



PERCONA  
Performance Consulting Experts

# XtraBackup



Kenny Gryp

[kenny.gryp@percona.com](mailto:kenny.gryp@percona.com)

Principal Consultant @ Percona

Collaborate 2011

# Percona

- MySQL/LAMP Consulting
- MySQL Support
- Percona Server (XtraDB)
- Percona XtraBackup
- InnoDB Recovery Toolkit, tcpstat
- (maatkit, innotop, aspersa, mysql-mmm, mysql-cacti-templates)
- ...
- <http://www.percona.com>
- <http://www.mysqlperformanceblog.com>

# What Is XtraBackup?

- Hot backup utility for InnoDB/XtraDB
- Support for MyISAM
- Open Source (GPL)
- Runs on Linux, FreeBSD, Windows (limited)
- Commercial Support Available
- A lot of features!

# Topics

- Comparison of MySQL Backup solutions
- How XtraBackup works
- Features

# Topics

---

- **Comparison of MySQL Backup solutions**
- How XtraBackup works
- Features

# Available Backup Solutions

InnoDB	Impact	Warmth	Backup Time	Restore Time
cold backup	veryhigh	cold	<b>very fast</b>	<b>fast</b>
mysqldump	medium	<b>hot</b>	medium	slow
snapshotting	high/ medium	<b>hot/ warm</b>	fast	<b>fast</b>
MySQL Enterprise Backup	<b>low</b>	<b>hot</b>	fast	<b>fast</b>
XtraBackup	<b>low</b>	<b>hot</b>	fast	<b>fast</b>

# Available Backup Solutions

InnoDB/MyISAM	Impact	Warmth	Backup Time	Restore Time
cold backup	veryhigh	cold	<b>very fast</b>	<b>fast</b>
mysqldump	high	warm	medium	slow
snapshotting	high/ medium	<b>hot/ warm</b>	fast	<b>fast</b>
MySQL Enterprise Backup	<b>low/ medium</b>	warm	fast	<b>fast</b>
XtraBackup	<b>low/ medium</b>	warm	fast	<b>fast</b>

# XtraBackup <-> MEB

- Shared Features:

- non-blocking, MyISAM support, compressed backups, partial backups, throttling, incremental support

- XtraBackup Only:

- Exporting/Importing individual tables
- Streaming Support
- Statistics/Analyze Data
- Parallel copying of data files

