



Fixing Replication with Percona toolkit

Martin Arrieta
June 2012

martin.arrieta@percona.com
Twitter: [@martinarrietac](https://twitter.com/martinarrietac)

Agenda

- General review of Percona Toolkit
- Why do we need those tools?
- Possible scenarios
- Understanding the DSN in Percona-Toolkit
- General review of pt-slave-find
- General review of pt-table-checksum
- General review of pt-table-sync
- Practical example
- Questions

Percona Toolkit

- Percona Toolkit is a collection of advanced command-line tools used by Percona support staff to perform a variety of MySQL and system tasks that are too difficult or complex to perform manually.
- Percona Toolkit is derived from Maatkit and Aspersa, two of the best-known toolkits for MySQL server administration. It is developed and supported by Percona Inc. For more information and other free, open-source software developed by Percona, visit <http://www.percona.com/software/>.

Percona Toolkit

- Written in Perl
- Dependencies:
 - perl-Time-HiRes
 - perl-DBI
 - perl-DBD-MySQL
- To install all tools, you can use our yum/apt repo, but you can also download directly from our website.

```
# wget http://percona.com/get/tool\_name
```

- For example:

```
# wget http://percona.com/get/pt-table-sync
```

Percona Toolkit

But wait!

Replication can be broken?

Percona Toolkit

YES!

How my replication get out of sync?

- Someone changed the information in the slave directly
- Statement Based Replication
 - Non deterministic SQL (UPDATE/DELETE with LIMIT and without ORDER BY)
 - Triggers and Store procedures.
- Master in MyISAM and slave in InnoDB (deadlock)
- --replicate-ignore-db with fully qualified queries
- And 10k other reasons :)

Other reasons.

- Binlog corruption in the master server
- PURGE BINARY LOGS was issued and you don't have enough binlogs files to update the slave.
- read_buffer_size larger than max_allowed_packet
- MySQL bugs
- Version incompatibilities

DSN in Percona Toolkit

- Uses DSNs to specify how to create a DBD connection to a MySQL server. A DSN is a comma-separated string of key=value parts.
- Some options are:

Some options

- **h** - Hostname
- **u** - Username
- **P** - Port
- **D** - Default DB
- **S** - Socket
- **F** - Default file

Examples

```
h=localhost,u=root,p=mysecret
```

```
S=/tmp/mysql/mysql.sock,u=root
```

```
h=10.10.1.100,u=ptuser,P=3307
```

Finding the slaves with pt-slave-find

- **Goal:**
 - Find and print the slaves information from a given master.
- **Risk:**
 - Low, it is a read only tool
- **User privileges needed:**
 - REPLICATION CLIENT
 - REPLICATION SLAVE
 - PROCESS
- **Usage:**
 - `./pt-slave-find master-host`
- **More information:**
 - <http://www.percona.com/doc/percona-toolkit/pt-slave-find.html>

pt-slave-find example

```
# ./pt-slave-find h=10.10.0.100,u=ptslavefind,p=s3cret
10.10.0.100
Version          5.5.24-55-log
Server ID        1
Uptime           20:45:06 (started 2012-06-13T17:07:15)
Replication      Is not a slave, has 1 slaves connected, is not read_only
Filters
Binary logging   STATEMENT
Slave status
Slave mode        STRICT
Auto-increment   increment 1, offset 1
InnoDB version   1.1.8-rel26.0
+- 10.10.0.101
  Version          5.5.25
  Server ID        2
  Uptime           20:29:57 (started 2012-06-13T17:22:24)
  Replication      Is a slave, has 0 slaves connected, is not read_only
  Filters
  Binary logging   STATEMENT
  Slave status     0 seconds behind, running, no errors
  Slave mode        STRICT
  Auto-increment   increment 1, offset 1
  InnoDB version   1.1.8
```

Finding the differences with pt-table-checksum

- **Goal:**
 - Performs an online replication consistency check by executing checksum queries on the master, which produces different results on replicas that are inconsistent with the master.
- **Risk:**
 - Adds load to the servers but it is a read-only tool
- **User privileges needed:**
 - For tables to make the checksums
 - REPLICATION CLIENT
 - REPLICATION SLAVE
 - PROCESS
 - SELECT
 - SUPER
 - For --replicate table you will need ALL PRIVILEGES
- **Usage:**
 - `./pt-table-checksum master-host`
- **More information:**
 - <http://www.percona.com/doc/percona-toolkit/pt-table-checksum.html>

