

# Introduction to InnoDB Hot Backup

## Fast, Consistent, Online Backups for MySQL

Heikki Tuuri  
CEO, Innobase Oy  
VP, Development  
Oracle

Ken Jacobs  
VP Product Strategy  
Oracle

**INNOBASE**





# Today's Topics

- What is InnoDB Hot Backup?
- Backing up InnoDB data - how it works
- Backing up selected InnoDB tables
- Backing up MyISAM data w/ Innobackup
- Restoring a backup
- Hot Backup performance
- Availability and compatibility
- Comparison with other backup tools
- Future directions

**INNOBASE**




# What is InnoDB Hot Backup?

- Backup tool for InnoDB tables and indexes
  - And with free Innobackup script, can backup MyISAM data too!
- Backs up running databases without locking InnoDB data; allows normal database access
- Backups stored on local disk
- Command-line user interface
- Features:
  - Compressed backups
  - Partial backups
  - Point-in-time recovery
- Easy to install and use

**INNOBASE**

# How InnoDB Hot Backup Works



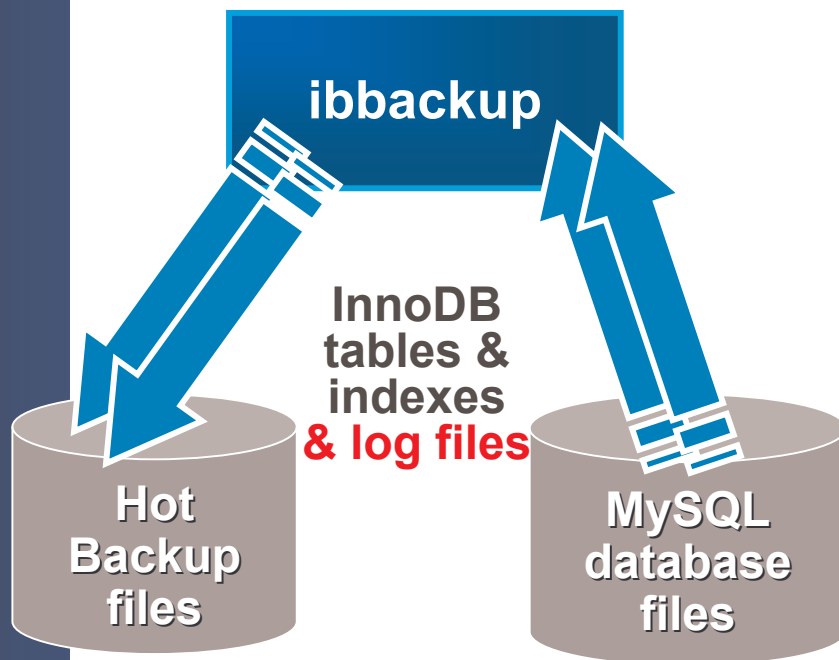
- Copies and compresses InnoDB data files (ibdata & .ibd), even as they are changing

➔ Produces “fuzzy” backup

- backup does not correspond to any specific InnoDB log sequence number (LSN), because different database pages were copied at different times

**INNOBASE**


# How InnoDB Hot Backup Works



- Copies and compresses InnoDB data files (ibdata & .ibd), even as they are changing
- ➔ Produces “fuzzy” backup
  - backup does not correspond to any specific InnoDB log sequence number (LSN), because different database pages were copied at different times
- Copies InnoDB log records generated during data file copy
  - saves all redo records with LSNs covering duration of data file copy
  - redo used later to restore copy to a consistent state

