MariaDB Observability

Peter Zaitsev,
CEO, Percona

MariaDB Day Brussels
February 2nd, 2020
Brussels, Belgium

PERCONA

Why Observability

Many Non Easily Repeatable Issues in Complex Systems

Resolving and Preventing Issues Requires Internals Understanding

PERCONA

Data Capture

ONGOING DATA CAPTURE (MONITORING)

TEMPORARY DATA CAPTURE (DEBUGGING)

Comprehensive View Needed

You Can’t just use at MariaDB Alone

OS Issues, Hardware Issues are often root cause

Application issues can’t be ignored

Background Load, Noisy Neighbors

MariaDB

PERCONA

MariaDB 10.4 Key Data Sources

SHOW STATUS

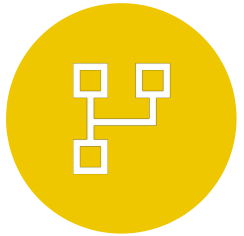

INFORMATION SCHEMA

PERFORMANCE SCHEMA

LOGS

EXPLAIN

OPTIMIZER TRACE

SHOW [GLOBAL] STATUS

Existed Forever

Shows 500+ status variables

Most are counters, some are gauges, some text

Session and Global Scope

Also available as Information Schema Table

Global And Session

```
MariaDB [(none)]> select * from information_schema.global_status
where variable_name='Questions';
+---------------+----------------+
| VARIABLE_NAME | VARIABLE_VALUE |
+---------------+----------------+
| QUESTIONS     | 82457893       |
+---------------+----------------+
1 row in set (0.0001 sec)


MariaDB [(none)]> select * from information_schema.session_status
where variable_name='Questions';
+---------------+----------------+
| VARIABLE_NAME | VARIABLE_VALUE |
+---------------+----------------+
| QUESTIONS     | 153            |
+---------------+----------------+
1 row in set (0.0000 sec)
```

Trust but Verify

```
MariaDB [(none)]> select * from
information_schema.session_status where variable_name like
"%rows%";

+----------------------------+----------------+
| VARIABLE_NAME              | VARIABLE_VALUE |
+----------------------------+----------------+
| INNODB_ROWS_DELETED        | 14553828       |
| INNODB_ROWS_INSERTED       | 15840851       |
| INNODB_ROWS_READ           | 1927758552     |
| INNODB_ROWS_UPDATED        | 29290781       |
…
+----------------------------+----------------+
14 rows in set (0.0000 sec)
```

VMSTAT Like output for MariaDB

```
root@mariadb104:~# mariadb-admin extended -i1 -r | grep Questions
| Questions                                        | 82525506

| Questions                                        | 378

| Questions                                        | 519

| Questions                                        | 591
```

PERCONA

INFORMATION_SCHEMA

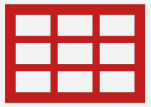

70+ Tables

Some are Schema Related

Others are Performance Statistics

Schema Info

```
MariaDB [information_schema]> select * from tables limit 1 \G
*************************** 1. row ***************************
   TABLE_CATALOG: def
    TABLE_SCHEMA: information_schema
      TABLE_NAME: ALL_PLUGINS
      TABLE_TYPE: SYSTEM VIEW
          ENGINE: Aria
         VERSION: 11
      ROW_FORMAT: Page
      TABLE_ROWS: NULL
  AVG_ROW_LENGTH: 0
     DATA_LENGTH: 8192
 MAX_DATA_LENGTH: 4503599627288576
    INDEX_LENGTH: 8192
       DATA_FREE: 0
  AUTO_INCREMENT: NULL
     CREATE_TIME: 2020-02-01 21:30:36
     UPDATE_TIME: 2020-02-01 21:30:36
      CHECK_TIME: NULL
 TABLE_COLLATION: utf8_general_ci
        CHECKSUM: NULL
  CREATE_OPTIONS: max_rows=18882
   TABLE_COMMENT:
MAX_INDEX_LENGTH: 137438945280
       TEMPORARY: Y
1 row in set (0.007 sec)
```

PERCONA

Innodb Metrics

```
MariaDB [information_schema]> select * from innodb_metrics where
status='enabled' limit 1 \G
*************************** 1. row ***************************
           NAME: lock_deadlocks
      SUBSYSTEM: lock
          COUNT: 0
      MAX_COUNT: NULL
      MIN_COUNT: NULL
      AVG_COUNT: 0
    COUNT_RESET: 0
MAX_COUNT_RESET: NULL
MIN_COUNT_RESET: NULL
AVG_COUNT_RESET: NULL
   TIME_ENABLED: 2020-02-01 20:34:32
  TIME_DISABLED: NULL
   TIME_ELAPSED: 3606
     TIME_RESET: NULL
         STATUS: enabled
           TYPE: counter
        COMMENT: Number of deadlocks
1 row in set (0.001 sec)
```

