



PERCONA
Performance Consulting Experts

Percona Server and XtraBackup: painless operations

Date, time, place:

MySQL Conference & Expo 2011
13-Apr-2011

Reporter:

Vadim Tkachenko
Co-founder, CTO,
Percona Inc

-
- Setup slave in couple minutes
 - Copy only part of your databases
 - It is all possible

Percona Server

- Percona Server is an enhanced drop-in replacement for MySQL
 - Performance: better and stable
 - Diagnostics
 - Operations



PERCONA
SERVER

Percona XtraBackup

- Percona XtraBackup
 - O'Reilly award MySQL Application of Year
 - The world's only open-source
 - Free MySQL hot backup software
 - non-blocking backups for InnoDB and XtraDB databases
 - Works with MySQL / MariaDB / Drizzle / Percona Server
 - Release 1.6 supports MySQL 5.5 and Percona Server 5.5




PERCONA
XTRABACKUP

What's in this talk

- XtraBackup basic and advance usage
- Backup features integrated with Percona Server

Basic idea

1. Copy tables, log changes into xtrabackup_logfile



changes → **xtrabackup_logfile**

2. Prepare – apply log to copied tables

Basic idea, once again

1. Copy tables, like cp
 - We copy files, but there are background changes
 - We log all changes during backup into log file
 - Necessary for consistent data
2. Prepare – apply log to copied tables

XtraBackup binaries and scripts

- **xtrabackup**
 - Percona Server XtraDB/InnoDB Plugin
- **xtrabackup_51**
 - MySQL 5.0/5.1 Built-in InnoDB
- **xtrabackup_55**
 - Percona Server 5.5 / MySQL 5.5
- **innobackupex** (originally from InnoDB Hot Backup)
 - wrapper to copy .frm files, triggers, MyISAM tables...
- **tar4ibd**: patched tar version, for streaming backups
 - reads innodb pages
 - capable of doing page checksumming
 - retries if page is corrupt
 - extract with 'tar -i'

Basic usage

- Copy tables - take backup
 - innobackupex /backup/dir
- Prepare – apply log to copied tables
 - innobackupex --apply-log /backup/dir
- Recovery:
 - Copy /backup/dir to datadir

Stream

- Unix-way to handle data
- `innobackupex | script1 | scrip2 | ..`

Stream

- `innobackupex --stream=tar ./ > backup.tar`
- Tar compatible stream (`-i` required to extract)
 - Tar was replaced by `libtar` and `tar4ibd`
- Compress
 - `innobackupex --stream=tar ./ | gzip - > backup.tar.gz`
 - `innobackupex --stream=tar ./ | qpress -io xtrabackup.tar > backup.tar.qpress`
- Copy to remote host
 - `innobackupex --stream=tar ./ | ssh vladim@desthost "cat - > /data/vol1/mysqluc/backup.tar"`
 - `ssh user@desthost "(nc -l 9999 | qpress -dio | tar -I /data/backups/backup.tar &)" && innobackupex --stream=tar ./ | qpress -io backup.tar | nc desthost 9999`
- The same but with throttling
 - `innobackupex --stream=tar | pv -q -L10m | ssh vladim@desthost "tar xfi -"`

Slave in ~~one~~ two clicks

- **Copy directly to slave**
 - `innobackupex --stream=tar /tmp/ --slave-info | ssh user@DESTSERVER "tar xfi - -C /DESTDIR"`
- **Prepare**
 - `innobackupex --apply-log /DESTDIR`
- **Start mysqld and SET MASTER**
 - From `xtrabackup_binlog_info` file
- **Effective for scale-out, fast cloud deployment**

Moving forward

- Parallel backup
- Partial backup
- Incremental backup
- Export/Import

Parallel backup

- Only local copy (not stream)
- innobackupex `-parallel=N`
 - Table1, table2 ,..., tableN are copied in parallel

Partial backup

- You have databases:
 - Customer1 Customer2 ... Customer65535
- You need:
 - Setup new server only with Customer10
- Use case: multi-tenant / shards
 - Our you know that some databases just not needed...
- `innobackupex --tables=[REGEXP]`
 - `innobackupex --tables=Customer10.*`
- `innobackupex --tables_file=<file with list of tables>`
- It works on MySQL / Percona Server, but:

Partial backup


- `innobackupex --tables=[REGEXP]`
 - `innobackupex --tables=Customer10.*`
- `innobackupex --databases=<list of databases>`
 - `innobackupex --databases=Customer10`
- `innobackupex --tables_file=<file with list of tables>`
- It works on MySQL / Percona Server, but:

Partial backup

- ibdata1 (and data dictionary) also copied
 - MySQL will complain <.. Table customer1.table1 has missed .ibd> ... <.. Table customer65535.table1 has missed .ibd>
 - Percona Server cleans data dictionary, missed entries just removed

Incremental backup

- Copy only changes since last backup / given point



Incremental backup

- Full Backup
 - innobackupex /backup/dir/full
- Incremental backup
 - innobackupex /backup/dir/inc –slave-info –incremental-basedir= /backup/dir/full
- Full prepare
 - innobackupex –apply-log –redo-only /backup/dir/full
 - innobackupex –apply-log /backup/dir/full –incremental-dir=/backup/dir/inc

Incremental problems

- Still whole tables scan to find changes
- Facebook uses special setup to find changes
 - http://www.facebook.com/note.php?note_id=10150098033318920
- We are on the way to get “real” incremental
 - Requires changes to Percona Server
 - Detect and what pages changed in real time

Combination with Percona Server

- Export / import tables
- You have databases:
 - Customer1 Customer2 ... Customer65535
- You need:
 - Copy Customer10 from one server to another
- Use case: multi-tenant / shards