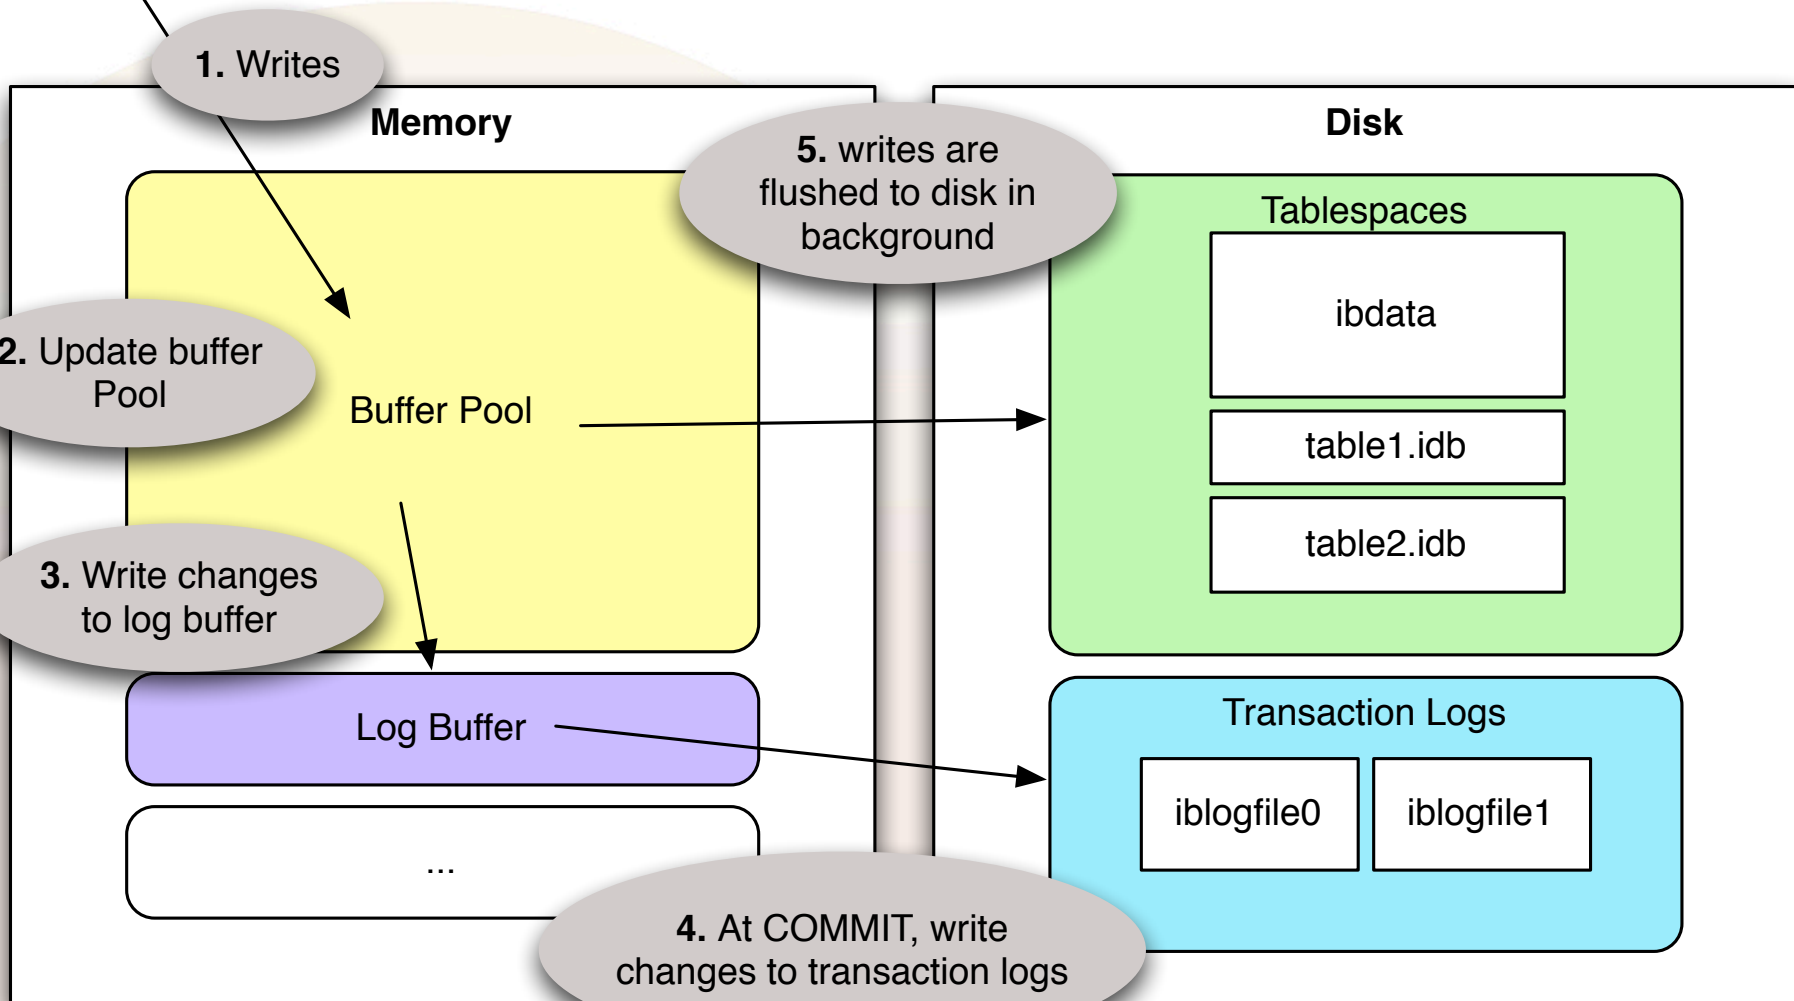
# Topics

- Comparison of MySQL Backup solutions
- **How XtraBackup works**
- Features

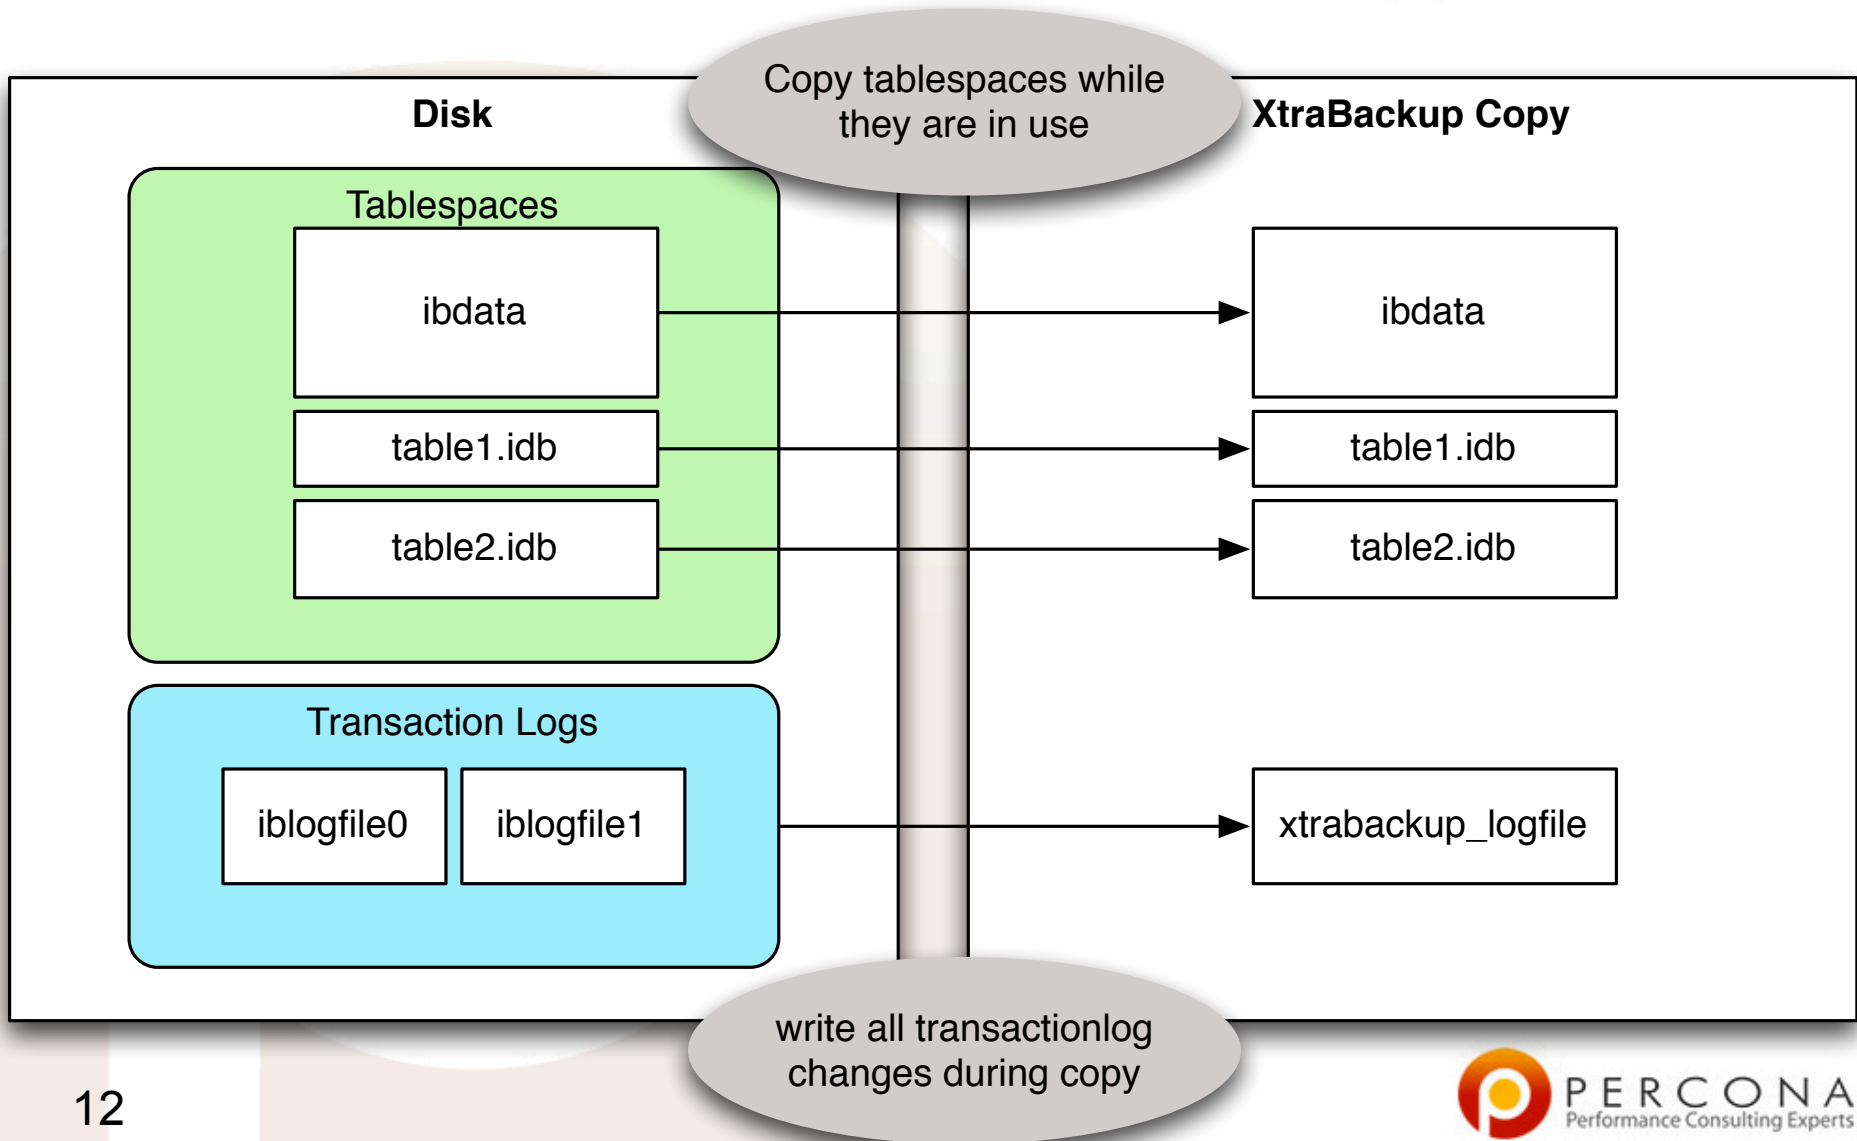
# XtraBackup Binaries

- xtrabackup: XtraDB/InnoDB Plugin
- xtrabackup\_51: MySQL 5.0/5.1 Built in InnoDB
- xtrabackup\_55: 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'