Word of Caution

While innodb_metrics look similar to SHOW STATUS many of them are NOT enabled by default.

innodb_monitor_enable=all

Innodb Mutexes

```
MariaDB [information_schema]> select * from
innodb_mutexes;
+------+-------------+-------------+----------+
| NAME | CREATE_FILE | CREATE_LINE | OS_WAITS |
+------+-------------+-------------+----------+
|      | log0log.cc  |         578 |        1 |
|      | btr0sea.cc  |         243 |      232 |
+------+-------------+-------------+----------+
2 rows in set (0.002 sec)
```

Extended Processlist

```
MariaDB [information_schema]> select * from processlist limit 1 \G
*************************** 1. row ***************************
ID: 107
USER: root
HOST: localhost
DB: information_schema
COMMAND: Query
TIME: 0
STATE: Filling schema table
INFO: select * from processlist limit 1
TIME_MS: 0.295
STAGE: 0
MAX_STAGE: 0
PROGRESS: 0.000
MEMORY_USED: 106240
MAX_MEMORY_USED: 2214408
EXAMINED_ROWS: 0
QUERY_ID: 17799397
INFO_BINARY: select * from processlist limit 1
TID: 5398
1 row in set (0.001 sec)
```

Performance Shema

- **Some of the Most Avanced Instrumentation Available**
- **50+ Tables**
- **Disabled by Default in MariaDB**
- **Can enable more instrumentation permanently or temporary**
- **Overhead can be high**

Enabling Performance Schema

performance_schema=1

(Restart Required)

Performance Schma Configuration

```
MariaDB [performance_schema]> show tables
like "setup%";
+-------------------------------------+
| Tables_in_performance_schema (setup%) |
+-------------------------------------+
| setup_actors                        |
| setup_consumers                     |
| setup_instruments                   |
| setup_objects                       |
| setup_timers                        |
+-------------------------------------+
5 rows in set (0.001 sec)
```

Actors – Who is going to be profiled

```
MariaDB [performance_schema]> select * from
setup_actors;
+------+------+------+
| HOST | USER | ROLE |
+------+------+------+
| %    | %    | %    |
+------+------+------+
1 row in set (0.000 sec)
```

Consumers – What Summaries are going to be built

```
MariaDB [performance_schema]> select * from setup_consumers;
+--------------------------------+---------+
| NAME                           | ENABLED |
+--------------------------------+---------+
| events_stages_current          | NO      |
| events_stages_history          | NO      |
| events_stages_history_long     | NO      |
| events_statements_current      | YES     |
| events_statements_history      | NO      |
| events_statements_history_long | NO      |
| events_waits_current           | NO      |
| events_waits_history           | NO      |
| events_waits_history_long      | NO      |
| global_instrumentation         | YES     |
| thread_instrumentation         | YES     |
| statements_digest              | YES     |
+--------------------------------+---------+
12 rows in set (0.000 sec)
```

Instrumentation Points

- **700+ Intrumentation Points**
- **250+ Enabled and Timed when Performance Shema is Enabled**

```
MariaDB [performance_schema]> select * from setup_instruments where
enabled='yes' limit 3 \G
*************************** 1. row ***************************
   NAME: wait/io/file/sql/map
ENABLED: YES
  TIMED: YES
*************************** 2. row ***************************
   NAME: wait/io/file/sql/binlog
ENABLED: YES
  TIMED: YES
*************************** 3. row ***************************
   NAME: wait/io/file/sql/binlog_index
ENABLED: YES
  TIMED: YES
3 rows in set (0.000 sec)
```

What Objects do we want to Instrument

```
MariaDB [performance_schema]> select * from setup_objects;
+-------------+--------------------+-------------+---------+-------+
| OBJECT_TYPE | OBJECT_SCHEMA      | OBJECT_NAME | ENABLED | TIMED |
+-------------+--------------------+-------------+---------+-------+
| TABLE       | mysql              | %           | NO      | NO    |
| TABLE       | performance_schema | %           | NO      | NO    |
| TABLE       | information_schema | %           | NO      | NO    |
| TABLE       | %                  | %           | YES     | YES   |
+-------------+--------------------+-------------+---------+-------+
4 rows in set (0.000 sec)
```

