The Future of MySQL Quality Assurance: Introducing pquery

Ramesh Sivaraman
Jr. QA Engineer

Roel Van de Paar
Sr. QA Lead

MYSQL WORLDWIDE CONFERENCE & EXPO

PERCONA LIVE

QA: The Past

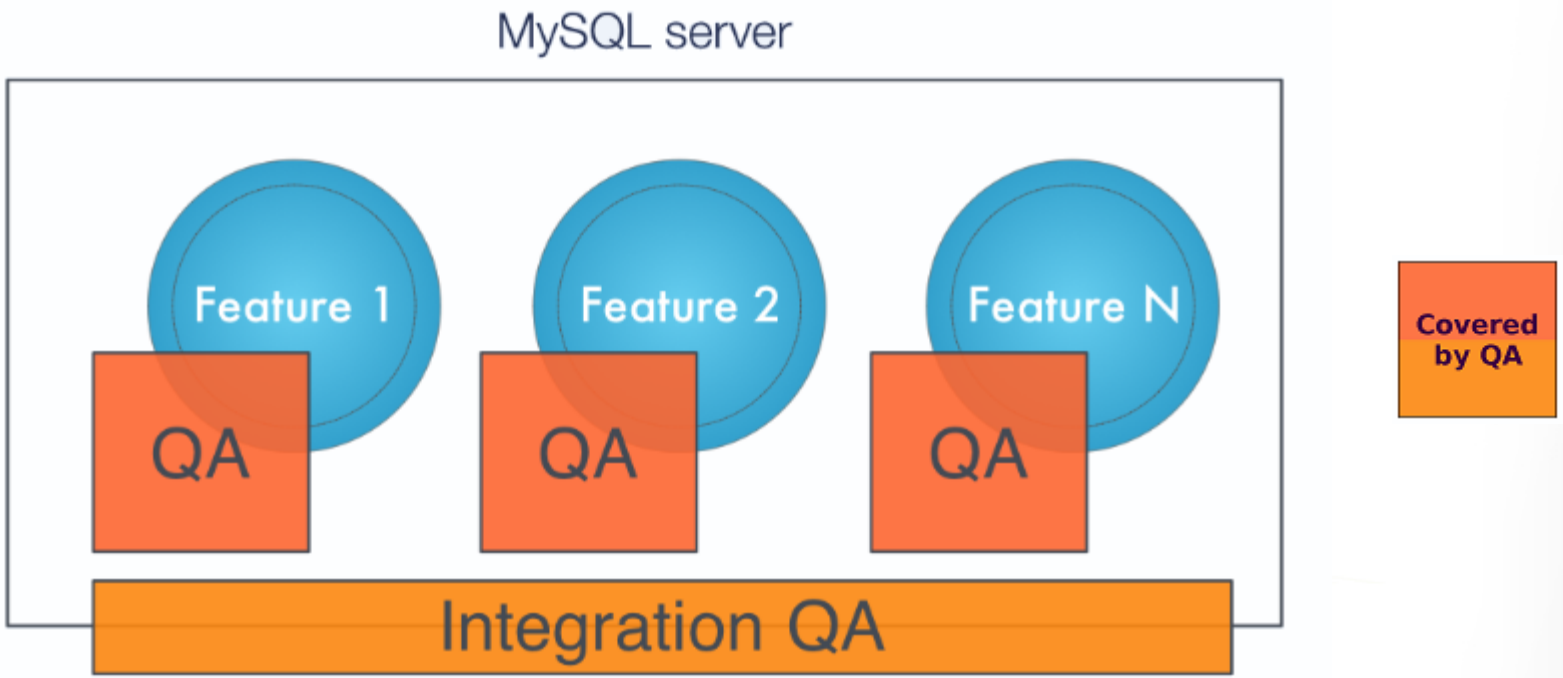
MySQL server
Feature 1
Feature 2
Feature N
QA
QA
QA
Integration QA
Covered by QA