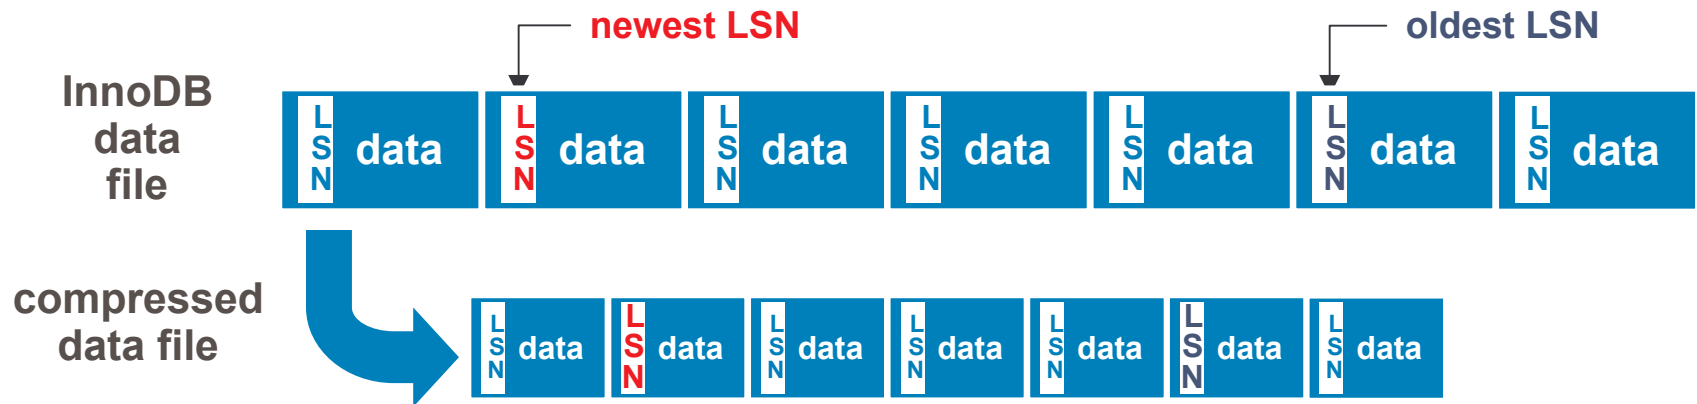
**INNOBASE**

# Backing up InnoDB Data Files



- Innobackup copies and compresses InnoDB data files
  - Includes system tablespace (ibdata) files and single-table tablespaces (.ibd) files
- Omits unused storage in each block, empty pages
- *Typically, backup file is only 30% the size of the original data*