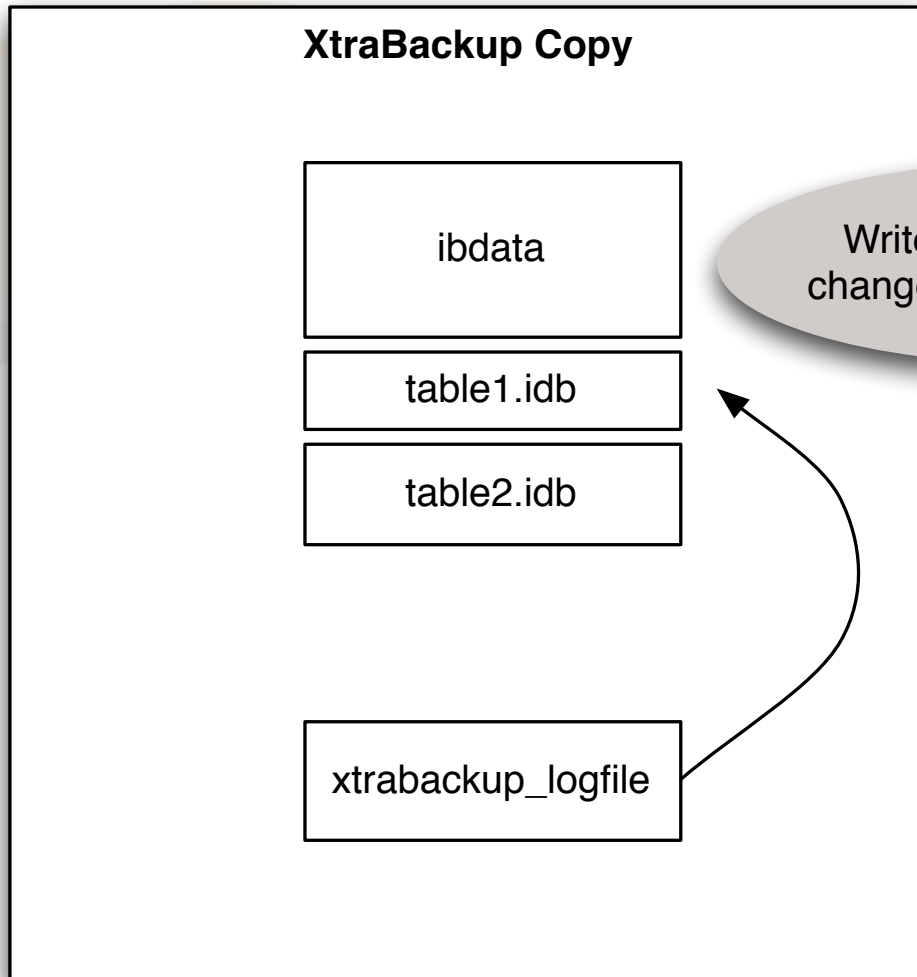
# (Very Basic) InnoDB



# How? (1) Copying



# How? (2) Prepare



Write all transaction log changes to the tablespaces

# How?

- When using innobackupex:
  - After the tablespaces have been copied:
    - FLUSH TABLES WITH READ LOCK;
    - Get binlog position
    - Copy all .frm .MYD .MYI .TRG .TRN files...
    - Stop copying transaction/redo logs
    - UNLOCK TABLES;

# How? Steps with xtrabackup

- Taking the backup:  
**# xtrabackup --backup --datadir=/var/lib/mysql/ --target-dir=/data/backups/mysql/**
- Preparing:  
**# xtrabackup --prepare --target-dir=/data/backups/mysql/**
- Only copies innodb data, not .frm files...

# How? Steps with innobackupex

- Taking the backup:  
**# innobackupex /data/backups**
- Preparing:  
**# innobackupex --apply-log /data/backups/  
2010-03-13\_02-42-44/**
- Restoring:  
**# innobackupex --copy-back /data/backups/  
2010-03-13\_02-42-44/**

# Sample Output (1): Backup

## # innobackupex /data/backups

InnoDB Backup Utility v1.5.1-xtrabackup; Copyright 2003, 2009 Innobase Oy.

All Rights Reserved.

This software is published under

the GNU GENERAL PUBLIC LICENSE Version 2, June 1991.

110203 21:26:02 innobackupex-1.5.1: Starting mysql with options: --unbuffered --

110203 21:26:02 innobackupex-1.5.1: Connected to database with mysql child process (pid=14647)

110203 21:26:08 innobackupex-1.5.1: Connection to database server closed

IMPORTANT: Please check that the backup run completes successfully.

At the end of a successful backup run innobackupex-1.5.1  
prints "completed OK!".

innobackupex-1.5.1: Using mysql Ver 14.14 Distrib 5.1.54, for debian-linux-gnu (x86\_64) using EditLine wrapper

innobackupex-1.5.1: Using mysql server version Copyright (c) 2000, 2010, Oracle and/or its affiliates. All rights reserved.

innobackupex-1.5.1: Created backup directory /data/backups/2011-02-03\_21-26-08

110203 21:26:08 innobackupex-1.5.1: Starting mysql with options: --unbuffered --

110203 21:26:08 innobackupex-1.5.1: Connected to database with mysql child process (pid=14671)

110203 21:26:12 innobackupex-1.5.1: Connection to database server closed

110203 21:26:12 innobackupex-1.5.1: Starting ibbackup with command:

## **xtrabackup --backup --suspend-at-end --target-dir=/data/ backups/2011-02-03\_21-26-08**

innobackupex-1.5.1: Waiting for ibbackup (pid=14681) to suspend

innobackupex-1.5.1: Suspend file '/data/backups/2011-02-03\_21-26-08/xtrabackup\_suspended'

xtrabackup Ver 1.4 Rev 193 for 5.1.47 unknown-linux-gnu (x86\_64)

xtrabackup: uses posix\_fadvise().

xtrabackup: cd to /var/lib/mysql

xtrabackup: Target instance is assumed as followings.

# Sample Output (2): Backup

>> log scanned up to (4953882516)

Copying ./ibdata1

to /data/backups/2011-02-03\_21-26-08/ibdata1

**>> log scanned up to (4953882516)**

...done

Copying ./sakila/film\_actor.ibd

to /data/backups/2011-02-03\_21-26-08/sakila/film\_actor.ibd

...done

Copying ./sakila/payment.ibd

to /data/backups/2011-02-03\_21-26-08/sakila/payment.ibd

...done

Copying ./sakila/rental.ibd

to /data/backups/2011-02-03\_21-26-08/sakila/rental.ibd

...done

Copying ./sakila/language.ibd

to /data/backups/2011-02-03\_21-26-08/sakila/language.ibd

...done

Copying ./sakila/customer.ibd

to /data/backups/2011-02-03\_21-26-08/sakila/customer.ibd

...done

Copying ./sakila/country.ibd

to /data/backups/2011-02-03\_21-26-08/sakila/country.ibd

...done

Copying ./sakila/store.ibd

to /data/backups/2011-02-03\_21-26-08/sakila/store.ibd

# Sample Output (3): Backup

```
110203 21:26:38 innobackupex-1.5.1: Continuing after ibbackup has suspended
110203 21:26:38 innobackupex-1.5.1: Starting mysql with options: --unbuffered --
110203 21:26:38 innobackupex-1.5.1: Connected to database with mysql child process (pid=14694)
>> log scanned up to (4953882516)
110203 21:26:42 innobackupex-1.5.1: Starting to lock all tables...
>> log scanned up to (4953882516)
>> log scanned up to (4953882516)
110203 21:26:53 innobackupex-1.5.1: All tables locked and flushed to disk
110203 21:26:53 innobackupex-1.5.1: Starting to backup .frm, .MRG, .MYD, .MYI,
innobackupex-1.5.1: .TRG, .TRN, .ARM, .ARZ and .opt files in
innobackupex-1.5.1: subdirectories of '/var/lib/mysql'
innobackupex-1.5.1: Backing up files '/var/lib/mysql/sakila/*.{frm,MYD,MYI,MRG,TRG,TRN,ARM,ARZ,opt,par}' (39 files)
innobackupex-1.5.1: Backing up files '/var/lib/mysql/mysql/*.{frm,MYD,MYI,MRG,TRG,TRN,ARM,ARZ,opt,par}' (65 files)
110203 21:26:54 innobackupex-1.5.1: Finished backing up .frm, .MRG, .MYD, .MYI, .TRG, .TRN, .ARM, .ARZ and .opt files
innobackupex-1.5.1: Resuming ibbackup

xtrabackup: The latest check point (for incremental): '4953882516'
>> log scanned up to (4953882516)
xtrabackup: Stopping log copying thread.
xtrabackup: Transaction log of lsn (4953882700) to (4953882700) was copied.
110203 21:26:56 innobackupex-1.5.1: All tables unlocked
110203 21:26:56 innobackupex-1.5.1: Connection to database server closed

innobackupex-1.5.1: Backup created in directory '/data/backups/2011-02-03_21-26-08'
innobackupex-1.5.1: MySQL binlog position: filename 'mysql-bin.000116', position 106
110203 21:26:56 innobackupex-1.5.1: completed OK!
```