QA: The Future

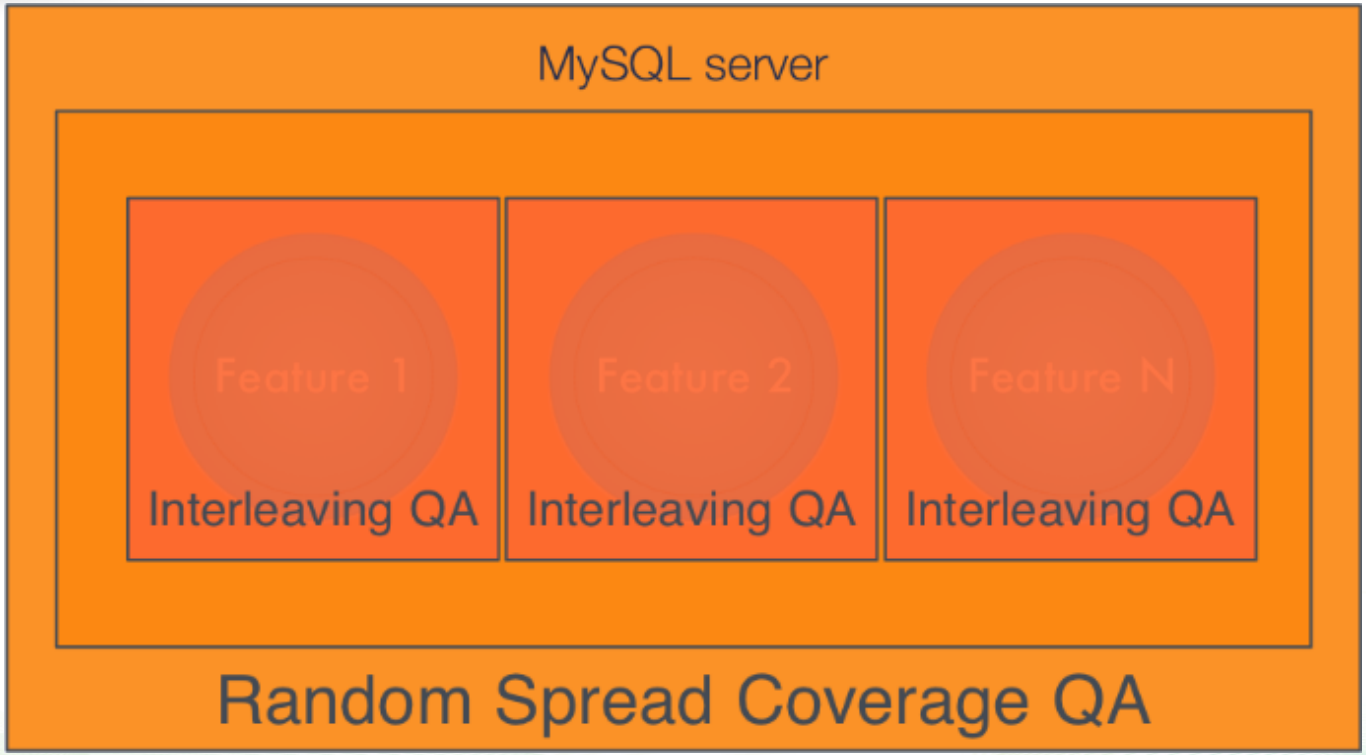

```
CREATE EVENT e_x1 ON SCHEDULE EVERY 1 SECOND DO DROP DATABASE db_x;
SELECT * FROM t2i FORCE INDEX(t2a) ORDER BY a;
update t1 set i= 10;
DROP TABLE t572;
INSERT INTO t1 VALUES (5,'yyy'), (6,'yyy');
set global init_connect="set @c=1";
select * from bigfailure where afield = 18446744071710965856+1;
select * from (select 2 from DUAL) b;
create event e_55 on schedule at 99990101000000 do drop table t;
select convert_tz('2003-12-31 04:00:00', NULL, 'UTC');
ALTER TABLE t4 CHANGE c1 c1 DOUBLE UNSIGNED NOT NULL;
select char_length('\n\t\r\b\0\_\%\\');
SET @@character_set_server = UTF8;
CREATE TABLE test.t1 (a VARCHAR(255), PRIMARY KEY(a));
SET GLOBAL innodb_cleaner_free_list_lwm = @start_value;
flush status;
SELECT VIEW_DEFINITION FROM INFORMATION_SCHEMA.VIEWS where TABLE_NAME = 'v1';
SELECT * FROM t1 WHERE MATCH (j) AGAINST ('innodb');
select count(*) from mysql.host;
DROP TABLE bug56680_2;
SELECT * FROM t14169459_1;
SHOW CHARACTER SET LIKE 'latin1%';
INSERT into gis_geometry SELECT * FROM gis_line;
select SQL_BIG_RESULT a,count(b), sum(b), avg(b), std(b), min(b), max(b), bit_and(b), bit_or(b), bit_xor(b) from t1 group by a;
DROP PROCEDURE IF EXISTS sp40;
SELECT * FROM t2 WHERE c2 BETWEEN '-838:59:59' AND '10:00:00' ORDER BY c1,c2 DESC;
INSERT INTO t783 VALUES(1);
CREATE TABLE t1 (a enum ('a','b','c')) CHARACTER SET utf16le;
purge master logs to 'master-bin.000002';
CREATE TABLE t1 (s1 VARCHAR (60) CHARACTER SET UTF8 COLLATE UTF8_UNICODE_520_CI) ENGINE = MyISAM;
call mtr.add_suppression("Info table is not ready to be used.");
DROP PROCEDURE IF EXISTS t_ps_proc;
CREATE TABLE m3(c1 BIT NULL, c2 VARCHAR(25) NOT NULL, c3 INTEGER(4) NULL, c4 VARCHAR(15) NOT NULL PRIMARY KEY, c5 BIT NOT NULL UNIQUE KEY,c6 DEC(10,8) NOT N
ULL DEFAULT 3.141592);
SELECT * FROM t2 WHERE c1 BETWEEN 0 AND 16777215 ORDER BY c1,c6;
CREATE TABLE m3(c1 YEAR NULL, c2 CHAR(25) NOT NULL, c3 TINYINT(4) NULL, c4 CHAR(15) NOT NULL PRIMARY KEY, c5 YEAR NOT NULL UNIQUE KEY,c6 DECIMAL(10,8) NOT N
ULL DEFAULT 3.141592);
insert into server_stopword values("root"),("properly");
CREATE TABLE t1 (a_dec DECIMAL(-1,1));
SELECT argument FROM mysql.general_log WHERE argument = 'UPDATE t1 SET a = \'aa1\' WHERE a = \'aa1-updated\'';
SELECT * FROM t1 WHERE c1 = '00:00:00' ORDER BY c1 DESC LIMIT 2;
SET GLOBAL innodb_sched_priority_io=1e1;
select char(53647 using utf8mb4);
SET @start_global_value = @@global.rpl_semi_sync_slave_trace_level;
CREATE TABLE t2(c1 BIT(16), c2 BIT(16), c3 BIT(16));
UPDATE t1 SET c1 = NOW() WHERE c1 >= ADDTIME(NOW(),'2 02:01:01');
UPDATE t2 SET a = 10;
DELETE FROM t1 WHERE a = 2 LIMIT 9;
SELECT EXTRACT(MICROSECOND FROM CAST(20010101235959.456 AS DATETIME(6)));
grant select on `mysqltest`.* to mysqltest_1@localhost;
update t1 set name='U+0419 CYRILLIC CAPITAL LETTER SHORT I' where ujis=0xA7AB;
CREATE TABLE t5(c1 DOUBLE(5,4) NOT NULL);
DELETE FROM user WHERE host='localhost' AND user='@#@';
SELECT _latin2'mysql' COLLATE latin2_general_ci='MySql';
SELECT @@session.max_heap_table_size = 16777216;
SELECT @start_binlog_cache_size;
```