Export/Import

- innobackupex --tables=Customer10.* /backup/dir
- innobackupex --apply-log --export /backup/dir
 - Will produce .exp files along with regular tables
 - Contains information about index “root” pages

```
xtrabackup: export metadata of table 'sakila/payment' to file `./sakila/
payment.exp` (4 indexes)
xtrabackup:      name=PRIMARY, id.low=20078, page=3
xtrabackup:      name=idx_fk_staff_id, id.low=20079, page=4
xtrabackup:      name=idx_fk_customer_id, id.low=20080, page=5
xtrabackup:      name=fk_payment_rental, id.low=20081, page=6
xtrabackup: export metadata of table 'sakila/rental' to file `./sakila/
rental.exp` (5 indexes)
xtrabackup:      name=PRIMARY, id.low=20082, page=3
xtrabackup:      name=rental_date, id.low=20083, page=4
xtrabackup:      name=idx_fk_inventory_id, id.low=20084, page=5
xtrabackup:      name=idx_fk_customer_id, id.low=20085, page=6
xtrabackup:      name=idx_fk_staff_id, id.low=20086, page=7
```

Export/Import cont.

- On destination server (Percona Server):
 - set global `innodb_expand_import=1;` /
`innodb_import_table_from_xtrabackup=1` (5.5)
 - For each table from Customer10:
 - alter table tableN discard tablespace;
- Copy files `/backup/dir/Customer10` to destination datadir
 - For each table from Customer10:
 - alter table tableN import tablespace;

In development:

- Problem: in compressed backup database.tar.gz you can't extract specific tables
- Solution: compress per table
- Now:
 - Database.tar.gz
- Will be:
 - Customer1.table1.ibd.quicklz
 - Customer2.table1.ibd.quicklz
 - ...
 - Customer65536.table1.ibd.quicklz
- All this in parallel and streaming to backup server

Statistics

- We have backup, we can analyze it:

```
> xtrabackup -stats -target-dir=/backup/dir
<INDEX STATISTICS>
table: art/link_out104, index: PRIMARY, space id: 12, root page 3
estimated statistics in dictionary:
  key vals: 25265338, leaf pages 497839, size pages 498304
real statistics:
  level 2 pages: pages=1, data=5395 bytes, data/pages=32%
  level 1 pages: pages=415, data=6471907 bytes, data/pages=95%
    leaf pages: recs=25958413, pages=497839, data=7492026403
bytes, data/pages=91%
table: art/link_out104, index: domain_id_2, space id: 12, root
page 4
estimated statistics in dictionary:
  key vals: 27755790, leaf pages 23125, size pages 26495
real statistics:
  level 2 pages: pages=1, data=510 bytes, data/pages=3%
  level 1 pages: pages=30, data=393125 bytes, data/pages=79%
    leaf pages: recs=25958413, pages=23125, data=337459369
bytes, data/pages=89%
```

Statistics in pretty-format

- Script <http://www.percona.com/docs/wiki/percona-xtrabackup:xtrabackup:statistics>

| TABLE | INDEX | TOT_PAGES | FREE_PAGES | PCT_FULL |
|-----------------|-----------------|-----------|------------|----------|
| art.link_out104 | | 832383 | 38561 | 86.8% |
| art.link_out104 | PRIMARY | 498304 | 49 | 91.9% |
| art.link_out104 | domain_id | 49600 | 6230 | 76.9% |
| art.link_out104 | domain_id_2 | 26495 | 3339 | 89.1% |
| art.link_out104 | from_message_id | 28160 | 142 | 96.3% |
| art.link_out104 | from_site_id | 38848 | 4874 | 79.4% |
| art.link_out104 | revert_domain | 153984 | 19276 | 71.4% |
| art.link_out104 | site_message | 36992 | 4651 | 83.4% |

Fast index creation

- Long waited feature in InnoDB-plugin
- Creates index “by sort” → very compact
 - But works only with ALTER TABLE add KEY command
- Does not play nice with mysqldump
 - MySQL dump
 - CREATE TABLE `stock` (`s_i_id` int(11) NOT NULL, ... `s_data` varchar(50) DEFAULT NULL, PRIMARY KEY (`s_w_id`,`s_i_id`), KEY `fkey_stock_2` (`s_i_id`))
 - We need CREATE TABLE without indexes; and ALTER TABLE after CREATE TABLE
 - Mysqldump from Percona Server:
 - **--innodb-optimize-keys**

Fast index creation

- Does not play nice with mysqldump
 - mysqldump
 - CREATE TABLE `stock` (`s_i_id` int(11) NOT NULL, ... `s_data` varchar(50) DEFAULT NULL, PRIMARY KEY (`s_w_id`,`s_i_id`), KEY `fkey_stock_2` (`s_i_id`))
 - We need CREATE TABLE without indexes; and ALTER TABLE after CREATE TABLE
- Mysqldump from Percona Server:
 - New option **--innodb-optimize-keys**

Third-party tools

- Zmanda
 - Enterprise backup management tool
 - Integration with XtraBackup in progress
- XtraBackup manager
 - Tool to manage backups: schedule, reclaim old copies etc
 - <http://code.google.com/p/xtrabackup-manager/>

References

- Project pages
 - <http://www.percona.com/software/percona-server/>
 - <http://www.percona.com/software/percona-xtrabackup/>
- Bugs
 - <https://bugs.launchpad.net/percona-server/>
 - <https://bugs.launchpad.net/percona-xtrabackup/>
- Mailling list
 - <http://groups.google.com/group/percona-discussion>

The end

- Slides will be online. <http://www.percona.com/about-us/presentations/>
- vadim@percona.com
- Your questions ?
- We are hiring!