
# XtraBackup

Vadim Tkachenko
CTO, Co-Founder, Percona Inc
http://percona.com
http://MySQLPerformanceBlog.com

Percona Live SF
Feb 2011

# Speaker

- CTO and co-founder of Percona Inc.
- Lead of Percona Server/XtraDB and XtraBackup development
- Co-author of “High Performance MySQL” 2nd edition
- Author on MySQLPerformanceBlog.com and SSDPerformanceBlog.com
- More reading on http://www.percona.com/about-us/our-team/vadim-tkachenko/
- Email: vadim@percona.com

# 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

- Couple features in today's talk are in source trunk only, next XtraBackup-1.6 binaries soon

# 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
- Oracle MySQL Enterpise Backup
  - Commercial, $4000 per-server

# 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

# XtraBackup

**Non-blocking online backup of XtraDB / InnoDB / InnoDB–plugin tables**

- Handles MySQL and Percona Server

**Blocking for MyISAM**

- Need to apply FLUSH TABLES WITH READ LOCK

# Basic design

1. Copy tables, log changes into xtrabackup_logfile

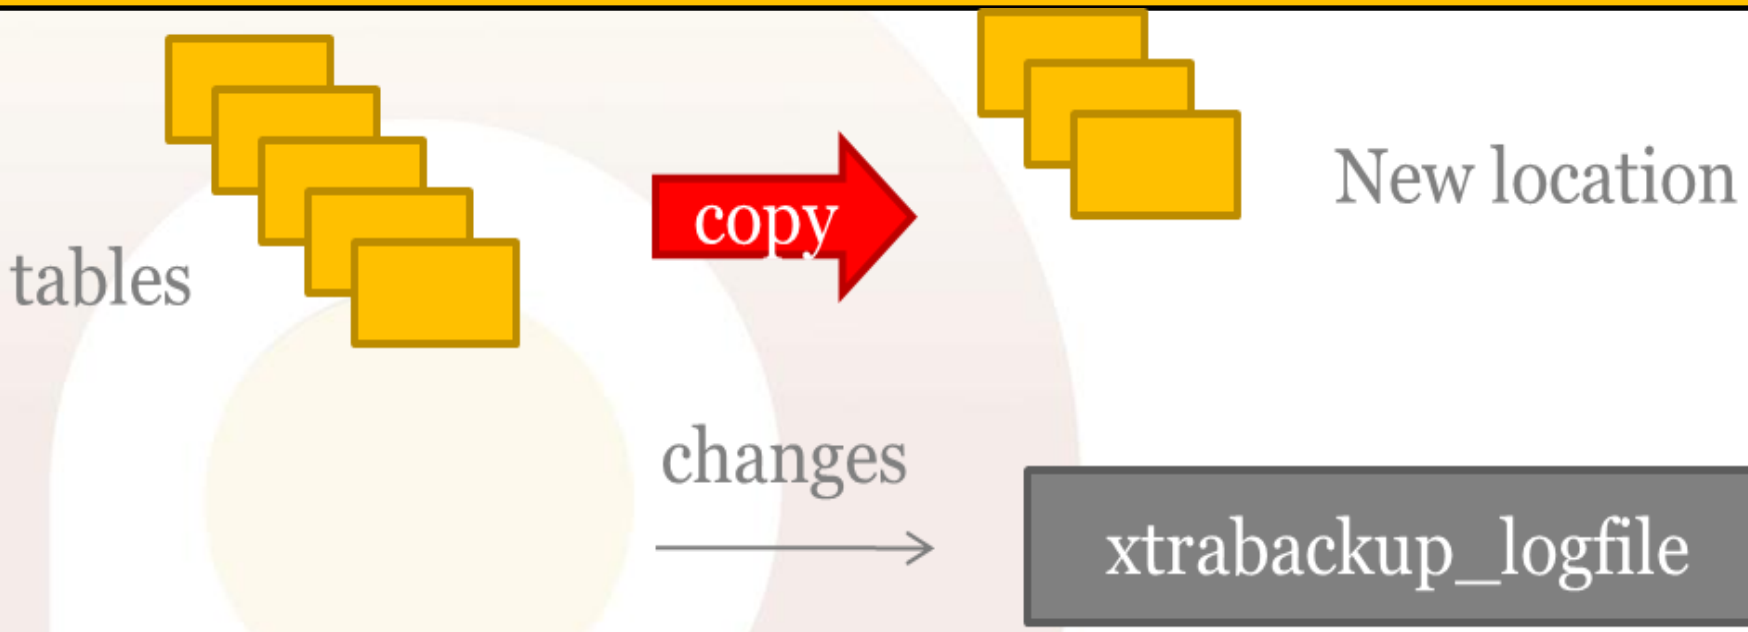


2. Prepare – apply log to copied tables

# Tools

## Binaries

- To handle .ibd & xtrabackup_logfile
- Xtrabackup
  - Percona Server 5.1 / InnoDB-plugin
- Xtrabackup_51
  - MySQL 5.1/5.0
- Xtrabackup_55
  - Percona Server 5.5 / MySQL 5.5
- tar4ibd

## Script

- Innobackupex
  - Taken from InnoDB Hot Backup (GPL) , heavily modified
  - Copy .frm, myisam, locking etc

# Two modes

## Localcopy

- You should have enough space mounted to server

## Stream

- Direct copy

# Basic usage

```
innobackupex  /destination/dir
```

```
innobackupex  --apply-log /backup/dir
```

- --use-memory (1GB defaults)