Example of Performance Schema

```
MariaDB [performance_schema]> select * from
events_statements_current  limit 1 \G

*************************** 1. row ***************************

              THREAD_ID: 53

               EVENT_ID: 226

           END_EVENT_ID: NULL

             EVENT_NAME: statement/sql/select

                 SOURCE: mysqld.cc:1179

            TIMER_START: 1484013787233000

              TIMER_END: 1484014492287000

             TIMER_WAIT: 705054000

              LOCK_TIME: 145000000

                SQL_TEXT: select * from events_statements_current
limit 1

                  DIGEST: NULL

             DIGEST_TEXT: NULL

          CURRENT_SCHEMA: performance_schema

             OBJECT_TYPE: NULL

           OBJECT_SCHEMA: NULL

             OBJECT_NAME: NULL

  OBJECT_INSTANCE_BEGIN: NULL

             MYSQL_ERRNO: 0

       RETURNED_SQLSTATE: NULL

            MESSAGE_TEXT: NULL

                  ERRORS: 0

                WARNINGS: 0

           ROWS_AFFECTED: 0

               ROWS_SENT: 0

           ROWS_EXAMINED: 0

CREATED_TMP_DISK_TABLES: 0

     CREATED_TMP_TABLES: 0

       SELECT_FULL_JOIN: 0

SELECT_FULL_RANGE_JOIN: 0

         SELECT_RANGE: 0

    SELECT_RANGE_CHECK: 0

           SELECT_SCAN: 1

   SORT_MERGE_PASSES: 0

         SORT_RANGE: 0

          SORT_ROWS: 0

          SORT_SCAN: 0

      NO_INDEX_USED: 1

 NO_GOOD_INDEX_USED: 0

   NESTING_EVENT_ID: NULL

 NESTING_EVENT_TYPE: NULL

1 row in set (0.001 sec)
```


PERCONA

IO Statistics

```
MariaDB [performance_schema]> select * from
file_summary_by_instance limit 5,1 \G

*************************** 1. row
***************************

               FILE_NAME: /var/lib/mysql/ibdata1

              EVENT_NAME:
wait/io/file/innodb/innodb_data_file

    OBJECT_INSTANCE_BEGIN: 140586480369856

              COUNT_STAR: 161

          SUM_TIMER_WAIT: 34773929190

          MIN_TIMER_WAIT: 13463450

          AVG_TIMER_WAIT: 215987135

          MAX_TIMER_WAIT: 2412890480

              COUNT_READ: 158

          SUM_TIMER_READ: 34702257590

          MIN_TIMER_READ: 76386310

          AVG_TIMER_READ: 219634415

          MAX_TIMER_READ: 2412890480

 SUM_NUMBER_OF_BYTES_READ: 4702208

             COUNT_WRITE: 0

         SUM_TIMER_WRITE: 0

         MIN_TIMER_WRITE: 0

         AVG_TIMER_WRITE: 0

         MAX_TIMER_WRITE: 0

SUM_NUMBER_OF_BYTES_WRITE: 0

              COUNT_MISC: 3

          SUM_TIMER_MISC: 71671600

          MIN_TIMER_MISC: 13463450

          AVG_TIMER_MISC: 23890230

          MAX_TIMER_MISC: 30040010

1 row in set (0.000 sec)
```

Logs

Error Log

General Query Log

Slow Query Log

Slow Query Log [Default]

```
# User@Host: root[root] @ localhost []
# Thread_id: 111  Schema: sbtest  QC_hit: No
# Query_time: 0.000356  Lock_time: 0.000149  Rows_sent: 1  Rows_examined: 1
# Rows_affected: 0  Bytes_sent: 190
SET timestamp=1580596196;
SELECT c FROM sbtest1 WHERE id=767650;
```

Slow Query Log [With Explain]

```
# Time: 200201 22:32:37

# User@Host: root[root] @ localhost []

# Thread_id: 113  Schema: sbtest  QC_hit: No

# Query_time: 0.000220  Lock_time: 0.000091  Rows_sent: 1  Rows_examined: 1

# Rows_affected: 0  Bytes_sent: 190

#

# explain: id   select_type     table   type    possible_keys   key     key_len ref
  rows    r_rowsfiltered r_filtered      Extra

# explain: 1    SIMPLE  sbtest1 const   PRIMARY PRIMARY 4       const   1       NULL
  100.00  NULL

#

SET timestamp=1580596357;

SELECT c FROM sbtest1 WHERE id=101985;
```