# Sample Output (3): Backup

```
110203 21:26:38 innobackupex-1.5.1: Continuing after ibbackup has suspended
110203 21:26:38 innobackupex-1.5.1: Starting mysql with options: --unbuffered --
110203 21:26:38 innobackupex-1.5.1: Connected to database with mysql child process (pid=14694)
>> log scanned up to (4953882516)
```

```
110203 21:26:42 innobackupex-1.5.1: Starting to lock all tables...
```

```
>> log scanned up to (4953882516)
>> log scanned up to (4953882516)
```

```
110203 21:26:53 innobackupex-1.5.1: All tables locked and flushed to disk
```

```
110203 21:26:53 innobackupex-1.5.1: Starting to
backup .frm, .MRG, .MYD, .MYI,
innobackupex-1.5.1: .TRG, .TRN, .ARM, .ARZ and .opt files
in
```

```
innobackupex-1.5.1: subdirectories of '/var/lib/mysql'
innobackupex-1.5.1: Backing up files '/var/lib/mysql/sakila/*.{frm,MYD,MYI,MRG,TRG,TRN,ARM,ARZ,opt,par}' (39 files)
innobackupex-1.5.1: Backing up files '/var/lib/mysql/mysql/*.{frm,MYD,MYI,MRG,TRG,TRN,ARM,ARZ,opt,par}' (65 files)
110203 21:26:54 innobackupex-1.5.1: Finished backing up .frm, .MRG, .MYD, .MYI, .TRG, .TRN, .ARM, .ARZ and .opt files
innobackupex-1.5.1: Resuming ibbackup
```

```
xtrabackup: The latest check point (for incremental): '4953882516'
```

```
>> log scanned up to (4953882516)
```

```
xtrabackup: Stopping log copying thread.
```

```
xtrabackup: Transaction log of lsn (4953882700) to (4953882700) was copied.
```

```
110203 21:26:56 innobackupex-1.5.1: All tables unlocked
```

# Sample Output (3): Backup

```
110203 21:26:38 innobackupex-1.5.1: Continuing after ibbackup has suspended
110203 21:26:38 innobackupex-1.5.1: Starting mysql with options: --unbuffered --
110203 21:26:38 innobackupex-1.5.1: Connected to database with mysql child process (pid=14694)
>> log scanned up to (4953882516)
110203 21:26:42 innobackupex-1.5.1: Starting to lock all tables...
>> log scanned up to (4953882516)
>> log scanned up to (4953882516)
110203 21:26:53 innobackupex-1.5.1: All tables locked and flushed to disk
110203 21:26:53 innobackupex-1.5.1: Starting to backup .frm, .MRG, .MYD, .MYI,
innobackupex-1.5.1: .TRG, .TRN, .ARM, .ARZ and .opt files in
innobackupex-1.5.1: subdirectories of '/var/lib/mysql'
innobackupex-1.5.1: Backing up files '/var/lib/mysql/sakila/*.{frm,MYD,MYI,MRG,TRG,TRN,ARM,ARZ,opt,par}' (39 files)
innobackupex-1.5.1: Backing up files '/var/lib/mysql/mysql/*.{frm,MYD,MYI,MRG,TRG,TRN,ARM,ARZ,opt,par}' (65 files)
110203 21:26:54 innobackupex-1.5.1: Finished backing up .frm, .MRG, .MYD, .MYI, .TRG, .TRN, .ARM, .ARZ and .opt files
innobackupex-1.5.1: Resuming ibbackup
```

```
xtrabackup: The latest check point (for incremental): '4953882516'
```

**>> log scanned up to (4953882900)**

```
xtrabackup: Stopping log copying thread.
```

xtrabackup: **Transaction log of Isn (4953882516) to (4953883000)**  
**was copied.**

```
110203 21:26:56 innobackupex-1.5.1: All tables unlocked
```

```
110203 21:26:56 innobackupex-1.5.1: Connection to database server closed
```

```
innobackupex-1.5.1: Backup created in directory '/data/backups/2011-02-03_21-26-08'
```

# Sample Output (4): Prepare

## # innobackupex --apply-log /data/backups/ 2011-02-03\_21-26-08/

InnoDB Backup Utility v1.5.1-xtrabackup; Copyright 2003, 2009 Innobase Oy.  
All Rights Reserved.

```
...
110203 21:49:40 innobackupex-1.5.1: Starting ibbackup with command: xtrabackup --prepare --target-dir=/data/backups/
2011-02-03_21-26-08
xtrabackup Ver 1.4 Rev 193 for 5.1.47 unknown-linux-gnu (x86_64)
xtrabackup: cd to /data/backups/2011-02-03_21-26-08
xtrabackup: This target seems to be not prepared yet.
xtrabackup: xtrabackup_logfile detected: size=2097152, start_lsn=(4953882516)
xtrabackup: Temporary instance for recovery is set as followings.
xtrabackup: innodb_data_home_dir = ./
xtrabackup: innodb_data_file_path = ibdata1:10M:autoextend
xtrabackup: innodb_log_group_home_dir = ./
xtrabackup: innodb_log_files_in_group = 1
xtrabackup: innodb_log_file_size = 2097152
xtrabackup: Starting InnoDB instance for recovery.
xtrabackup: Using 104857600 bytes for buffer pool (set by --use-memory parameter)
...
110203 21:49:40 Percona XtraDB (http://www.percona.com) 1.0.8-11.2 started; log sequence number 4953882700
[notice (again)]
  If you use binary log and don't use any hack of group commit,
  the binary log position seems to be:
InnoDB: Last MySQL binlog file position 0 106, file name /var/log/mysql/mysql-bin.000116
xtrabackup: starting shutdown with innodb_fast_shutdown = 1
110203 21:49:46 InnoDB: Starting shutdown...
```

# Sample Output (4): Prepare

```
# innobackupex --apply-log /data/backups/2011-02-03_21-26-08/
```

```
InnoDB Backup Utility v1.5.1-xtrabackup; Copyright 2003, 2009 Innobase Oy.
```

```
All Rights Reserved.
```

```
...
```

```
110203 21:49:40 innobackupex-1.5.1: Starting ibbackup with command: xtrabackup --prepare --
```

```
target-dir=/data/backups/2011-02-03_21-26-08
```

```
xtrabackup Ver 1.4 Rev 193 for 5.1.47 unknown-linux-gnu (x86_64)
```

```
xtrabackup: cd to /data/backups/2011-02-03_21-26-08
```

```
xtrabackup: This target seems to be not prepared yet.
```

```
xtrabackup: xtrabackup_logfile detected: size=2097152,  
start_lsn=(4953882516)
```

```
xtrabackup: Temporary instance for recovery is set as followings.
```

```
xtrabackup: innodb_data_home_dir = ./
```

```
xtrabackup: innodb_data_file_path = ibdata1:10M:autoextend
```

```
xtrabackup: innodb_log_group_home_dir = ./
```

```
xtrabackup: innodb_log_files_in_group = 1
```

```
xtrabackup: innodb_log_file_size = 2097152
```

```
xtrabackup: Starting InnoDB instance for recovery.
```

```
xtrabackup: Using 104857600 bytes for buffer pool (set by --use-memory parameter)
```

```
...
```

```
110203 21:49:40 Percona XtraDB (http://www.percona.com) 1.0.8-11.2 started; log sequence number 4953882700
```

```
[notice (again)]
```

```
If you use binary log and don't use any hack of group commit,
```

```
the binary log position seems to be:
```

```
InnoDB: Last MySQL binlog file position 0 106, file name /var/log/mysql/mysql-bin.000116
```

