



PERCONA
Performance Consulting Experts

XtraBackup: Hot Backups and More

Vadim Tkachenko

Morgan Tocker

<http://percona.com>

<http://MySQLPerformanceBlog.com>

MySQL CE

Apr 2010

Introduction

- Vadim Tkachenko
 - Percona Inc, CTO and Lead of Development
- Morgan Tocker
 - Percona Inc, Director of Training

Comparing Backup Options

	Engines	Warmth	Flex Backup	Flex Restore	Backup time	Restore time	Binlog coord	Impact
mysqldump	InnoDB/XtraDB	HOT	ROW	MOST	MED	WORST	YES	MEDIUM
mysqldump	ALL	WARM	ROW	MOST	MED	WORST	YES	MEDIUM
Snapshot	ALL	MOSTLY HOT	SYSTEM	SYSTEM	SLOW-MED	BEST	YES	HEAVY
XtraBackup	InnoDB/XtraDB	HOT	TABLE	TABLE	FAST	BEST	YES	LOW
InnoDB Hot Backup	InnoDB	HOT	TABLE	TABLE	FAST	BEST	YES	LOW
XtraBackup / InnoDB Hot Backup	ALL	WARM	TABLE	TABLE	FAST	BEST	YES	MEDIUM/LARGE
MySQLHotCopy	MyISAM	WARM	TABLE	TABLE	FAST	BEST	YES	LOW
Cold Backup	ALL	COLD	SYSTEM	SYSTEM	BEST	BEST	YES	VERY HEAVY

What is talk about

- XtraBackup
- Internals
- Basic and advance usage

Idea


- Copy XtraDB / InnoDB tables without stopping server (and locking tables)
- Challenge: Tables are being changed in background
- Alternative solutions
 - ZFS snapshot (Solaris / FreeBSD)
 - LVM, good if performance is not concern
 - 160 ios -> 25 ios
 - ~20% space reservation
 - InnoDB Hot Backup / Oracle Online Backup
 - Commercial, per-server, last release 2006 ?
 - Can be expensive for X,000 servers

Solution

- Copy tables and log changes
 - InnoDB transactional redo logs are doing the same
 - XtraBackup is mini instance of XtraDB / InnoDB
 - Idea taken from InnoDB Hot Backup

Basic design

- Copy tables, log changes into xtrabackup_logfile



- Prepare – apply log to copied tables

XtraBackup

- Non-blocking online backup of XtraDB / InnoDB / InnoDB–plugin tables
- Blocking for MyISAM

Tools

- Binaries
 - xtrabackup
 - Copy .ibd & xtrabackup_logfile
 - tar4ibd
- Script
 - innobackupex-1.5.1
 - Taken from InnoDB Hot Backup (GPL) , heavily modified
 - Copy .frm, myisam, locking etc

Basic usage

- `innobackupex-1.5.1 /destination/dir`
- `innobackupex-1.5.1 --apply-log /backup/dir`
 - `--use-memory` (1GB defaults)
- `innobackupex-1.5.1 --throttle=IOS /destination/dir`

Stream

- `innobackupex-1.5.1 --stream=tar ./ > backup.tar`
- Tar compatible stream (`-i` required to extract)
 - Tar was replaced by `libtar` and `tar4ibd`
- Compress
 - `innobackupex-1.5.1 --stream=tar ./ | gzip - > backup.tar.gz`
- Copy to remote host
 - `innobackupex-1.5.1 --stream=tar ./ | ssh vadim@desthost "cat - > /data/vol1/mysqluc/backup.tar"`
- The same but with throttling
 - `innobackupex --stream=tar | pv -q -L10m | ssh vadim@desthost "tar xfi -"`

Slave

- To setup slave from master
 - No specific actions needed
 - File `xtrabackup_binlog_info`
`mysql-bin.000020` `70206554`
- To backup from slave
 - `innobackupex-1.5.1 --slave-info /destination/dir`
- LOCK caveats
 - FLUSH TABLES WITH READ LOCK to get master/slave info and to copy MYISAM tables
 - 100GB of MyISAM tables will cause a problem


Incremental/Differential

- We are able to copy only changes from last backup

- `xtrabackup --backup --target-dir=/data/backups/incremental --incremental-basedir=/data/backups/2010-04-11_17-55-54`

- From output:

- ... `xtrabackup: The latest check point (for incremental): '2438774348' ...`



- `xtrabackup --prepare --target-dir=/data/vadim/mysqluc/2010-04-11_20-45-42 --incremental-dir=/data/vadim/mysqluc/incremental`

Export/import

- Export only table to another server
 - `xtrabackup --table=tpcc.order_line --backup --target-dir=/data/vadim/mysqluc/order_line/`
- Prepare
 - `xtrabackup --prepare --export --target-dir=/data/vadim/mysqluc/order_line/`
- Copy to remote host (it has to be XtraDB)
 - `alter table order_line discard tablespace;`
 - `set global innodb_expand_import=1;`
 - `alter table order_line import tablespace;`

Statistics

- Show details on tablespaces / indexes

```
table: tpcc/order_line, index: PRIMARY, space id: 25, root page: 3, zip size: 0
estimated statistics in dictionary:
  key vals: 32471816, leaf pages: 264267, size pages: 302592
real statistics:
  level 2 pages: pages=1, data=8406 bytes, data/pages=51%
  level 1 pages: pages=467, data=4756806 bytes, data/pages=62%
  leaf pages: recs=35116662, pages=264267, data=2410713870 bytes, data/pages=55%
```

Thank you!

- Questions?
- vadim@percona.com
- morgan@percona.com