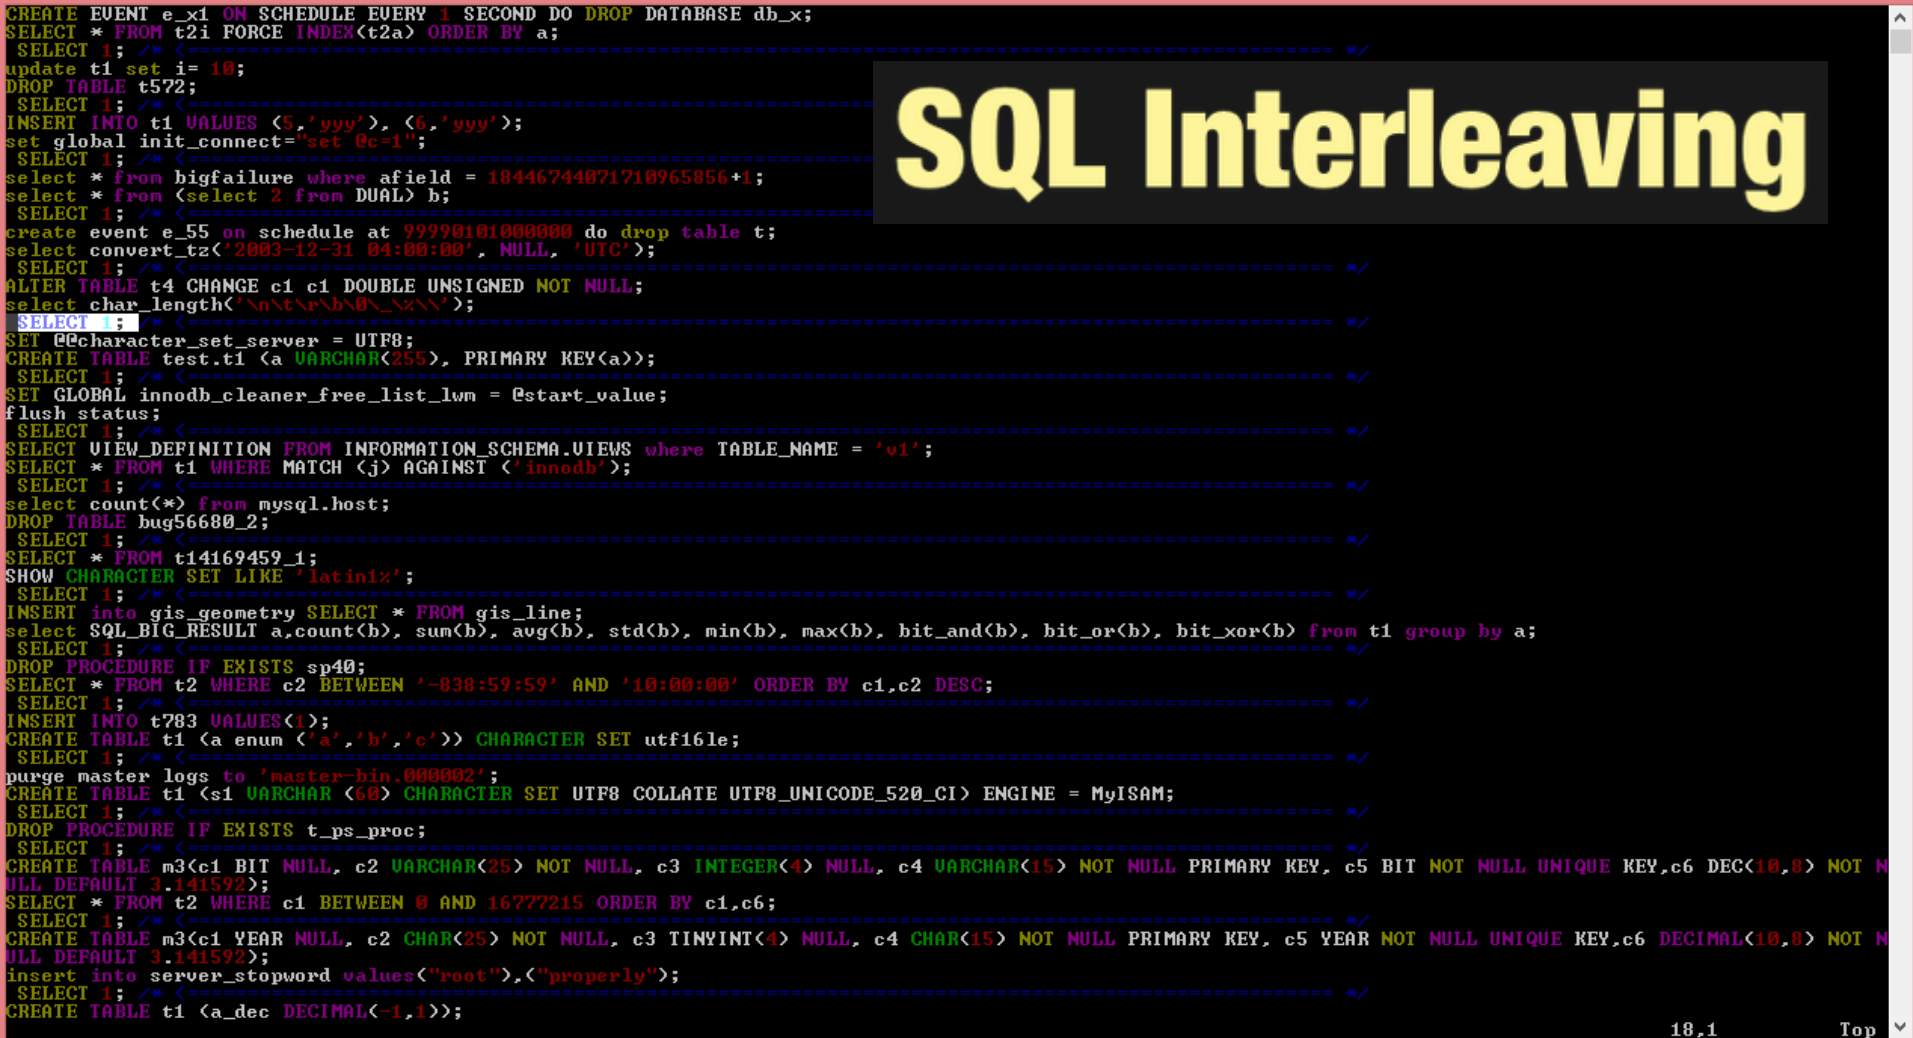


```
CREATE EVENT e_x1 ON SCHEDULE EVERY 1 SECOND DO DROP DATABASE db_x;
SELECT * FROM t2i FORCE INDEX(t2a) ORDER BY a;
 SELECT 1; /* <============================================================================== */
update t1 set i= 10;
DROP TABLE t572;
 SELECT 1; /* <============================================================================== */
INSERT INTO t1 VALUES (5,'yyy'), (6,'yyy');
set global init_connect="set @c=1";
 SELECT 1; /* <============================================================================== */
select * from bigfailure where afield = 18446744071710965856+1;
select * from (select 2 from DUAL) b;
 SELECT 1; /* <============================================================================== */
create event e_55 on schedule at 99990101000000 do drop table t;
select convert_tz('2003-12-31 04:00:00', NULL, 'UTC');
 SELECT 1; /* <============================================================================== */
ALTER TABLE t4 CHANGE c1 c1 DOUBLE UNSIGNED NOT NULL;
select char_length('\n\t\r\b\0\_\%\\');
 SELECT 1; /* <============================================================================== */
SET @@character_set_server = UTF8;
CREATE TABLE test.t1 (a VARCHAR(255), PRIMARY KEY(a));
 SELECT 1; /* <============================================================================== */
SET GLOBAL innodb_cleaner_free_list_lwm = @start_value;
flush status;
 SELECT 1; /* <============================================================================== */
SELECT VIEW_DEFINITION FROM INFORMATION_SCHEMA.VIEWS where TABLE_NAME = 'v1';
SELECT * FROM t1 WHERE MATCH (j) AGAINST ('innodb');
 SELECT 1; /* <============================================================================== */
select count(*) from mysql.host;
DROP TABLE bug56680_2;
 SELECT 1; /* <============================================================================== */