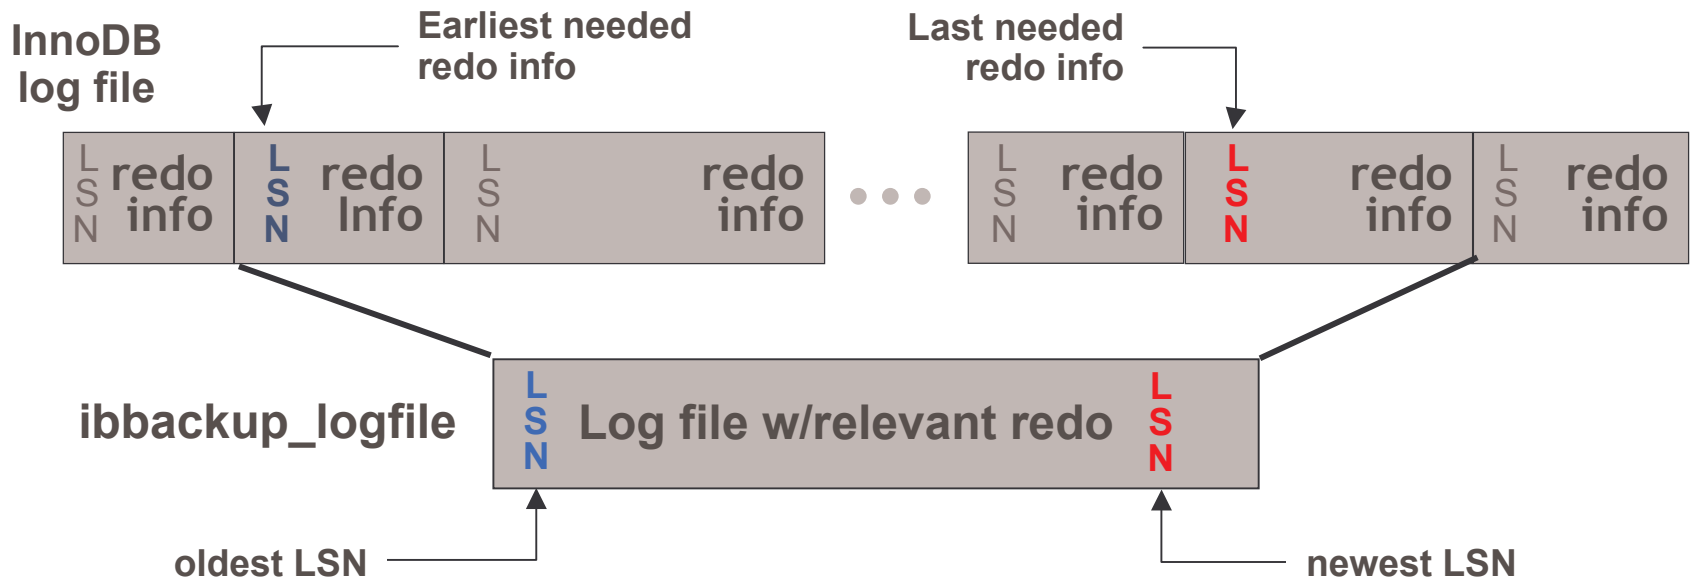
# Backing up InnoDB Data Files



- Innobackup copies and compresses InnoDB data files
  - Includes system tablespace (ibdata) files and single-table tablespaces (.ibd) files
- Omits unused storage in each block, empty pages
- *Typically, backup file is only 30% the size of the original data*
- Notes earliest and latest Log Sequence Number (LSN)

**INNOBASE**

# Backing up InnoDB Log Files




- Innobackup copies the portion of the log file that contains all required redo information
- Covers period from start to end of data file backup
- Can recover all data blocks modified after they were copied to the (compressed) data file

# Partial Backups of InnoDB Data

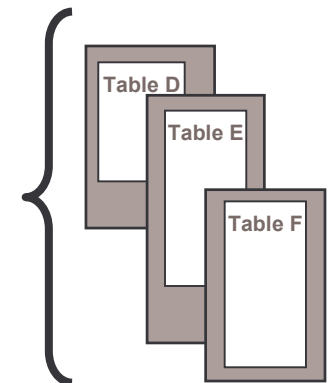
- When using “file per table”, you can backup a subset of InnoDB tables
- Tables included in the backup are specified with regular expressions
- Use the *-- include* option
- Backup contains all tables in system tablespace plus those separate tables that match the pattern

Multiple tables & indexes in the system tablespace (ibdata files)



All tables will be backed up

One table + its indexes per file (.ibd files)