# 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

Copy to remote host

- innobackupex --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
- innobackupex --stream=tar /tmp/ --slave-info | ssh user@SLAVE "tar xfi - -C /data/slave/mysql“
- File xtrabackup_binlog_info
  mysql-bin.000020 70206554

## To backup from slave

- innobackupex --slave-info /data/backup

## 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

## Differential

- Backup since last full backup

## Incremental

- Backup since last incremental / differential

# Incremental/Differential

We are able to copy only changes from last backup

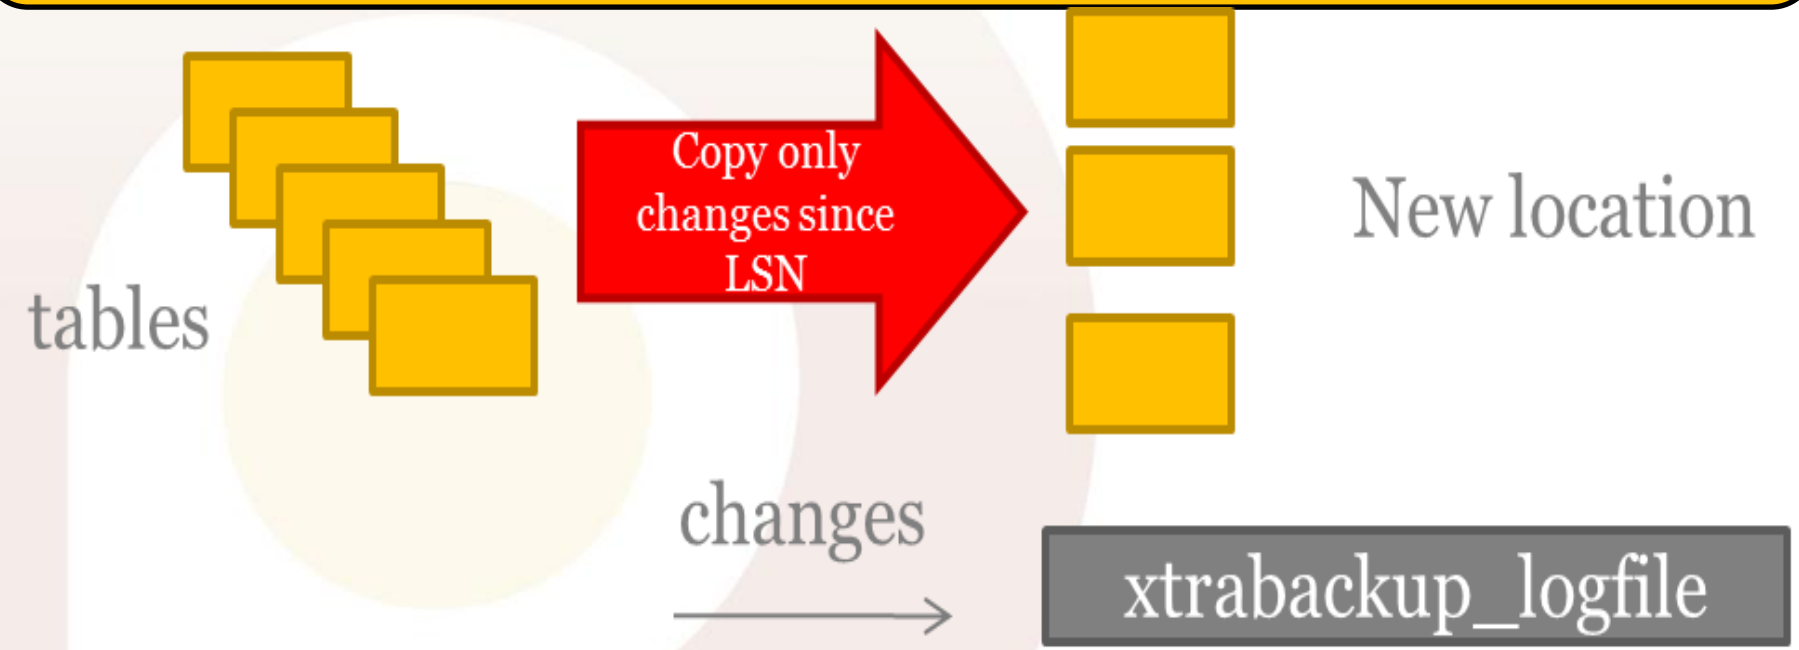

# Incremental/Differential

innobackupex /data/backups/incremental --incremental-basedir=/data/backups/2010-04-11_17-55-54

Or innobackupex /data/backups/incremental --incremental-lsn=2438774348

innobackupex –apply-log  --redo-only /data/backups/2011-02-11_20-45-42

innobackupex –apply-log  /data/backups/2011-02-11_20-45-42  --incremental-dir=/data/vadim/mysqluc/incremental

# Partial backup

## Copy only selected databases / tables

- Innobackupex –tables_file=file_with_list_of_tables /data/backup
- Innobackup –include=[regexp mask]
- Innobackup –databases=DB_list

# Parallel backup

## Copy files in parallel

- Innobackupex –paralell=8 /data/backup

# Export/import

## Export only table to another server

```
xtrabackup  --table=tpcc.order_line --backup --target-dir=/data/backup/order_line/
```

## Prepare

```
xtrabackup --prepare --export --target-dir=/data/backup/order_line/
```

## Copy to remote host (it has to be Percona Server)

```
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%

# Coming soon

## Table level compression

- With parallel compression support

## Incremental streaming

# Thank you!

- Questions?
- vadim@percona.com