SELECT * FROM t14169459_1;
SHOW CHARACTER SET LIKE 'latin1%';
 SELECT 1; /* <============================================================================== */
INSERT into gis_geometry SELECT * FROM gis_line;
select SQL_BIG_RESULT a,count(b), sum(b), avg(b), std(b), min(b), max(b), bit_and(b), bit_or(b), bit_xor(b) from t1 group by a;
 SELECT 1; /* <============================================================================== */
DROP PROCEDURE IF EXISTS sp40;
SELECT * FROM t2 WHERE c2 BETWEEN '-838:59:59' AND '10:00:00' ORDER BY c1,c2 DESC;
 SELECT 1; /* <============================================================================== */
INSERT INTO t783 VALUES(1);
CREATE TABLE t1 (a enum ('a','b','c')) CHARACTER SET utf16le;
 SELECT 1; /* <============================================================================== */
purge master logs to 'master-bin.000002';
CREATE TABLE t1 (s1 VARCHAR (60) CHARACTER SET UTF8 COLLATE UTF8_UNICODE_520_CI) ENGINE = MyISAM;
 SELECT 1; /* <============================================================================== */
DROP PROCEDURE IF EXISTS t_ps_proc;
 SELECT 1; /* <============================================================================== */
CREATE TABLE m3(c1 BIT NULL, c2 VARCHAR(25) NOT NULL, c3 INTEGER(4) NULL, c4 VARCHAR(15) NOT NULL PRIMARY KEY, c5 BIT NOT NULL UNIQUE KEY,c6 DEC(10,8) NOT N
ULL DEFAULT 3.141592);
SELECT * FROM t2 WHERE c1 BETWEEN 0 AND 16777215 ORDER BY c1,c6;
 SELECT 1; /* <============================================================================== */
CREATE TABLE m3(c1 YEAR NULL, c2 CHAR(25) NOT NULL, c3 TINYINT(4) NULL, c4 CHAR(15) NOT NULL PRIMARY KEY, c5 YEAR NOT NULL UNIQUE KEY,c6 DECIMAL(10,8) NOT N
ULL DEFAULT 3.141592);
insert into server_stopword values("root"),("properly");
 SELECT 1; /* <============================================================================== */
CREATE TABLE t1 (a_dec DECIMAL(-1,1));
```