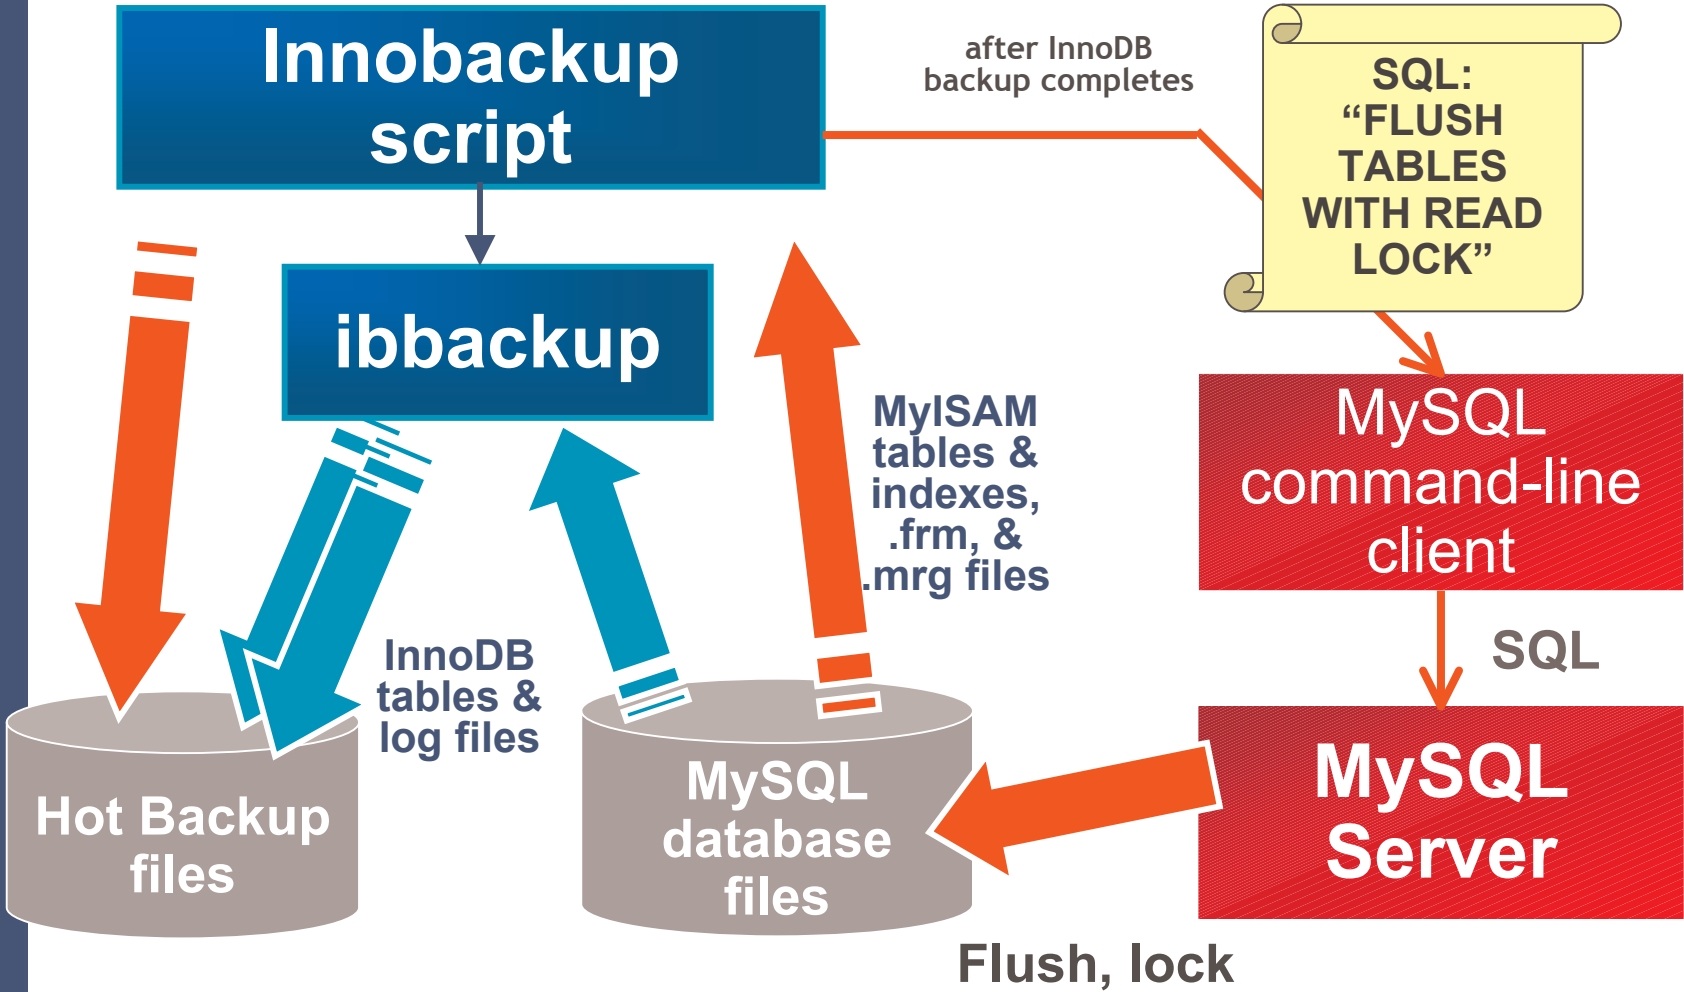
Selective backup possible



# MyISAM Backup with Innobackup

- Innobackup is a free, easy-to-use Perl script
- Used in combination with Hot Backup
  - InnoDB Hot Backup is run as a child process by Innobackup
- Backs up InnoDB data, .frm files, MERGE, and MyISAM data
- “Freezes” updates during MyISAM backup ... which happens just before backup of InnoDB tables completes

