

MySQL 5.6 New Replication Features

Ronald Bradford
New York & Boston
March 2012




DISCLAIMER

The presentation provides information that is publicly available for MySQL 5.6 GA. The content of this presentation has been created by an independent MySQL Expert and is not an official statement of Oracle Corporation or a reflection of any of its employees.

Slides available at <http://effectivemysql.com>

EFFECTIVE MYSQL

- Chapter 3 Sample PDF
- Number of points discussed



<http://www.mhprofessional.com/product.php?isbn=0071791868>

MySQL 5.6 Replication

<http://dev.mysql.com/tech-resources/articles/mysql-5.6-replication.html>

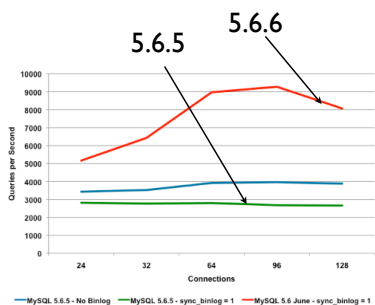
NEW 5.6 FEATURES

- Binary log group commit
- Multi-threaded slaves
- Crash-save slaves
- Global Transaction Identifier (GTID)
- Replication checksums
- Optimized row-based replication
- Time-delayed replication
- Remote binary log backup
- UUID

GROUP COMMIT

- Enabled by default
- Ideal for write-intensive workloads
- Produces less master overhead
- Allows sync_binlog=1 without significant impact

GROUP COMMIT



<http://mysqlmusings.blogspot.se/2012/06/binary-log-group-commit-in-mysql-56.html>

GROUP COMMIT

```
# my.cnf
[mysqld]
binlog_order_commits = ON | OFF
binlog_max_flush_queue_time = N
```

MULTI-THREADED SLAVES

- Enables multiple Slave SQL_THREADS
- Works on per schema
- Ideal for SaaS model

Dynamic



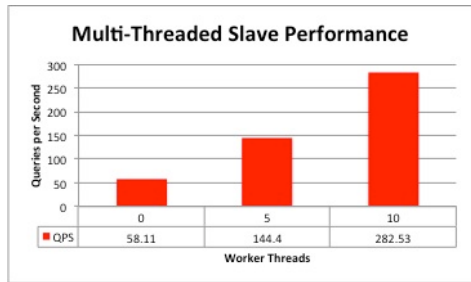
MULTI-THREADED SLAVES

```
# my.cnf
[mysqld]
slave_parallel_workers = N

mysql> SET GLOBAL slave_parallel_workers=N;
```



MULTI-THREADED SLAVES



https://blogs.oracle.com/MySQL/entry/benchmarking_mysql_replication_with_multi



MULTI-THREADED SLAVES

```
slave> SELECT * FROM mysql.slave_worker_info\G
...
***** 3. row *****
      Master_id: 3
      Worker_id: 2
      Relay_log_name: ./gamma-relay-bin.000009
      Relay_log_pos: 13641
      Master_log_name: alpha-bin.000003
      Master_log_pos: 21209
      Checkpoint_relay_log_name: ./gamma-relay-bin.000009
      Checkpoint_relay_log_pos: 13329
      Checkpoint_master_log_name: alpha-bin.000003
      Checkpoint_master_log_pos: 20897
      Checkpoint_seqno: 0
      Checkpoint_group_size: 64
      Checkpoint_group_bitmap:
```

Uses InnoDB



CRASH SAFE SLAVES

- Removes file based position
 - master.info
 - relay-log.info
- NOTE: Files removed when enabled
- Not enabled by default

Dynamic



CRASH SAFE SLAVES

```
# my.cnf
[mysqld]
master-info-repository = TABLE | FILE
relay-log-info-repository = TABLE | FILE

#sync_master_info
#sync_relay_log_info
```



CRASH SAFE SLAVES

```
mysql> SELECT * FROM mysql.slave_relay_log_info\G
Master_id: 3
Number_of_lines: 6
Relay_log_name: ./gamma-relay-bin.000006
Relay_log_pos: 346
Master_log_name: alpha-bin.000002
Master_log_pos: 391
Sql_delay: 0
Number_of_workers: 0
```