PERCONA

EXPLAIN

- **Understand Query Execution Plan**
- **Essential Skill for Developers and DBAs**
- **Multiple Output Formats**
- **Can get EXPLAIN plan for actual running query**

EXPLAIN

```
MariaDB [sbtest]> explain  select count(*) from sbtest1 s1,sbtest1 s2;

+------+-------------+-------+-------+---------------+------+---------+------+--------
+-----------------------------------------------+

| id   | select_type | table | type  | possible_keys | key  | key_len | ref  | rows   | Extra
                  |

+------+-------------+-------+-------+---------------+------+---------+------+--------
+-----------------------------------------------+

|    1 | SIMPLE      | s1    | index | NULL          | k_1  | 4       | NULL | 986499 | Using index
                |

|    1 | SIMPLE      | s2    | index | NULL          | k_1  | 4       | NULL | 986499 | Using index; Using join
buffer (flat, BNL join) |

+------+-------------+-------+-------+---------------+------+---------+------+--------
+-----------------------------------------------+

2 rows in set (0.001 sec)
```

PERCONA

Advanced Explain

- **EXPLAIN FORMAT=JSON**
- **SHOW EXPLAIN FOR <CONNECTION_ID>**
- **ANALYZE**

ANALYZE example

```
MariaDB [sbtest]> analyze select count(*)
from sbtest1 s1,sbtest1 s2 where
s1.id<100 \G

*************************** 1. row
***************************

           id: 1

  select_type: SIMPLE

        table: s1

         type: range

possible_keys: PRIMARY

          key: PRIMARY

      key_len: 4

          ref: NULL

         rows: 99

       r_rows: 99.00

     filtered: 100.00

   r_filtered: 100.00

        Extra: Using where; Using index
*************************** 2. row
***************************

           id: 1

  select_type: SIMPLE

        table: s2

         type: index

possible_keys: NULL

          key: k_1

      key_len: 4

          ref: NULL

         rows: 986499

       r_rows: 1000000.00

     filtered: 100.00

   r_filtered: 100.00

        Extra: Using index; Using join
buffer (flat, BNL join)
```

Looking for Open Source Database Observability Solution ?