Random Spread Coverage Testing

Sad QA Engineer: Mr. Thinkalot

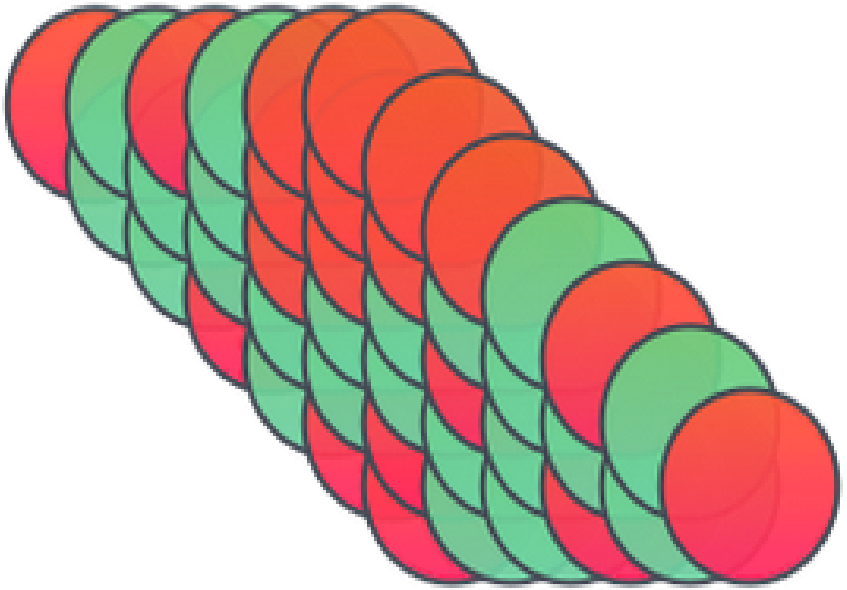

Random Spread Coverage Testing

Happy QA Engineer: Mr. Nuclear

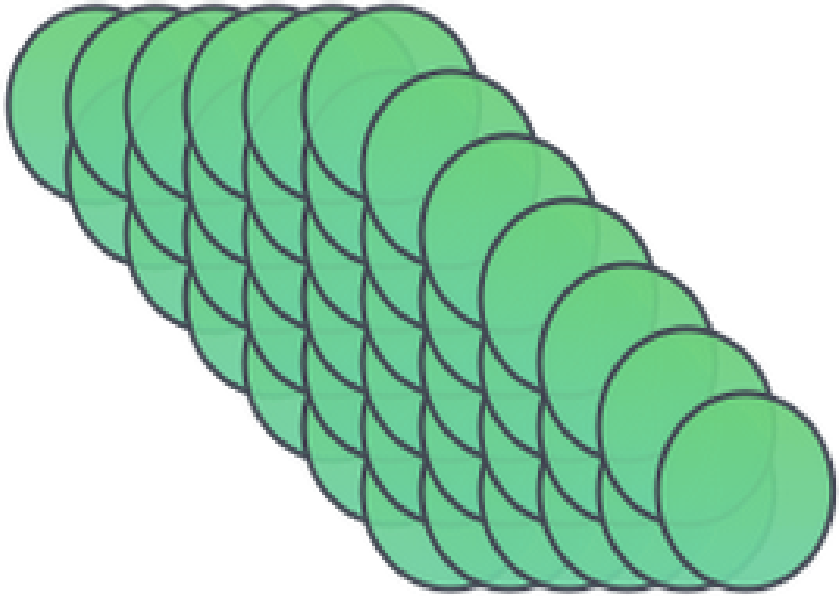

pquery-run.sh

```
[roel@localhost percona-qa]$ ./pquery-run.sh
[11:07:57] [0] Workdir: /sda/973294 | Rundir: /dev/shm/973294 | Basedir: /sda/mysql-5.6.23-linux-x86_64-debug
[11:07:57] [0] mysqld Start Timeout: 60 | Client Threads: 1 | Queries/Thread: 25000 | Trials: 100000 | Save coredump/valgrind is
sue trials only: TRUE
[11:07:57] [0] Valgrind run: FALSE | pquery timeout: 15 | SQL file used: /sda/main-ms-ps-md.sql
[11:07:57] [0] pquery Binary: /home/roel/percona-qa/pquery/pquery
[11:07:57] [0] MYSAFE: --no-defaults --event-scheduler=ON --maximum-bulk_insert_buffer_size=1M --maximum-join_buffer_size=1M --m
aximum-max_heap_table_size=1M --maximum-max_join_size=1M --maximum-myisam_max_sort_file_size=1M --maximum-myisam_mmap_size=1M --
maximum-myisam_sort_buffer_size=1M --maximum-optimizer_trace_max_mem_size=1M --maximum-preload_buffer_size=1M --maximum-query_al
loc_block_size=1M --maximum-query_prealloc_size=1M --maximum-range_alloc_block_size=1M --maximum-read_buffer_size=1M --maximum-r
ead_rnd_buffer_size=1M --maximum-sort_buffer_size=1M --maximum-tmp_table_size=1M --maximum-transaction_alloc_block_size=1M --max
imum-transaction_prealloc_size=1M --log-output=none --sql_mode=ONLY_FULL_GROUP_BY
[11:07:57] [0] Making a copy of the pquery binary used to /sda/973294 (handy for later re-runs/reference etc.)
[11:07:57] [0] Archiving a copy of this script (/home/roel/percona-qa/./pquery-run.sh) in the workdir /sda/973294 for later refe
rence & adding pquery- prefix to avoid pquery-prep-run issues...
[11:07:58] [0] Archiving a copy of the SQL input file used (/sda/main-ms-ps-md.sql) in the workdir /sda/973294 for later referen
ce...
[11:07:58] [0] Generating datadir template (using mysql_install_db)...