Uses InnoDB



CRASH SAFE SLAVES

```
slave> SELECT * FROM mysql.slave_master_info\G
Master_id: 3
Number_of_lines: 23
Master_log_name: alpha-bin.000003
Master_log_pos: 187
Host: alpha
User_name: repl
User_password: repl
Port: 3306
Connect_retry: 60
Enabled_ssl: 0
...
Heartbeat: 1800
Bind:
Ignored_server_ids: 0
Uuid: ba7ac732-b707-11e1-a1b3-0800275824dc
Retry_count: 86400
Ssl_crl:
```

Uses InnoDB



CRASH SAFE SLAVES

```
$ cat [data-dir]/master.info

mysql> SELECT * FROM mysql.slave_master_info;
Empty set (0.00 sec)
mysql> STOP SLAVE;
mysql> SET GLOBAL master_info_repository='TABLE';
mysql> SELECT * FROM mysql.slave_master_info\G
Number_of_lines: 23
Master_log_name: mysql-bin.000002
Master_log_pos: 120
Host: 127.0.0.1
....
```

GTID

- Unique event identifier across replication topology
 - All or nothing
- InnoDB only
- Some statement limitations
- Increased slave write load
- Workbench utilities

<http://dev.mysql.com/doc/refman/5.6/en/replication-gtids.html>

GTID

```
# my.cnf
[mysqld]
log-bin
log-slave-updates
gtid-mode=ON
disable-gtid-unsafe-statements

mysql> CHANGE MASTER TO ...
> MASTER_AUTO_POSITION = 1;
```

GTID

```
slave> SHOW SLAVE STATUS\G

Master_Host: master.example.com
Master_User: repl
...
Master_UUID: f52e965a-b18d-11e1-bdfe-00238b979631
Retrieved_Gtid_Set:
Executed_Gtid_Set:
8A94F357-AAB4-11DF-86AB-C80AA9429562:1-9,
F52E965A-B18D-11E1-BDFE-00238B979631:1-155
```

BINLOG CHECKSUM

- Writes checksum to binary log
- Default is ON



BINLOG CHECKSUM

```
# my.cnf
[mysqld]
binlog_checksum = NONE | CRC32
master_verify_checksum = ON | OFF

# slave my.cnf
[mysqld]
slave_sql_verify_checksum = ON | OFF
```



ROW IMAGE

- Optimizing before and after image
- Reduces disk & network

Dynamic



ROW IMAGE

```
# my.cnf
[mysqld]
binlog_format = ROW
binlog_row_image = FULL | MINIMAL | NOBLOB;
```



DELAYED REPLICATION

- Enables slave delay
- Testing slave lag
- Helps in potential damaging command

Dynamic

<http://dev.mysql.com/doc/refman/5.6/en/replication-delayed.html>
<http://datacharmer.blogspot.com/2011/01/first-look-at-delayed-replication-in.html>



DELAYED REPLICATION

```
mysql> STOP SLAVE;  
mysql> CHANGE MASTER TO  
  > MASTER_DELAY = N;  
mysql> START SLAVE;  
  
mysql> SHOW SLAVE STATUS\G  
...  
  Slave_IO_State: Waiting for master to send event  
    Master_Host: master.example.com  
      SQL_Delay: 10  
  SQL_Remaining_Delay: 6  
    Slave_SQL_Running_State: Waiting until MASTER_DELAY  
seconds after master executed event
```



REMOTE BINLOG

- Remote copy
- Does not require SSH
- Improves backup strategy

<http://dev.mysql.com/doc/refman/5.6/en/mysqlbinlog.html>



REMOTE BINLOG

```
$ mysqlbinlog -user=root \  
  --host=[host] \  
  --port=3306 \  
  --raw \  
  --read-from-remote-server \  
  --stop-never \  
  binary-logs.000001
```



UUID

- Universal Unique Identifier (UUID)
- MySQL Enterprise Monitor feature
- Part of GTID



UUID

```
master> SELECT @@GLOBAL.server_uuid;
+-----+
| @@GLOBAL.server_uuid |
+-----+
| cae53eb6-1aa8-11e1-a608-00238b979631 |
+-----+
slave> SHOW SLAVE STATUS\G
...
Master_Server_Id: 1
Master_UUID: cae53eb6-1aa8-11e1-a608-00238b979631

$ cat auto.cnf
[auto]
server-uuid= cae53eb6-1aa8-11e1-a608-00238b979631
```

Trap for virtual server clones



CONCLUSION

- Many new features
- Improved write performance
- Improved slave performance (especially SaaS)



Email: me@ronaldbradford.com

Slides available at <http://effectivemysql.com>

Effective MySQL Performance, Scalability & Business Continuity

