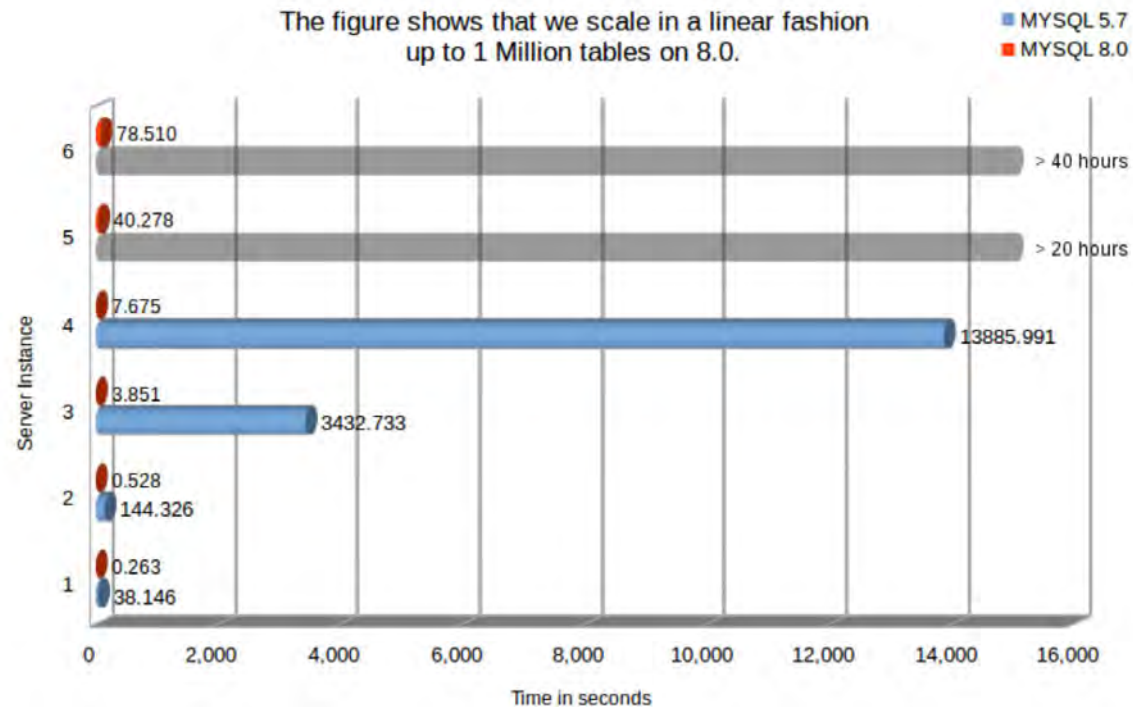
Atomic (Crash Save) DDLs

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

--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

Only Optimizer and Replication support so far

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**

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!



Persistent Global Variables

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



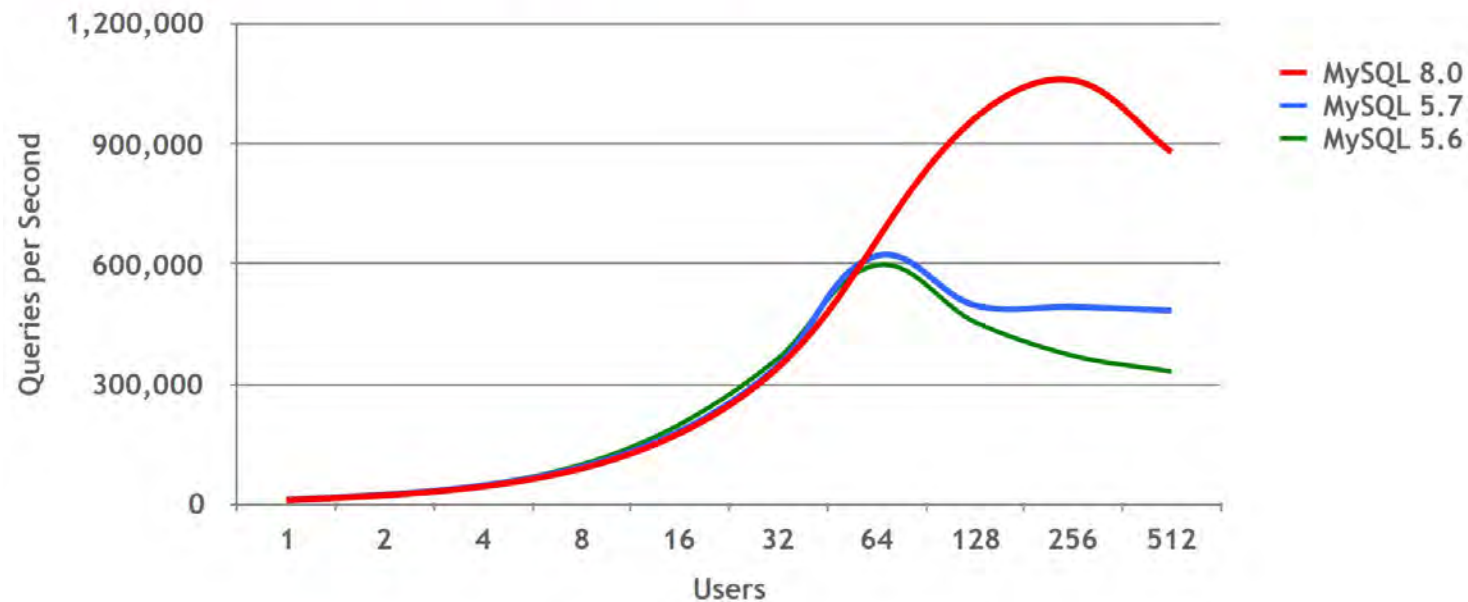
System Configuration :
Oracle Linux 7,
Intel(R) Xeon(R) CPU E7-4860 2.27GHz,
40cores-ht