[11:08:05] [0] Making a copy of mysqld used to /sda/973294/mysqld (handy for core file analysis and manual bundle creation)...
[11:08:05] [0] Storing a copy of ldd files for mysqld in same directory also...
[11:08:05] [0] Starting pquery testing iterations...
[11:08:05] [0] ====== TRIAL #1 ======
[11:08:05] [0] Ensuring there are no relevant servers running...
[11:08:05] [0] Clearing rundir...
[11:08:06] [0] Generating new trial workdir /dev/shm/973294/1...
[11:08:06] [0] Copying datadir from template...
[11:08:06] [0] Starting mysqld. Error log is stored at /dev/shm/973294/1/log/master.err
[11:08:06] [0] Waiting for mysqld (pid: 12258) to fully start...
[11:08:08] [0] Server started ok. Client:  /sda/mysql-5.6.23-linux-x86_64-debug/bin/mysql -uroot -S/dev/shm/973294/1/socket.sock
[11:08:08] [0] Starting pquery (log stored in /dev/shm/973294/1/pquery.log)...
[11:08:08] [0] pquery running (Max duration: 15s)...
./pquery-run.sh: line 299: 12258 Aborted                 (core dumped) $CMD > ${RUNDIR}/${TRIAL}/log/master.err 2>&1
[11:08:11] [0] Cleaning up...
[11:08:17] [0] pquery summary: 1779/2174 queries failed (18.17% were successful)
[11:08:17] [0] mysqld core detected at /dev/shm/973294/1/data/core.12258.1000.1000.6.1428973691.mysqld
[11:08:17] [0] Bug found (as per error log): t1->result_range
[11:08:18] [1] Copying rundir from /dev/shm/973294/1 to /sda/973294/1
[11:08:19] [1] ====== TRIAL #2 ======
[11:08:19] [1] Ensuring there are no relevant servers running...
[11:08:19] [1] Clearing rundir...
[11:08:19] [1] Generating new trial workdir /dev/shm/973294/2...
[11:08:19] [1] Copying datadir from template...
[11:08:19] [1] Starting mysqld. Error log is stored at /dev/shm/973294/2/log/master.err
[11:08:19] [1] Waiting for mysqld (pid: 12396) to fully start...
[11:08:22] [1] Server started ok. Client:  /sda/mysql-5.6.23-linux-x86_64-debug/bin/mysql -uroot -S/dev/shm/973294/2/socket.sock
[11:08:22] [1] Starting pquery (log stored in /dev/shm/973294/2/pquery.log)...
```