pt-table-checksum example

```
# ./pt-table-checksum h=10.10.0.100,u=pttablechecksum,p=s3cret
      TS ERRORS  DIFFS      ROWS  CHUNKS  SKIPPED    TIME TABLE
06-14T15:02:33      0      0        0       1       0    0.157 mysql.columns_priv
06-14T15:02:33      0      0         3       1       0    0.022 mysql.db
06-14T15:02:34      0      0         0       1       0    1.029 mysql.event
06-14T15:02:35      0      0         0       1       0    1.031 mysql.func
06-14T15:02:36      0      1         0       1       0    1.030 mysql.help_category
06-14T15:02:36      0      1         0       1       0    0.026 mysql.help_keyword
06-14T15:02:36 Skipping table mysql.help_relation because on the master it would be checksummed in one
chunk but on these replicas it has too many rows:
    1029 rows on replslave
The current chunk size limit is 556 rows (chunk size=278 * chunk size limit=2.0).
06-14T15:02:37      0      1         0       1       0    1.029 mysql.help_topic
06-14T15:02:37      0      0         0       1       0    0.024 mysql.host
06-14T15:02:38      0      0         0       1       0    1.036 mysql.ndb_binlog_index
06-14T15:02:38      0      0         0       1       0    0.031 mysql.plugin
06-14T15:02:39      0      0         0       1       0    1.030 mysql.proc
06-14T15:02:39      0      0         0       1       0    0.028 mysql.procs_priv
06-14T15:02:39      0      1         2       1       0    0.024 mysql.proxies_priv
06-14T15:02:39      0      0         0       1       0    0.024 mysql.servers
06-14T15:02:40      0      0         1       1       0    1.035 mysql.tables_priv
06-14T15:02:41      0      0         0       1       0    0.031 mysql.time_zone_transition_type
06-14T15:02:42      0      1        11       1       0    1.033 mysql.user
06-14T15:02:44      0      0      4079       4       0    1.254 world.City
06-14T15:02:44      0      0       239       1       0    0.030 world.Country
06-14T15:02:44      0      0       984       1       0    0.025 world.CountryLanguage
```

Synchronizing tables with pt-table-sync

- **Goal:**
 - Synchronizes data efficiently between MySQL tables.
- **Risk:**
 - High!! this tool changes data, so it is a good idea to backup your database before run this tool.
- **User privileges needed:**
 - For tables to make the checksums REPLICATION CLIENT, REPLICATION SLAVE, PROCESS, SUPER
 - For --replicate table you will need ALL PRIVILEGES
 - For tables that you want to sync SELECT, INSERT, UPDATE, DELETE
- **Usage:**
 - `./pt-table-sync master-host slave-host`
 - `./pt-table-sync --replicate master-host`
 - `./pt-table-sync --sync-to-master slave-host`
- **More information:**
 - <http://www.percona.com/doc/percona-toolkit/pt-table-sync.html>

pt-table-sync example

```
#./pt-table-sync h=10.10.0.101,u=pttablesinc,p=s3cret --databases=world --print --sync-  
to-master  
REPLACE INTO `world`.`Country`(`code`, `name`, `continent`, `region`, `surfacearea`,  
`indepyear`, `population`, `lifeexpectancy`, `gnp`, `gnpold`, `localname`,  
`governmentform`, `headofstate`, `capital`, `code2`) VALUES ('URY', 'Uruguay', 'South  
America', 'South America', '175016.00', '1828', '3337000', '75.2', '20831.00',  
'19967.00', 'Uruguay', 'Republic', 'Jorge Batlle Ib??ez', '3492', 'UY') /*percona-toolkit  
src_db:world src_tbl:Country src_dsn:P=3306,h=10.10.0.100,p=...,u=pttablesinc  
dst_db:world dst_tbl:Country dst_dsn:h=10.10.0.101,p=...,u=pttablesinc lock:1  
transaction:1 changing_src:1 replicate:0 bidirectional:0 pid:3718 user:root  
host:repmaster*/;  
#./pt-table-sync h=10.10.0.100,u=pttablesinc,p=s3cret --databases=world --print --  
replicate=percona.checksums  
REPLACE INTO `world`.`Country`(`code`, `name`, `continent`, `region`, `surfacearea`,  
`indepyear`, `population`, `lifeexpectancy`, `gnp`, `gnpold`, `localname`,  
`governmentform`, `headofstate`, `capital`, `code2`) VALUES ('URY', 'Uruguay', 'South  
America', 'South America', '175016.00', '1828', '3337000', '75.2', '20831.00',  
'19967.00', 'Uruguay', 'Republic', 'Jorge Batlle Ib??ez', '3492', 'UY') /*percona-toolkit  
src_db:world src_tbl:Country src_dsn:h=10.10.0.100,p=...,u=pttablesinc dst_db:world  
dst_tbl:Country dst_dsn:h=10.10.0.101,p=...,u=pttablesinc lock:1 transaction:1  
changing_src:percona.checksums replicate:percona.checksums bidirectional:0 pid:3728  
user:root host:repmaster*/;
```

Practical example.



Questions?