# MyISAM Backup with Innobackup



**INNOBASE**

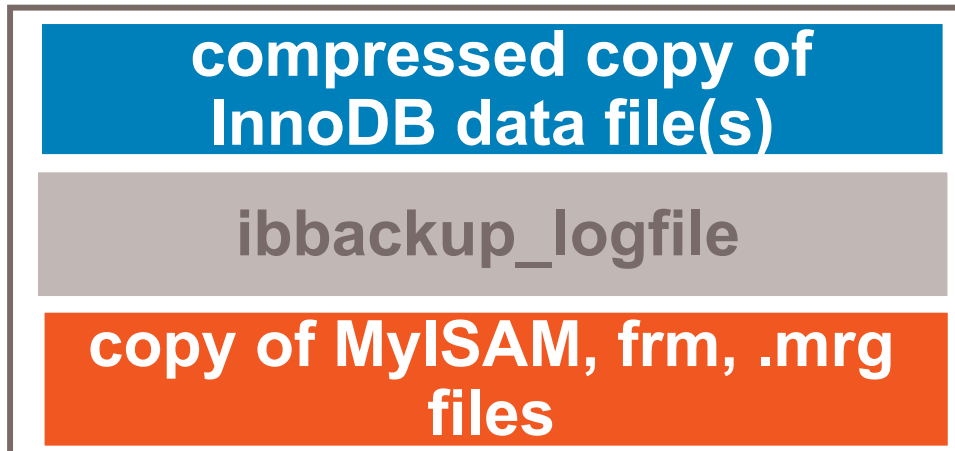


# Things to Think About

- InnoDB tables are fully accessible, including updates, during all phases of the backup
- MyISAM tables can be read and updated during backup of InnoDB tables
- MyISAM tables can be read, but not updated while they are being copied
  - Innobackup uses FLUSH TABLES WITH READ LOCK near the end of the backup run
- Innobackup works best if ...
  - you can wait for insert/update/delete transactions during the backup of MyISAM data
  - your MyISAM tables are small, thus copied quickly

**INNOBASE**

# Backup Files are a “Raw Backup”



Files  
comprising  
raw backup

- Hot Backup accesses the database files while in use to complete the “backup phase”
- Result is a “raw backup” that *cannot* be used by the MySQL server
- You can copy raw backup files to tape, if desired
- Use Innobackup to restore database before use

**INNOBASE**

# Restoring a Database for Use - The Restore (“Apply Log”) Phase

compressed copy of InnoDB  
data file(s)

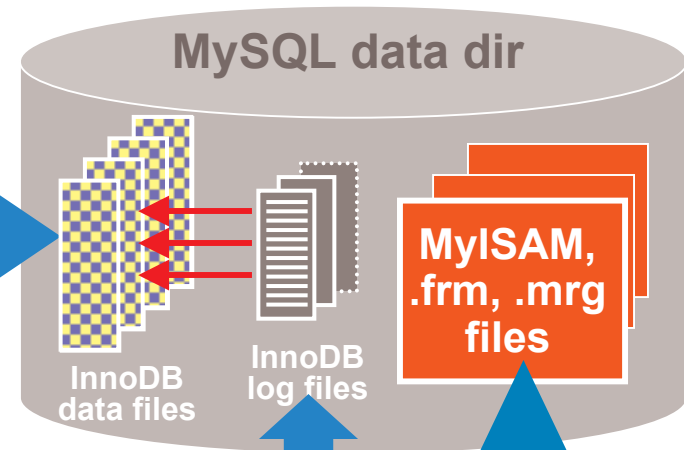
uncompresses InnoDB  
files to data dir

recreates InnoDB  
log files

applies log, so  
InnoDB files are  
consistent

restores MyISAM and  
other files

**INNOBASE**



ibbackup\_logfile

copy of MyISAM, frm, .mrg files



# Restoring a Database for Use - Additional Notes

- Hot Backup / Innobackup rolls forward data files to a common point in time (the time at the end of backup)
- After restore, Hot Backup prints the location in the binlog for the next SQL operation that executed after the backup completed
- Note: the restore phase need not run on database server host
  - You can perform recovery on any machine, and copy recovered files to your database server host