Plain Better Performance at Scale

Sysbench: OLTP_RO Point-Selects

2.1x Faster than MySQL 5.7

2.8x Faster than MySQL 5.6



Feature Requests

Better Single Thread Performance

Parallel Single Query Processing Please

MySQL 8 for Devs

InnoDB NO WAIT and SKIP LOCKED

New! Better Handling of Hot Row Contention

```
SELECT * FROM tickets  
WHERE id IN (1,2,3,4)  
AND order_id IS NULL  
FOR UPDATE  
NOWAIT;
```

Error
immediately if a
row is already
locked

```
SELECT * FROM tickets  
WHERE id IN (1,2,3,4)  
AND order_id IS NULL  
FOR UPDATE  
SKIP LOCKED;
```

Non
deterministically
skip over locked
rows

Descending Indexes

Descending flag in index definition is no more ignored

Allows efficient handling of ORDER BY A ASC, B DESC queries

Aggregation of Relational into JSON

- **JSON_ARRAYAGG()** and **JSON_OBJECTAGG()**

```
1  mysql> SELECT id, col FROM t1;
2  +-----+-----+
3  | id   | col                               |
4  +-----+-----+
5  |    1 | {"key1": "value1", "key2": "value2"} |
6  |    2 | {"keyA": "valueA", "keyB": "valueB"} |
7  +-----+-----+
8  2 rows in set (0.00 sec)
9
10 mysql> SELECT JSON_OBJECTAGG(id, col) FROM t1;
11 +-----+
12 | JSON_OBJECTAGG(id, col)
13 +-----+
14 | {"1": {"key1": "value1", "key2": "value2"}, "2": {"keyA": "valueA", "keyB": "valueB"}}
15 +-----+
16 1 row in set (0.00 sec)
```

JSON to Table Conversion (Labs)

```
SET @doc=(SELECT doc->"$.properties.amenities"  
FROM seats WHERE id = 28100);
```

```
SELECT * FROM json_table(@doc, "$[*]" columns (  
  id for ordinality, amenity_type varchar(100) path "$.type",  
  distance float path '$.distance_in_meters')  
) AS amenities  
WHERE amenity_type IN ('snacks', 'bar')  
ORDER BY distance;
```

```
+-----+-----+-----+  
| id    | amenity_type | distance |  
+-----+-----+-----+  
|     2 | bar         | 100.538 |  
|     3 | snacks     | 136.647 |  
+-----+-----+-----+  
2 rows in set (0.00 sec)
```

Labs

Better JSON Document Data Extraction

```
1 mysql> CREATE TABLE t1 (doc JSON);
2 Query OK, 0 rows affected (0.01 sec)
3
4 mysql> INSERT INTO t1 VALUES ('[1, 2, 3, 4, 5]');
5 Query OK, 1 row affected (0.00 sec)
6
7 mysql> SELECT doc->"$[1 to 3]" FROM t1;
8 +-----+
9 | doc->"$[1 to 3]" |
10 +-----+
11 | [2, 3, 4]      |
12 +-----+
13 1 row in set (0.00 sec)
14
15 mysql> SELECT doc->"$[last-2]" FROM t1;
16 +-----+
17 | doc->"$[last-2]" |
18 +-----+
19 | 3                |
20 +-----+
21 1 row in set (0.00 sec)
```


Common Table Expression

- Recursive and Non-Recursive

```
1 WITH RECURSIVE cte AS
2 (
3   SELECT category_id, name, 0 AS depth FROM category WHERE parent IS NULL
4   UNION ALL
5   SELECT c.category_id, c.name, cte.depth+1 FROM category c JOIN cte ON
6     cte.category_id=c.parent
7 )
8 SELECT * FROM cte ORDER BY depth;
```

category_id	name	depth
1	ELECTRONICS	0
2	TELEVISIONS	1
6	PORTABLE ELECTRONICS	1
7	MP3 PLAYERS	2
9	CD PLAYERS	2
10	2 WAY RADIOS	2
3	TUBE	2
4	LCD	2
5	PLASMA	2
8	FLASH	3

Window Functions

- Like GROUP BY, But Preserving Rows rather than collapsing them

```
1 mysql> CREATE TABLE t(i INT);
2 mysql> INSERT INTO t VALUES (1),(2),(3),(4);
3
4 mysql> SELECT SUM(i) AS sum FROM t;
5 +-----+
6 | sum |
7 +-----+
8 |  10 |
9 +-----+
10
11 mysql> SELECT i, SUM(i) OVER () AS sum FROM t;
12 +-----+-----+
13 | i   | sum |
14 +-----+-----+
15 |  1 |  10 |
16 |  2 |  10 |
17 |  3 |  10 |
18 |  4 |  10 |
19 +-----+-----+
```

Much Better GIS

- **“Matching or Exceeding PostgreSQL GIS Feature Set”**

5.7

- The world is flat
- The world is infinite
- Axes are unitless
- Axes are orthogonal
- Axis order is irrelevant
- Axis direction is irrelevant

8.0

- The world can be flat or ellipsoidal
- Geographic coordinate systems wrap around
- Axes have units
- Geographic axes are not orthogonal
- Geographic axis order matters
- Axis direction may be relevant

MySQL Document Store

Full Text Indexing

GeoJSON Support

Anyone Using Document Store ?

Summary

MySQL 8 looks like release to be excited about

Has a lot of new features both for Devs and Ops

Thank You!