# Sample Output (4): Prepare

```
# innobackupex-1.5.1 --apply-log /data/backups/2011-02-03_21-26-08/
```

```
InnoDB Backup Utility v1.5.1-xtrabackup; Copyright 2003, 2009 Innobase Oy.
```

```
All Rights Reserved.
```

```
...
```

```
110203 21:49:40 innobackupex-1.5.1: Starting ibbackup with command: xtrabackup --prepare --target-dir=/data/backups/2011-02-03_21-26-08
```

```
xtrabackup Ver 1.4 Rev 193 for 5.1.47 unknown-linux-gnu (x86_64)
```

```
xtrabackup: cd to /data/backups/2011-02-03_21-26-08
```

```
xtrabackup: This target seems to be not prepared yet.
```

```
xtrabackup: xtrabackup_logfile detected: size=2097152, start_lsn=(4953882516)
```

```
xtrabackup: Temporary instance for recovery is set as followings.
```

```
xtrabackup: innodb_data_home_dir = ./
```

```
xtrabackup: innodb_data_file_path = ibdata1:10M:autoextend
```

```
xtrabackup: innodb_log_group_home_dir = ./
```

```
xtrabackup: innodb_log_files_in_group = 1
```

```
xtrabackup: innodb_log_file_size = 2097152
```

```
xtrabackup: Starting InnoDB instance for recovery.
```

```
xtrabackup: Using 104857600 bytes for buffer pool (set by --use-memory parameter)
```

```
...
```

```
110203 21:49:40 Percona XtraDB (http://www.percona.com) 1.0.8-11.2 started; log sequence number 4953882700
```

```
[notice (again)]
```

```
If you use binary log and don't use any hack of group commit,  
the binary log position seems to be:
```

```
InnoDB: Last MySQL binlog file position 0 106, file name /var/log/
```

```
mysql/mysql-bin.000116
```

```
xtrabackup: starting shutdown with innodb_fast_shutdown = 1
```

# Sample Output (5): Prepare

110203 21:49:50 innobackupex-1.5.1: Restarting xtrabackup with command: **xtrabackup --prepare**

**--target-dir=/data/backups/2011-02-03\_21-26-08** for creating  
ib\_logfile\*

xtrabackup Ver 1.4 Rev 193 for 5.1.47 unknown-linux-gnu (x86\_64)  
xtrabackup: cd to /data/backups/2011-02-03\_21-26-08  
xtrabackup: This target seems to be already prepared.  
xtrabackup: notice: xtrabackup\_logfile was already used to '--prepare'.

...  
110203 21:49:50 InnoDB: Log file ./ib\_logfile0 did not exist: new to be created

InnoDB: Setting log file **./ib\_logfile0 size to 5 MB**

InnoDB: Database physically writes the file full: wait...

110203 21:49:50 InnoDB: Log file ./ib\_logfile1 did not exist: new to be created

InnoDB: Setting log file **./ib\_logfile1 size to 5 MB**

InnoDB: Database physically writes the file full: wait...

InnoDB: The log sequence number in ibdata files does not match

InnoDB: the log sequence number in the ib\_logfiles!

110203 21:49:51 InnoDB: Database was not shut down normally!

InnoDB: Starting crash recovery.

InnoDB: Reading tablespace information from the .ibd files...

InnoDB: Last MySQL binlog file position 0 106, file name /var/log/mysql/mysql-bin.000116

110203 21:49:51 Percona XtraDB (<http://www.percona.com>) 1.0.8-11.2 started; log sequence number 4953882936

[notice (again)]

If you use binary log and don't use any hack of group commit,  
the binary log position seems to be:

InnoDB: Last MySQL **binlog file position 0 106, file name /var/log/mysql/mysql-bin.000116**

# Sample Output (5): Prepare

```
110203 21:49:50 innobackupex-1.5.1: Restarting xtrabackup with command: xtrabackup --prepare --target-dir=/data/  
backups/2011-02-03_21-26-08 for creating ib_logfile*  
xtrabackup Ver 1.4 Rev 193 for 5.1.47 unknown-linux-gnu (x86_64)  
xtrabackup: cd to /data/backups/2011-02-03_21-26-08  
xtrabackup: This target seems to be already prepared.  
xtrabackup: notice: xtrabackup_logfile was already used to '--prepare'.  
...  
110203 21:49:50 InnoDB: Log file ./ib_logfile0 did not exist: new to be created  
InnoDB: Setting log file ./ib_logfile0 size to 5 MB  
InnoDB: Database physically writes the file full: wait...  
110203 21:49:50 InnoDB: Log file ./ib_logfile1 did not exist: new to be created  
InnoDB: Setting log file ./ib_logfile1 size to 5 MB  
InnoDB: Database physically writes the file full: wait...  
InnoDB: The log sequence number in ibdata files does not match  
InnoDB: the log sequence number in the ib_logfiles!  
110203 21:49:51 InnoDB: Database was not shut down normally!  
InnoDB: Starting crash recovery.  
InnoDB: Reading tablespace information from the .ibd files...  
InnoDB: Last MySQL binlog file position 0 106, file name /var/log/mysql/mysql-bin.000116  
110203 21:49:51 Percona XtraDB (http://www.percona.com) 1.0.8-11.2 started; log sequence number 4953882936  
[notice (again)]  
If you use binary log and don't use any hack of group commit,  
the binary log position seems to be:  
InnoDB: Last MySQL binlog file position 0 106, file name /var/log/mysql/mysql-bin.000116  
xtrabackup: starting shutdown with innodb_fast_shutdown = 1  
110203 21:49:51 InnoDB: Starting shutdown...  
110203 21:49:56 InnoDB: Shutdown completed; log sequence number 4953882636  
  
110203 21:49:56 innobackupex-1.5.1: completed OK!
```

# Sample Output (6): Restore

```
# innobackupex --copy-back /data/backups/  
2011-02-03_21-26-08/
```

```
InnoDB Backup Utility v1.5.1-xtrabackup; Copyright 2003, 2009 Innobase Oy.  
All Rights Reserved.
```

```
...
```

```
innobackupex-1.5.1: Starting to copy MyISAM tables, indexes,  
innobackupex-1.5.1: .MRG, .TRG, .TRN, .ARM, .ARZ, .opt, and .frm  
files
```

```
innobackupex-1.5.1: in '/data/backups/2011-02-03_21-26-08'  
innobackupex-1.5.1: back to original data directory '/var/lib/mysql'  
innobackupex-1.5.1: Copying directory '/data/backups/2011-02-03_21-26-08/test'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/xtrabackup_binlog_pos_innodb'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/xtrabackup_binary'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/xtrabackup_checkpoints'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/xtrabackup_binlog_info'  
innobackupex-1.5.1: Copying directory '/data/backups/2011-02-03_21-26-08/sakila'  
innobackupex-1.5.1: Copying directory '/data/backups/2011-02-03_21-26-08/mysql'
```

```
innobackupex-1.5.1: Starting to copy InnoDB tables and indexes  
innobackupex-1.5.1: in '/data/backups/2011-02-03_21-26-08'  
innobackupex-1.5.1: back to original InnoDB data directory '/var/lib/mysql'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/ibdata1'
```

```
innobackupex-1.5.1: Starting to copy InnoDB log files
```

# Sample Output (6): Restore

```
# innobackupex-1.5.1 --copy-back /data/backups/2011-02-03_21-26-08/  
InnoDB Backup Utility v1.5.1-xtrabackup; Copyright 2003, 2009 Innobase Oy.  
All Rights Reserved.
```

```
...
```

```
innobackupex-1.5.1: Starting to copy MyISAM tables, indexes,  
innobackupex-1.5.1: .MRG, .TRG, .TRN, .ARM, .ARZ, .opt, and .frm files  
innobackupex-1.5.1: in '/data/backups/2011-02-03_21-26-08'  
innobackupex-1.5.1: back to original data directory '/var/lib/mysql'  
innobackupex-1.5.1: Copying directory '/data/backups/2011-02-03_21-26-08/test'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/xtrabackup_binlog_pos_innodb'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/xtrabackup_binary'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/xtrabackup_checkpoints'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/xtrabackup_binlog_info'  
innobackupex-1.5.1: Copying directory '/data/backups/2011-02-03_21-26-08/sakila'  
innobackupex-1.5.1: Copying directory '/data/backups/2011-02-03_21-26-08/mysql'
```

```
innobackupex-1.5.1: Starting to copy InnoDB tables and indexes
```

```
innobackupex-1.5.1: in '/data/backups/2011-02-03_21-26-08'  
innobackupex-1.5.1: back to original InnoDB data directory '/var/lib/mysql'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/ibdata1'
```

```
innobackupex-1.5.1: Starting to copy InnoDB log files
```

```
innobackupex-1.5.1: in '/data/backups/2011-02-03_21-26-08'  
innobackupex-1.5.1: back to original InnoDB log directory '/var/lib/mysql'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/ib_logfile1'  
innobackupex-1.5.1: Copying file '/data/backups/2011-02-03_21-26-08/ib_logfile0'  
innobackupex-1.5.1: Finished copying back files.
```

# Topics

---

- Comparison of MySQL Backup solutions
- How XtraBackup works
- **Features**

# Features

- Streaming
- Compression
- Incremental/Delta
- Filtering
- Throttling
- Parallel Copying
- Statistics
- Export/Import Tables

# Features

- **Streaming**
- Compression
- Incremental/Delta
- Filtering
- Throttling
- Parallel Copying
- Statistics
- Export/Import Tables

# Streaming Backups

- Uses tar4ibd by default
- Example:  

```
# innobackupex --stream=tar ./ > backup.tar
```
- Or to a remote host:  

```
# innobackupex --stream=tar ./ | ssh user@desthost \  
"cat - > /data/backups/backup.tar"
```
- Remote copy using netcat:  

```
# ssh user@desthost "( nc -l 9999 > /data/backups/  
backup.tar & )" && innobackupex --stream=tar ./ | \  
nc desthost 9999
```
- Extract with 'tar -i'!