Open Source Database Focused Observability Solution from Percona

100% Free and Open Source

http://per.co.na/PMM

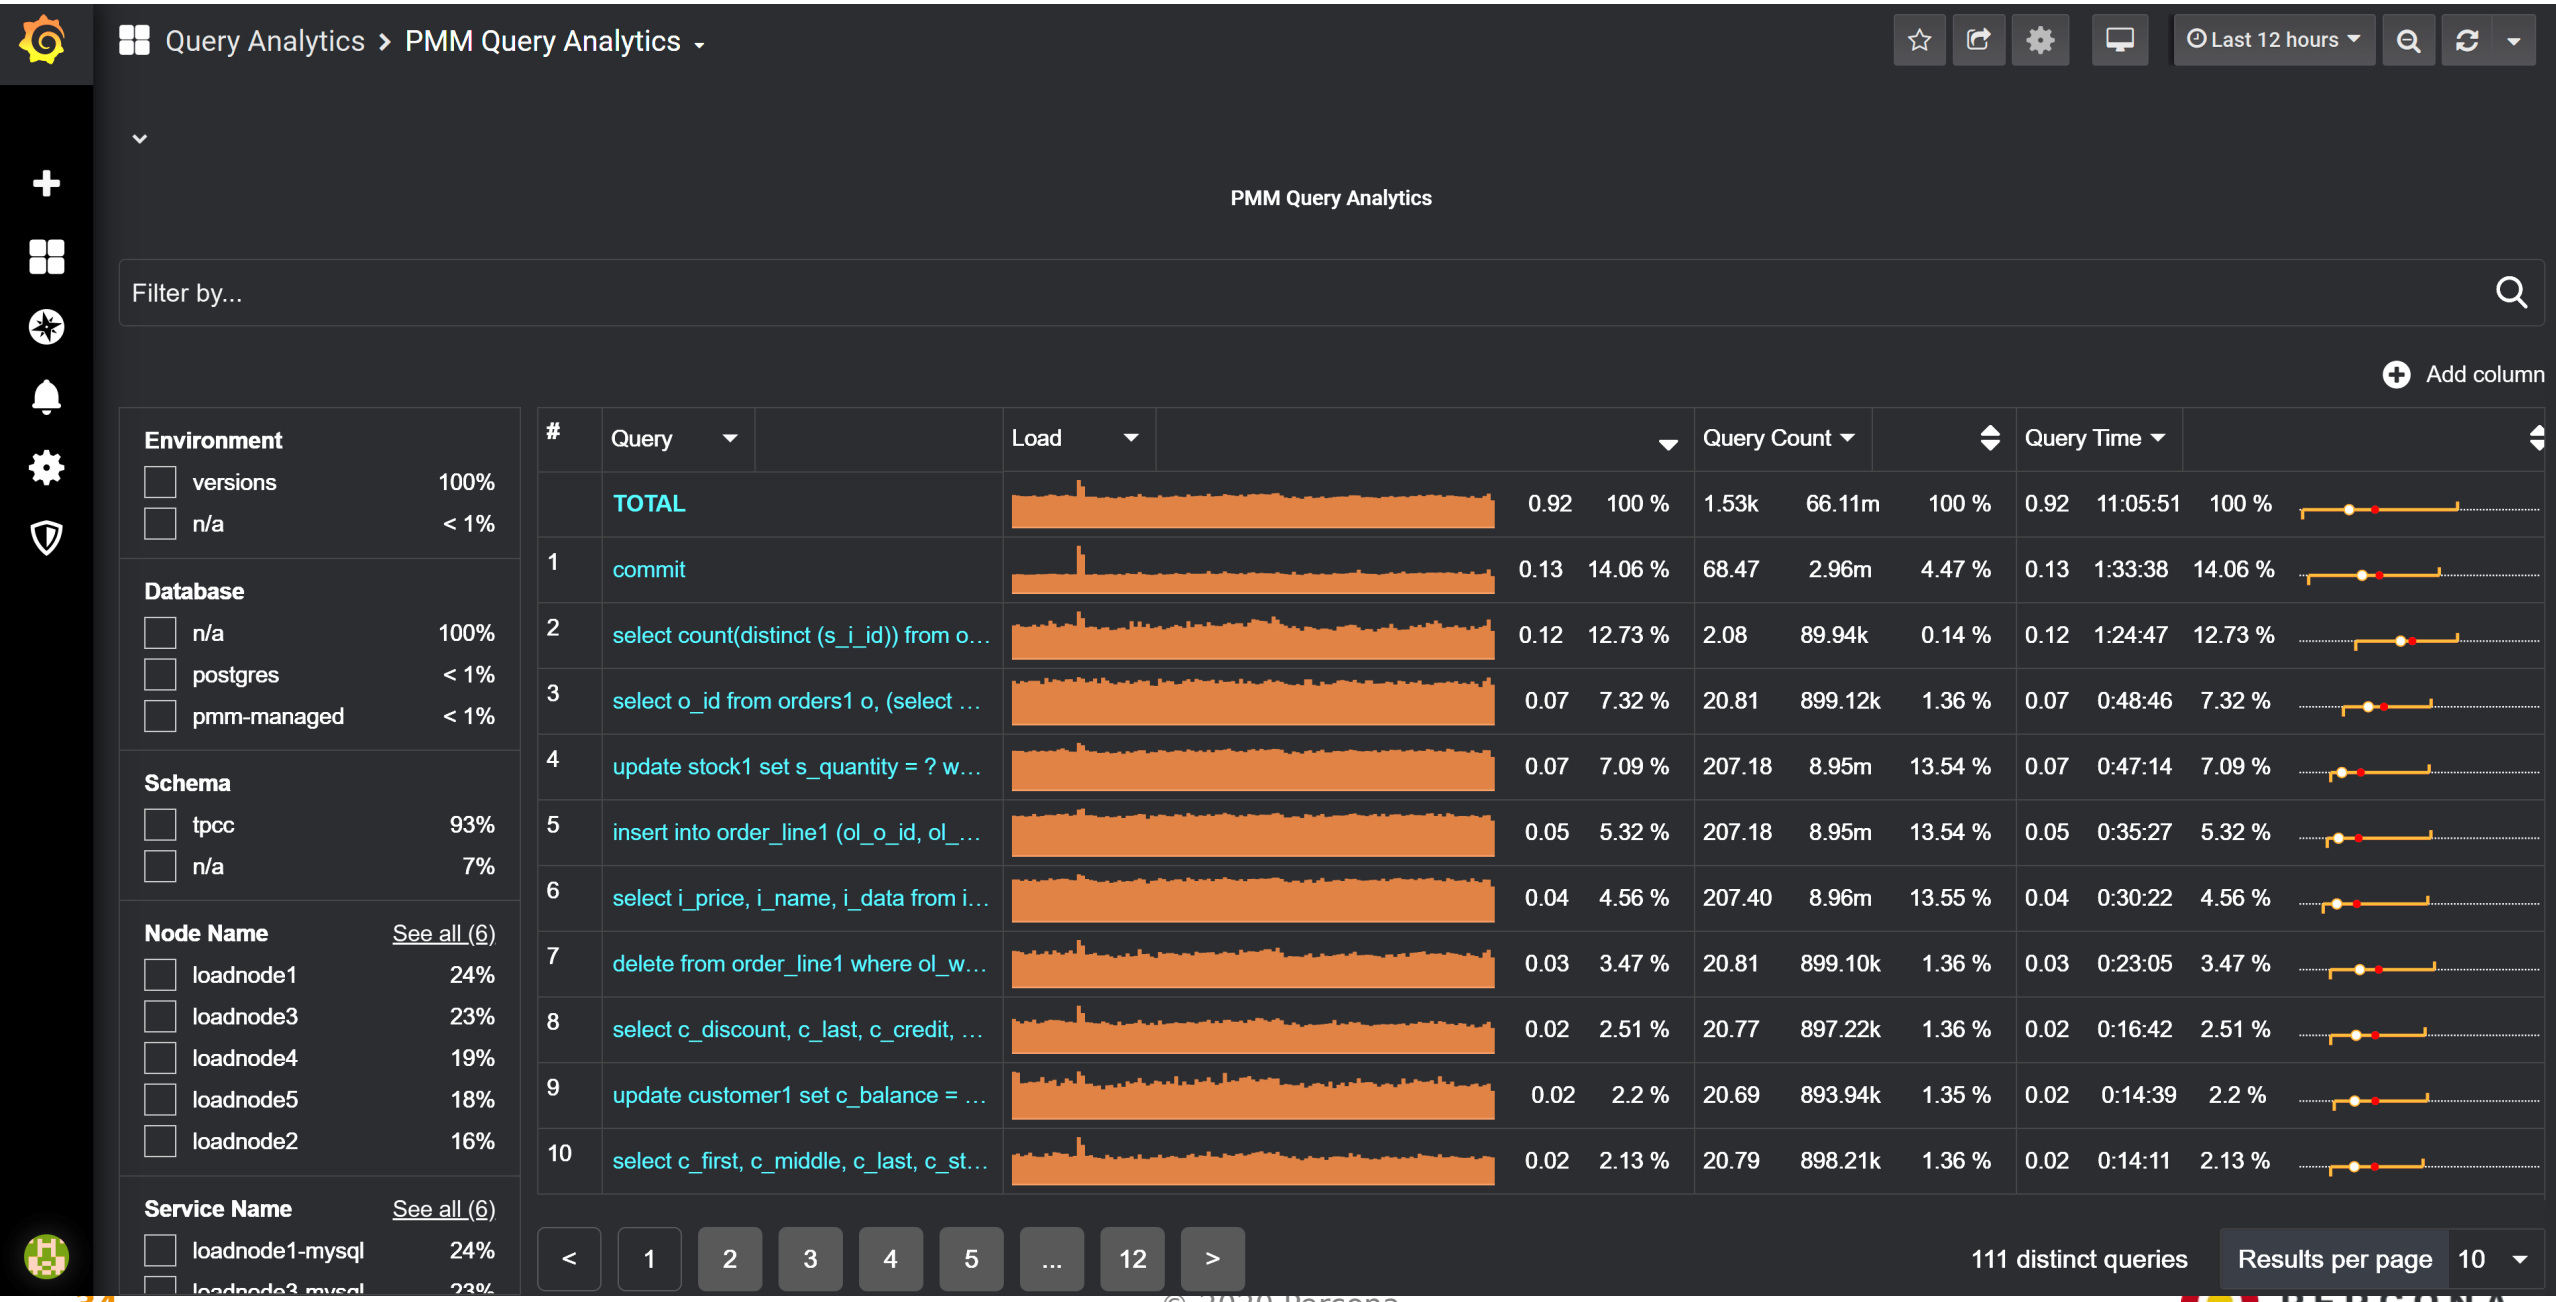

Thank you, Let’s Connect!

https://www.linkedin.com/in/peterzaitsev/

https://twitter.com/PeterZaitsev