pquery-prep-red.sh

```
MYEXTRA:
MYSAFE: --no-defaults --event-scheduler=ON --maximum-bulk_insert_buffer_size=1M --maximum-join_buffer_size=1M --maximum-max_heap
_table_size=1M --maximum-max_join_size=1M --maximum-myisam_max_sort_file_size=1M --maximum-myisam_mmap_size=1M --maximum-myisam_
sort_buffer_size=1M --maximum-optimizer_trace_max_mem_size=1M --maximum-preload_buffer_size=1M --maximum-query_alloc_block_size=
1M --maximum-query_prealloc_size=1M --maximum-range_alloc_block_size=1M --maximum-read_buffer_size=1M --maximum-read_rnd_buffer_
size=1M --maximum-sort_buffer_size=1M --maximum-tmp_table_size=1M --maximum-transaction_alloc_block_size=1M --maximum-transactio
n_prealloc_size=1M --log-output=none --sql_mode=ONLY_FULL_GROUP_BY
==========================================================================================================================
======================
If you agree that these look correct, hit enter twice. If something looks wrong, press CTRL+C to abort.
(Note that MYSAFE was only introduced around 13-10-2014, so it would be empty for earlier runs before this date
(To learn more, you may want to read some of the info/code in this script (pquery-prep-red.sh) in the 'Current location checks'
var checking section.)
Btw, note that only MYEXTRA is used by reducer, so MYSAFE string will be merged into MYEXTRA for the resulting reducer<nr>.sh sc
ripts.
==========================================================================================================================
======================
Press ENTER or CTRL-C now... 1...
Press ENTER or CTRL-C now... 2...
Ok, using MYEXTRA/MYSAFE as listed above
========== Processing pquery trial 1
* Obtaining quer(y)(ies) from the trial's coredump (core: ./1/data/core.12258.1000.1000.6.1428973691.mysqld)
  > 1 quer(y)(ies) added 3x to the SQL trace
* Obtaining quer(y)(ies) from the trial's mysqld error log (if any)
  > 1 quer(y)(ies) added 3x to the SQL trace
TEXT variable set to: "t1->result_range"

Done!! Start reducer scripts like this: './reducerTRIAL.sh' where TRIAL stands for the trial number you would like to reduce
Both reducer and the SQL trace file have been pre-prepped with all the crashing queries and settings, ready for you to use witho
ut further options!
```

pquery-results.sh

```
[roel@localhost 973294]$ ~/percona-qa/pquery-results.sh
================ Sorted unique issue strings (Approx query_thread-0.sql_out trials executed, 1 remaining reducer scripts)
t1->result_range                                (Seen   1 times: reducers 1)
================
[roel@localhost 973294]$ cd ..
[roel@localhost sda]$ cd MS-159219/
/sda/MS-159219
[roel@localhost MS-159219]$ ~/percona-qa/pquery-results.sh
ls: cannot access [0-9]*: No such file or directory
================ Sorted unique issue strings (Approx  trials executed, 37 remaining reducer scripts)
btr0cur.cc line 700                             (Seen   1 times: reducers 21)
fsp0fsp.cc line 1887                            (Seen   3 times: reducers 22,42,62)
fts0fts.cc line 2228                            (Seen   1 times: reducers 167)
.length % 4                                     (Seen   5 times: reducers 109,138,14,175,179)
.slen % 2                                       (Seen   5 times: reducers 101,119,134,153,170)
.slen % 4                                       (Seen   3 times: reducers 158,34,99)
t1->result_range                                (Seen  10 times: reducers 160,161,164,28,30,40,53,57,72,73)
.tlen % 2                                       (Seen   3 times: reducers 176,64,91)
.tlen % 4                                       (Seen   2 times: reducers 103,132)
Z20commit_try_norebuildP18Alter_inplace_infoP23ha_innobase_inplace_ctxPK5TABLEP5trx_tPKc(Seen   1 times: reducers 94)
Z34convert_charset_partition_constantP4ItemPK15charset_info_st(Seen   1 times: reducers 8)
ZN8Protocol13end_statementEv                    (Seen   2 times: reducers 106,2)
================
[roel@localhost MS-159219]$ _
```

reducer.sh features (available from lp:randgen)

- Fully automatic ultra-fast testcase reduction
- Near-100% issue reproducibility
- Compatible with sporadic issues
- True multi-threaded testcase reduction