# Features

- Streaming
- **Compression**
- Incremental/Delta
- Filtering
- Throttling
- Parallel Copying
- Statistics
- Export/Import Tables

# Compression

- Use streaming and pipe to compression script
- Utilities: pigz, qpress, LZOP, pbzip2, LZMA
- Example:

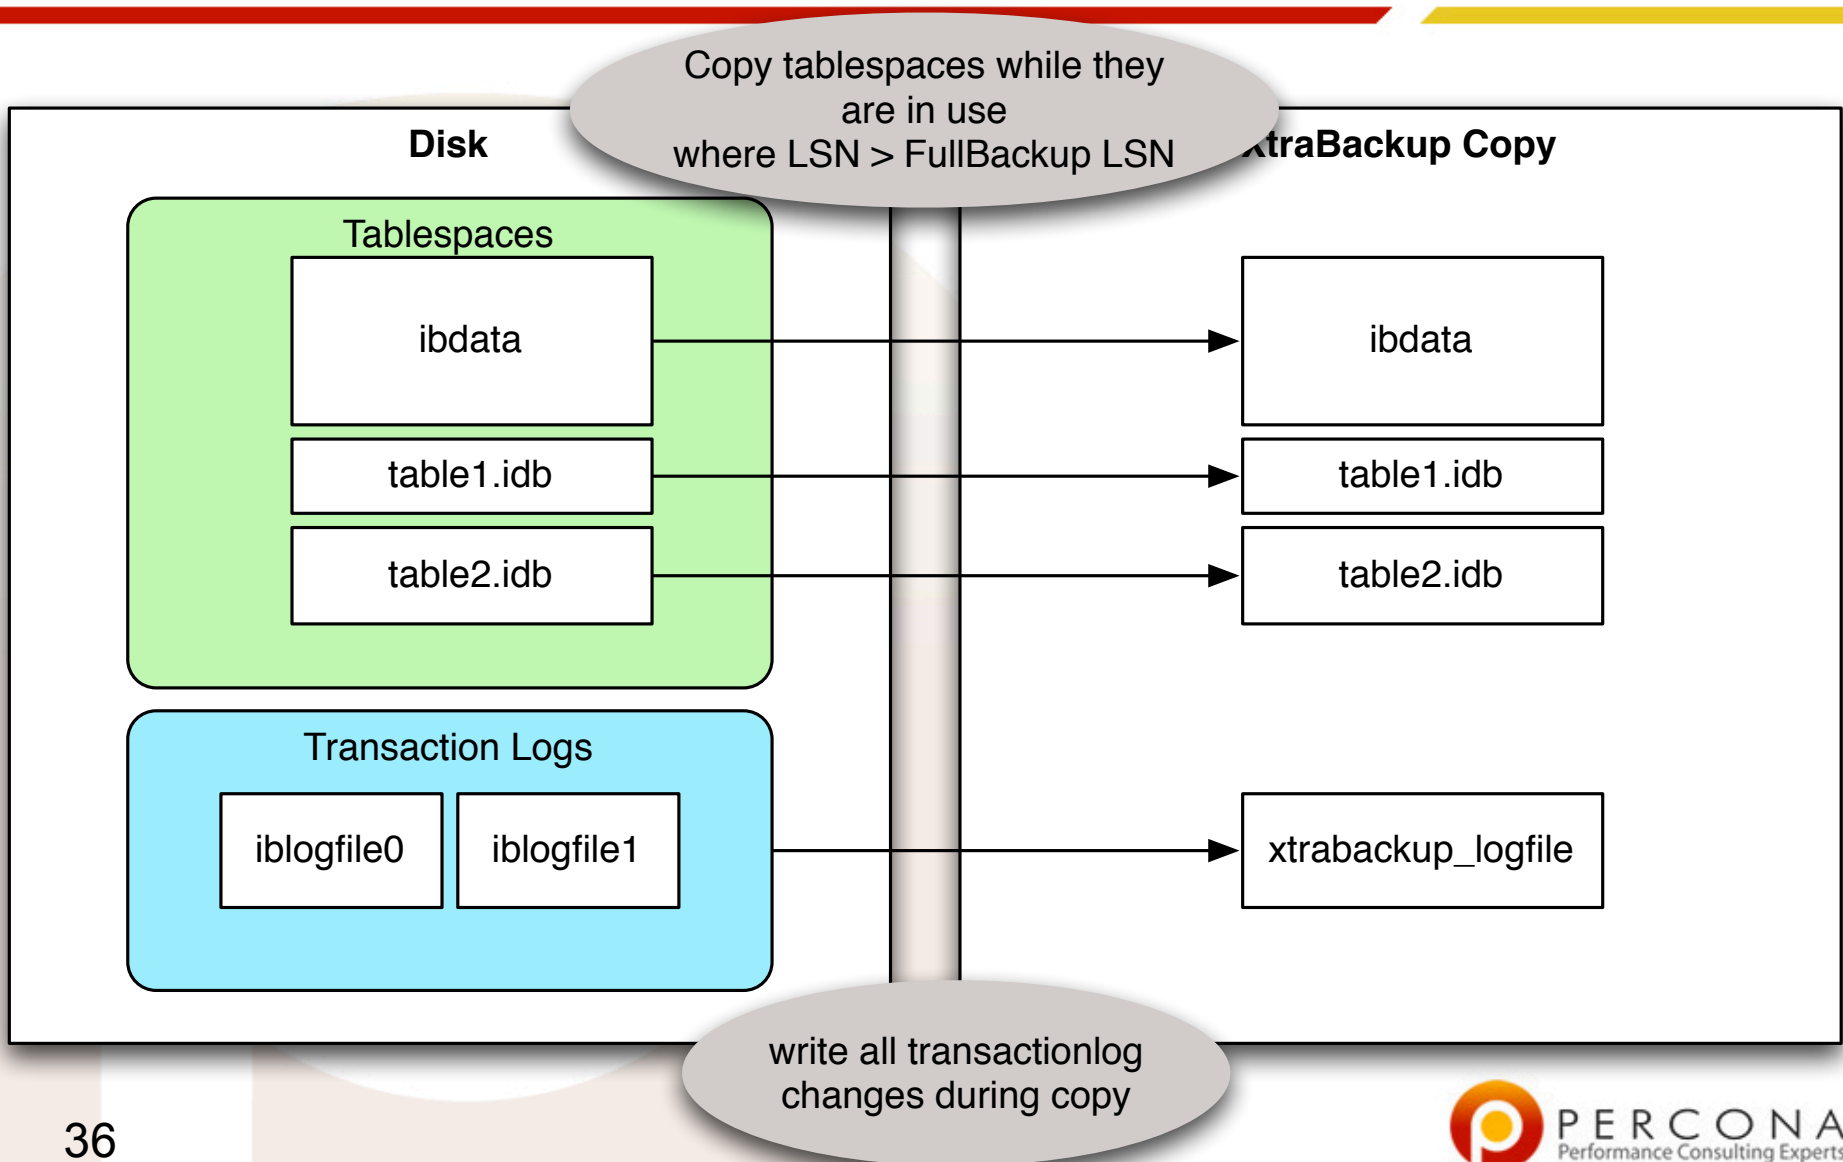
```
# innobackupex --stream=tar ./ | qpress -io \  
xtrabackup.tar > backup.tar.qpress
```
- Can be combined to stream to other host:

```
# ssh user@desthost "( nc -l 9999 | qpress -dio | tar -i  
> /data/backups/ backup.tar & )" && \  
innobackupex --stream=tar ./ | qpress -io backup.tar | \  
nc desthost 9999
```