# Restoring a Database for Use - Additional Notes

- To bring the database to the present, apply MySQL binlog after restore phase is complete
- To restore the database to a point before the present, truncate the MySQL binlog at the appropriate SQL statement before applying it
  - Direct the output of mysqlbinlog to a file, instead of piping it to MySQL
  - Output file contains timestamps for all SQL statements in the binlog
  - Edit output file to truncate at desired point
  - Redirect the truncated binlog file to MySQL



# InnoDB Hot Backup Performance

- Backup time is comparable to copying data and log files via shell
- Restore (“apply log”) phase is usually faster than the backup phase
- Performance is disk-bound; all sequential i/o
  - Typical backup speeds are 20 - 40 MB/s
  - Max CPU usage is typically 30% on a uniprocessor
- No practical limit to db size, number of tables
  - Hot Backup now used w/ dbs of 100s of gigabytes
- Best to run backup when update rate is low
  - Because logs are circular, backup must finish before log file wraps



# Availability and Compatibility

- InnoDB Hot Backup v3.0 is available for:
  - Microsoft Windows
  - GNU/Linux on i386, AMD64 (x86\_64, EM64T), and PowerPC
  - FreeBSD 5.x on i386
  - HP-UX 11 on PA-RISC 1.1
  - IBM AIX 4.3 and 5.2
  - Sun Solaris 8 & 10 on 32-bit SPARC (also works on Solaris 9)
  - Sun Solaris 10 on i386
- InnoDB Hot Backup v3.0 is compatible with every MySQL/InnoDB version from 3.23 to 5.1

**INNOBASE**

# Licensing InnoDB Hot Backup

- Licensed per server, except embedded
- Free 30-day evaluation license available
- Please visit [www.innodb.com/hot-backup](http://www.innodb.com/hot-backup)
- Pay by credit card or bank xfer

	Euros	USD
1-year	€ 390	\$520
Perpetual	€ 990	\$1320
1-year email support	€ 590	\$790
Embedded	Request quote	

(Prices as of April 17, 2007)

**INNOBASE**

# Backup Tool Comparison

	InnoDB Hot Backup	MySQL dump	MySQL HotCopy	Zmanda RM for MySQL	Linux LVM	Veritas
Allows updates	<b>Yes</b> (InnoDB tables only)	<b>No</b>	<b>No</b>	<b>No</b>	<b>Yes</b>	<b>Yes</b>
Backs up .frm files	<b>Yes</b>	<b>Yes</b>	<b>Yes</b>	<b>Yes</b>	<b>Yes</b>	<b>Yes</b>
Supported table types	InnoDB, MyISAM, MERGE	All	MyISAM, ARCHIVE	All	All	All
Backup Speed	Very good	Poor	Good	Good	Good	Good
Restore Speed	Very good	Very poor	Poor	Good	Good	Good

**INNOBASE**

# Backup Tool Comparison, cont'd

	InnoDB Hot Backup	MySQL dump	MySQL HotCopy	Zmanda RM for MySQL	Linux LVM	Veritas
Supported platforms	Linux Windows Solaris FreeBSD HP-UX IBM AIX	All	Unix Netware	Linux	Linux	All
Licensing	\$	Free	Free	\$	Free	\$

**INNOBASE**



# INNOBASE

developer of

**INNODB<sup>®</sup>**

and

**INNODB<sup>®</sup> Hot Backup**

**ORACLE<sup>®</sup>**

Innbase is an Oracle company