Result: ~2-8 lines of SQL per bug

reducer{nr}.sh

```
[roel@localhost 973294]$ ./reducer1.sh
2015-04-14 11:22:40 [Init] Workdir: /tmp/1428974560
2015-04-14 11:22:40 [Init] Temporary storage directory (TMP environment variable) set to /tmp/1428974560/tmp
2015-04-14 11:22:40 [Init] Input file: /sda/973294/1/pquery_thread-0.sql
2015-04-14 11:22:40 [Init] Server (When MULTI mode is not active): /sda/mysql-5.6.23-linux-x86_64-debug/bin/mysqld (as roel)
2015-04-14 11:22:40 [Init] Client (When MULTI mode is not active): /sda/mysql-5.6.23-linux-x86_64-debug/bin/mysql -uroot -S/tmp/
1428974560/socket.sock
2015-04-14 11:22:40 [Init] Querytimeout: 90 seconds (ensure this is at least 1.5x what was set in RQG using the --querytimeout o
ption)
2015-04-14 11:22:40 [Init] Passing the following additional options to mysqld:  --event-scheduler=ON --maximum-bulk_insert_buffe
r_size=1M --maximum-join_buffer_size=1M --maximum-max_heap_table_size=1M --maximum-max_join_size=1M --maximum-myisam_max_sort_fi
le_size=1M --maximum-myisam_mmap_size=1M --maximum-myisam_sort_buffer_size=1M --maximum-optimizer_trace_max_mem_size=1M --maximu
m-preload_buffer_size=1M --maximum-query_alloc_block_size=1M --maximum-query_prealloc_size=1M --maximum-range_alloc_block_size=1
M --maximum-read_buffer_size=1M --maximum-read_rnd_buffer_size=1M --maximum-sort_buffer_size=1M --maximum-tmp_table_size=1M --ma
ximum-transaction_alloc_block_size=1M --maximum-transaction_prealloc_size=1M --log-output=none --sql_mode=ONLY_FULL_GROUP_BY
2015-04-14 11:22:40 [Init] Setting up standard working subdirectories
2015-04-14 11:22:49 [Init] Loading timezone data into mysql database
2015-04-14 11:22:50 [Info] If you see a [GLIBC] crash above, change reducer to use a non-Valgrind-instrumented build of mysql_tz
info_to_sql (Ref. BUG#13498842)
2015-04-14 11:23:10 [Init] Run mode: MODE=3: mysqld error log
2015-04-14 11:23:10 [Init] Looking for this string: 't1->result_range' in mysqld error log output (@ /tmp/1428974560/error.log.o
ut when MULTI mode is not active)
2015-04-14 11:23:10 [Info] Leading [] = No bug/issue found yet | [*] = Bug/issue at least seen once
2015-04-14 11:23:10 [Init] Number of lines in input file: 3963
2015-04-14 11:23:10 [] [Stage V] Verifying the bug/issue exists and is reproducible by reducer (duration depends on initial inpu
t file size)
2015-04-14 11:23:10 [] [Stage V] [MULTI] Starting 10 verification subreducer threads to verify if the issue is sporadic (/tmp/14
28974560/subreducer/)
2015-04-14 11:23:10 [] [Stage V] [MULTI] Ensuring any old subreducer processes are terminated
2015-04-14 11:23:18 [] [Stage V] [MULTI] Forking subreducer threads [PIDs]: #1 [6326] #2 [6389] #3 [6509] #4 [6604] #5 [6710] #6
 [6826] #7 [6937] #8 [7058] #9 [7204] #10 [7364]
2015-04-14 11:23:18 [] [Stage V] [MULTI] Waiting for all forked verification subreducer threads to finish/terminate
2015-04-14 11:27:51 [] [Stage V] [MULTI] Finished/Terminated verification subreducer threads: #1 #2 #3 #4 #5 #6 #7 #8 #9 #10
2015-04-14 11:27:51 [] [Stage V] [MULTI] All verification subreducer threads have finished/terminated
2015-04-14 11:27:51 [*] [Stage V] [MULTI] Threads which reproduced the issue: #1 #2 #3 #4 #5 #6 #7 #8 #9 #10
2015-04-14 11:27:51 [*] [Stage V] [MULTI] All threads reproduced the issue: this issue is not sporadic
2015-04-14 11:27:51 [*] [Stage V] [MULTI] Note: if this issue proves sporadic in actual reduction (slow/stalling reduction), use
 the FORCE_SPORADIC=1 setting
2015-04-14 11:27:51 [*] [Stage V] [MULTI] Ensuring any rogue subreducer processes are terminated
2015-04-14 11:27:52 [*] [Stage V] [MULTI] Deciding which verified output file to keep out of 10 threads
2015-04-14 11:27:52 [Init] Initial number of lines in restructured input file: 3963
2015-04-14 11:27:52 [Info] Linecounts for restructured files are usually higher as INSERT lines are broken up etc.
2015-04-14 11:27:52 [*] [Stage 1] Now executing first trial in stage 1 (duration depends on initial input file size)
2015-04-14 11:27:52 [*] [Stage 1] [Trial 1] Now filtering line(s) 414 to 3584 (Current chunk size: 3171)
2015-04-14 11:27:56 [*] [Stage 1] [Trial 1] [*ErrorLogOutputBug*] [0] Swapping files & saving last known good mysqld error log o
```

pquery core & framework

- 80 to 120 coredumps (crashes/asserts) per hour
- C++ core
- 15 Seconds per trial run time
  (w/ Thousands of SQL lines executed per trial)
- High end automation
- Full framework, 3 binaries (MS,PS,MD)
- PXC & third party tools/SE testing compatible

pquery Track Record

- In the first ~2 months of it's life, over 200 bugs were logged with Oracle, Percona and TokuTek, most with high quality testcases
- We're now up to ~300 bugs logged, most with 2-8 line SQL testcases
- Community showing interest & Oracle MySQL fixing the upstream bugs
- 13 Part QA (inc pquery) series in progress (http://goo.gl/wIZDjw)

Checkout the Full pquery Framework Today!

Twitter: @RoelVandePaar @rameshvs02