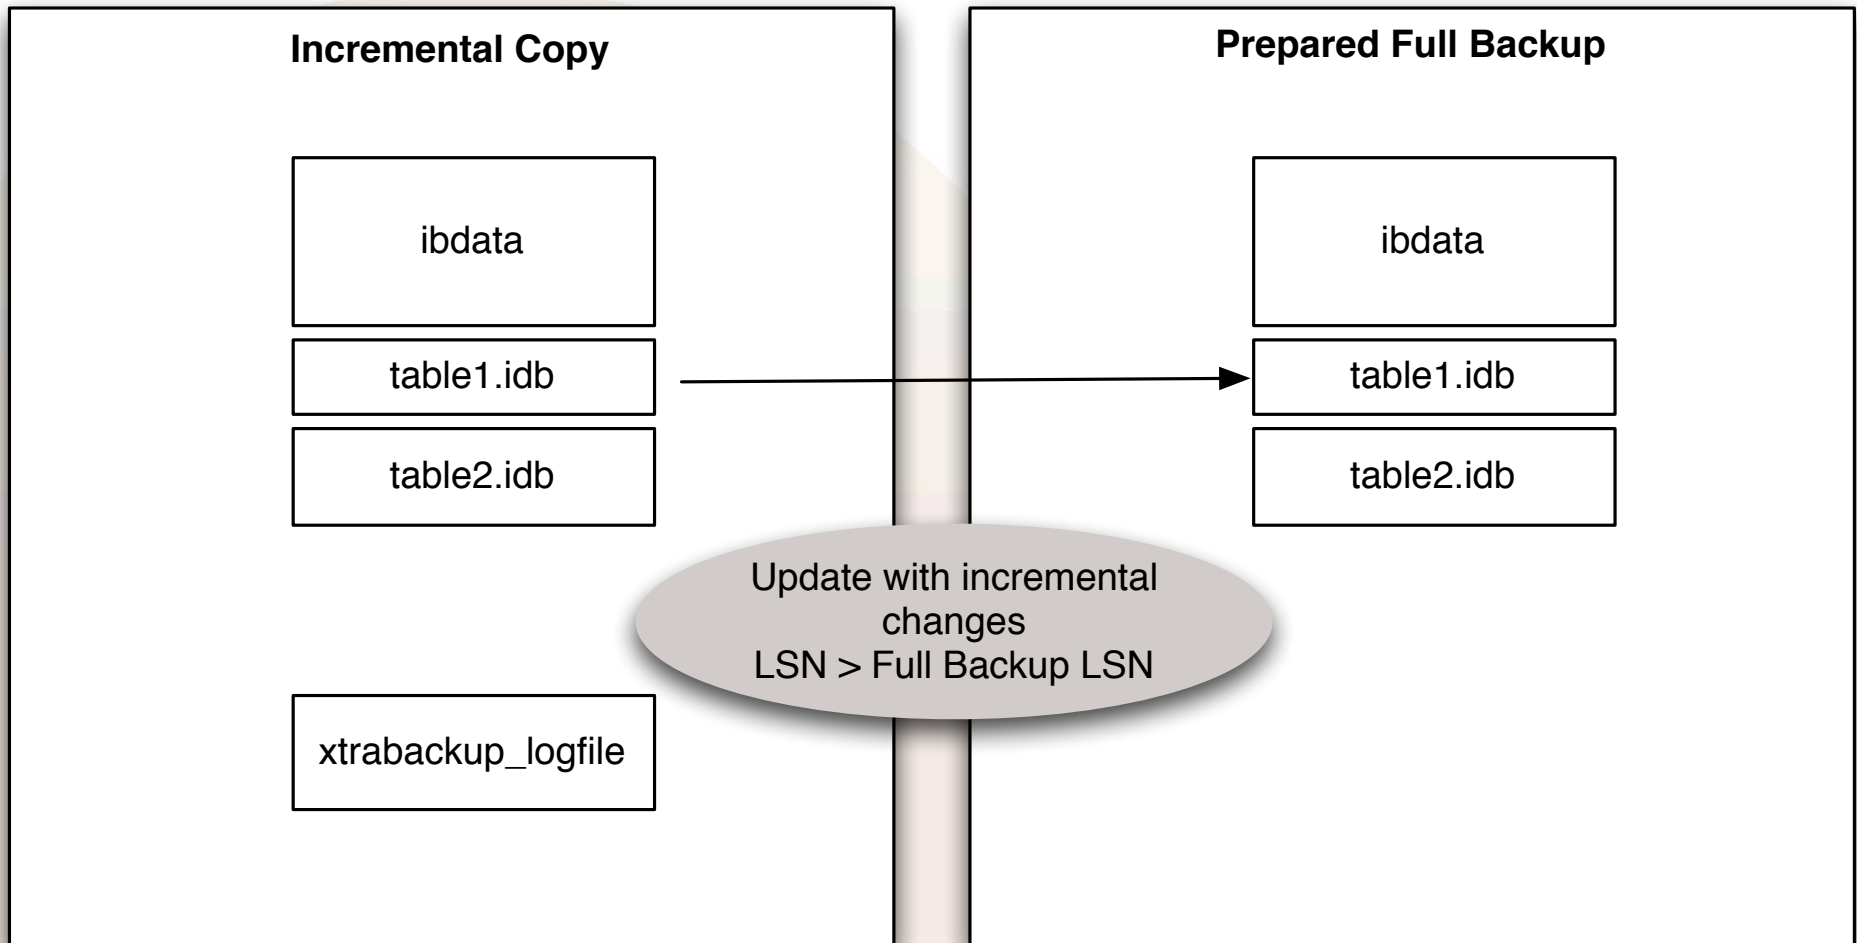
# Features

- Streaming
- Compression
- **Incremental/Delta**
- Filtering
- Throttling
- Parallel Copying
- Statistics
- Export/Import Tables

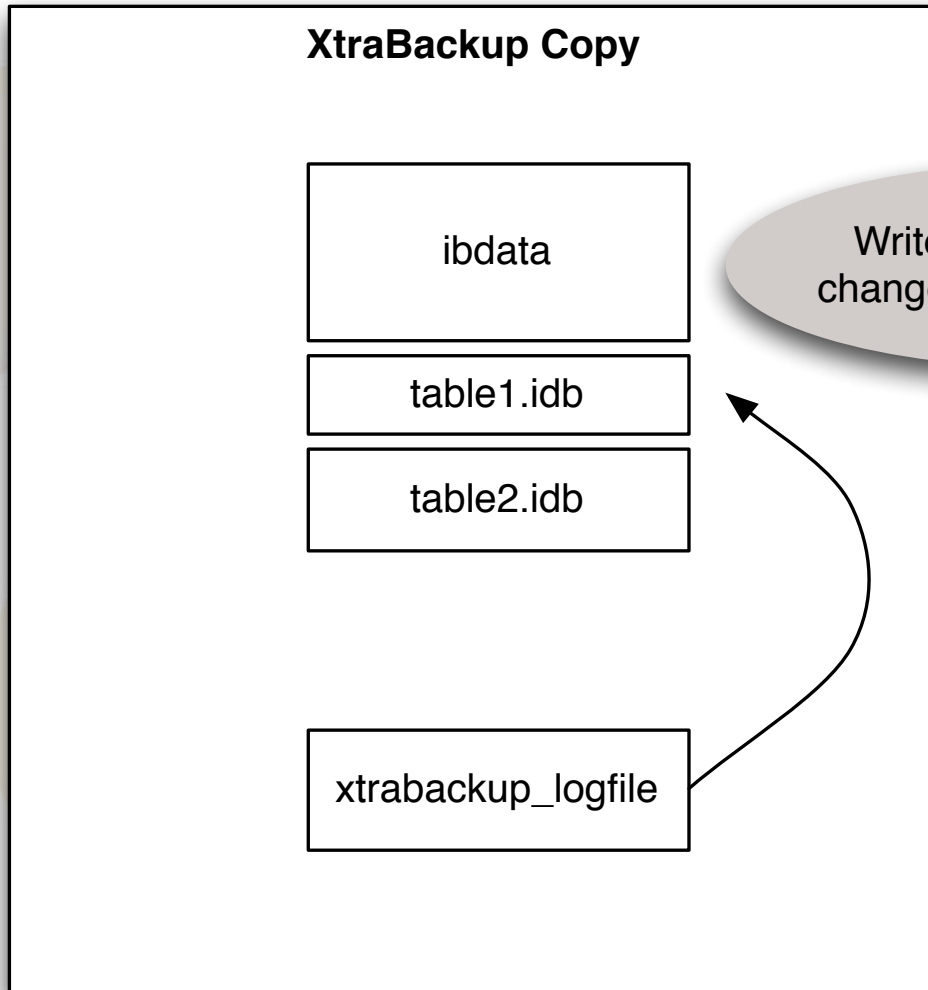
# Incremental/Delta Backups



# Incremental/Delta Backups



# Incremental/Delta Backups



Write all transaction log changes to the tablespaces

# Incremental/Delta Backups

- Example:

```
# xtrabackup --backup --target-dir=/data/backups/inc1 \  
--incremental-basedir=/data/backups/base \  
--datadir=/var/lib/mysql/
```

```
# xtrabackup --prepare --apply-log-only \  
--target-dir=/data/backups/base \  
--incremental-dir=/data/backups/inc1
```

# Incremental/Delta Backups

- Copy only parts of the tablespaces where LSN is newer than the previous backup
- Notes:
  - Currently only possible with xtrabackup, not innobackupex
  - Schema changes are not allowed
  - base backup should be prepared, but without rollback phase
    - innobackupex --apply-log --redo-only
    - xtrabackup --prepare --apply-log-only

# Features

- Streaming
- Compression
- Incremental/Delta
- **Filtering**
- **Throttling**
- **Parallel Copying**
- Statistics
- Export/Import Tables

# Filtering/Throttling/Parallel Copying

- Filtering, only copying certain tablespaces:  
**--tables, --databases, --tables-file**
- Throttling:
  - Limit to 100 read/writes IOs:  
**--throttle=100**
  - In streaming backup, use **pv** to limit to 10MB/s  
# innobackupex --stream=tar ./ | \  
**pv -q -L10m | \  
ssh user@desthost "cat - > /data/backups/backup.tar"**
- Parallel copying of the tablespaces (beta):  
**--parallel 2 # 2 threads**

# Features

- Streaming
- Compression
- Incremental/Delta
- Filtering
- Throttling
- Parallel Copying
- **Statistics**
- Export/Import Tables

# Statistics

From CREATE TABLE:

```
PRIMARY KEY (`link_id`),  
UNIQUE KEY `domain_id_2` (`domain_id`,`link_id`),
```

# xtrabackup --stats

<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

From CREATE TABLE:

```
PRIMARY KEY (`link_id`),  
UNIQUE KEY `domain_id_2` (`domain_id`,`link_id`),
```

# xtrabackup --stats

## <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%

45 leaf pages: recs=**25958413**, pages=23125, data=**337459369 bytes**, data/pages=**89%**

# Statistics

- Script available to make output more readable:

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%

# Features

- Streaming
- Compression
- Incremental/Delta
- Filtering
- Throttling
- Parallel Copying
- Statistics
- **Export/Import Tables**

# Export/Import Tables

- Only possible with `innodb_file_per_table=1`
- Export on both InnoDB/XtraDB
- Import only on XtraDB with `innodb_expand_import=1`

# Export Tables

```
# xtrabackup --prepare --export --innodb-file-per-table --  
target-dir=/data/backups/2011-02-03_21-26-08/
```

```
...
```

```
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
```

```
...
```

```
# ls -l {payment,rental}.{exp,ibd,frm}
```

```
16K payment.exp
```

```
12K payment.frm
```

```
10M payment.ibd
```

```
16K rental.exp
```

```
12K rental.frm
```

```
10M rental.ibd
```

# Import Tables

```
mysql> CREATE TABLE `rental` (  
  `rental_id` int(11) NOT NULL AUTO_INCREMENT,  
  `rental_date` datetime NOT NULL,  
  `inventory_id` mediumint(8) unsigned NOT NULL,  
  `customer_id` smallint(5) unsigned NOT NULL,  
  `return_date` datetime DEFAULT NULL,  
  `staff_id` tinyint(3) unsigned NOT NULL,  
  `last_update` timestamp NOT NULL DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP,  
  PRIMARY KEY (`rental_id`),  
  UNIQUE KEY `rental_date` (`rental_date`,`inventory_id`,`customer_id`),  
  KEY `idx_fk_inventory_id` (`inventory_id`),  
  KEY `idx_fk_customer_id` (`customer_id`),  
  KEY `idx_fk_staff_id` (`staff_id`),  
  CONSTRAINT `fk_rental_customer` FOREIGN KEY (`customer_id`) REFERENCES `customer` (`customer_id`) ON  
    UPDATE CASCADE,  
  CONSTRAINT `fk_rental_inventory` FOREIGN KEY (`inventory_id`) REFERENCES `inventory` (`inventory_id`) ON  
    UPDATE CASCADE,  
  CONSTRAINT `fk_rental_staff` FOREIGN KEY (`staff_id`) REFERENCES `staff` (`staff_id`) ON UPDATE CASCADE  
) ENGINE=InnoDB AUTO_INCREMENT=16050 DEFAULT CHARSET=utf8;  
mysql> SET FOREIGN_KEY_CHECKS=0;  
mysql> SET GLOBAL innodb_expand_import=1;  
mysql> ALTER TABLE sakila.rental DISCARD TABLESPACE;  
# cp /data/backups/2011-02-03_21-26-08/sakila/rental.{exp,ibd} /var/lib/mysql/sakila/  
# chown -R mysql:mysql /var/lib/mysql/sakila/rental.{exp,ibd}  
mysql> ALTER TABLE sakila.rental IMPORT TABLESPACE;  
mysql> select count(*) from sakila.rental\G  
***** 1. row *****  
count(*): 16044  
mysql> SET FOREIGN_KEY_CHECKS=1;
```

# Import Tables

```
mysql> CREATE TABLE `rental` (  
  `rental_id` int(11) NOT NULL AUTO_INCREMENT,  
  ...  
) ENGINE=InnoDB AUTO_INCREMENT=16050 DEFAULT CHARSET=utf8;  
mysql> SET FOREIGN_KEY_CHECKS=0;  
mysql> SET GLOBAL innodb_expand_import=1;  
mysql> ALTER TABLE sakila.rental DISCARD  
  TABLESPACE;  
# cp /data/backups/2011-02-03_21-26-08/sakila/rental.  
  {exp,ibd} /var/lib/mysql/sakila/  
# chown -R mysql:mysql /var/lib/mysql/sakila/rental.  
  {exp,ibd}  
mysql> ALTER TABLE sakila.rental IMPORT  
  TABLESPACE;  
mysql> SET FOREIGN_KEY_CHECKS=1;
```

# XtraBackup

- Complete backup solution for InnoDB/XtraDB
- Open Source (GPL)
- Commercial Support Available:  
<http://www.percona.com>
- A lot of Features:  
Streaming, Compression, Incremental, Filtering,  
Throttling, Parallel Copying, Statistics, Export/Import

# XtraBackup

- Website: <http://www.percona.com/software/percona-xtrabackup/>
- Downloads: <http://www.percona.com/downloads/>
- Apt/Yum Repos: <http://www.percona.com/docs/wiki/repositories:start>
- Bugs/Code: <https://launchpad.net/percona-xtrabackup>
- IRC: Freenode #percona
- Mailinglist: <http://groups.google.com/group/percona-discussion>
- <http://www.percona.com>
- <http://www.mysqlperformanceblog.com>

# Percona Live MySQL Conference NYC

May 26, 2011

<http://www.percona.com/live/nyc-2011/>



PERCONA  
LIVE