



INNODB[®]

Efficient, Reliable Transactional Storage for MySQL

Ken Jacobs
Vice President, Product Strategy
Server Technologies Division,
Manager, InnoDB group
Oracle Corporation

INNOBASE





Today's Topics

- Innobase and InnoDB
 - Who is Innobase?
 - What is InnoDB?
 - Innobase and MySQL
 - Innobase and Oracle
- Introduction to InnoDB
 - Architecture and features
 - Example InnoDB applications
 - Comparison with other storage engines
- InnoDB Hot Backup

INNOBASE



Who is Innobase Oy?

- Founded by Dr. Heikki Tuuri in 2000
- Headquartered in Helsinki, Finland
- Distributed development organization
- Develops InnoDB, which is distributed and supported by MySQL AB
- Develops and distributes InnoDB Hot Backup
- Acquired by Oracle in October, 2005

INNOBASE



What is InnoDB?

- A pluggable storage engine for MySQL
- Provides high performance and reliability for mission-critical applications
 - storage management, transactions and referential integrity
- A standard part of MySQL Enterprise Server and MySQL Community Server editions
- Distributed and supported by MySQL AB
- Most popular & proven transactional engine

INNOBASE

Innodb Oy and MySQL AB



- Both are Nordic companies
 - Innobase: Finland
 - MySQL: Sweden
- Close relationship, since 2000
- InnoDB has been distributed and supported by MySQL AB since 2001
- MySQL Platinum Partner
- Gold Sponsor of the 2007 MySQL User Conference

INNOBASE

Innoblase OY and Oracle

- ✓ Oracle acquires Innoblase OY, in October 2005
- ✓ “Oracle intends to continue developing InnoDB”
- ✓ “... and expand our commitment to open source”
- ✓ “Oracle fully expects to negotiate an extension of the relationship between Innoblase and MySQL”
 - OEM contract renewed in April 2006



“Business As Usual”

INNOBASE

InnoDB - Transactions and Reliability for MySQL



- Full transaction support, high concurrency
- High performance memory and i/o architecture
- Flexible storage options
- Efficient indexing and declarative referential integrity to protect foreign-key relationships
- Online backup with InnoDB Hot Backup

INNOBASE

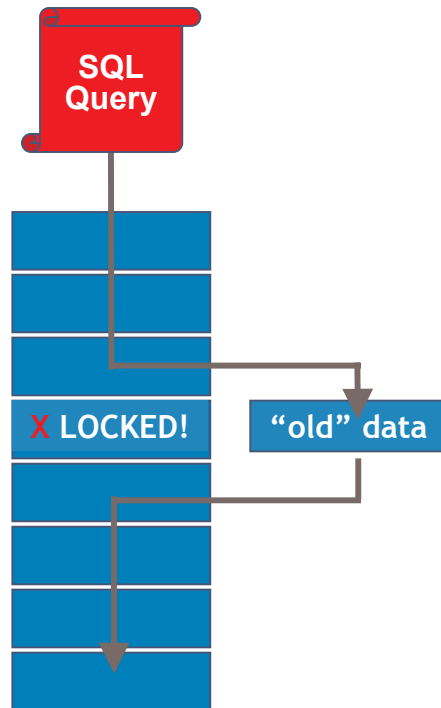


InnoDB Architecture

- Modeled on Gray & Reuter's text *Transaction Processing: Concepts & Techniques*
- Emulates Oracle architecture
 - Multi-version concurrency control provides higher throughput
 - Undo info in the database, not the logs
 - Tablespaces for data, index storage

INNOBASE

InnoDB Transactions & Locking



- Full transaction support
 - atomicity, consistency, isolation, durability
- SQL-standard isolation levels
 - SERIALIZABLE
 - REPEATABLE READ
 - READ COMMITTED
 - RED UNCOMMITTED
- Unlimited row-level locking
- Multi-version read-consistency
- Automatic deadlock detection

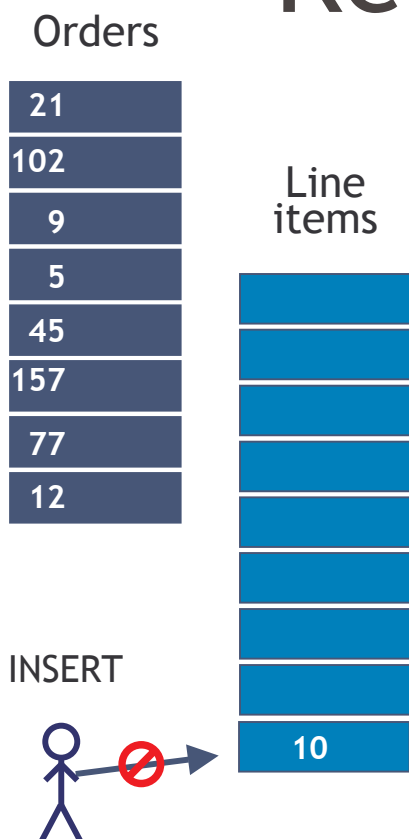
INNOBASE



InnoDB Transactions

- **A**tomic - all changes are either committed as a group, or all are rolled back as a group
- **C**onsistent - transactions operate on a consistent view of the data, leaving the data in a consistent state (by transaction's end)
- **I**solated - each transaction “thinks” it is running by itself - effects of other transactions are invisible until it commits
- **D**urable - once committed, all changes persist, even if there are system failures

InnoDB and Referential Integrity



- Use SQL to define relationships between tables (foreign keys)
- InnoDB prevents changes that violate those relationships
- Supports SQL clauses in CREATE TABLE
 - ON UPDATE RESTRICT | CASCADE | SET NULL
 - ON DELETE RESTRICT | CASCADE | SET NULL
- Protects data integrity

PAYCOM.net Relies on InnoDB



on-line credit card and
check processing

“Without InnoDB’s transactions and efficient row locking, we would not be able to use MySQL as the basis for our company’s operations. ... When the inevitable hardware issues do arise, with the InnoDB storage engine and the InnoDB Hot Backup we can recover our database in a matter of minutes, and be back online ready to work.”

**Clarke Retzer,
Chief Systems Engineer**

- 100’s of millions of rows
- all InnoDB data is stored in a single 1/2 terabyte file

INNOBASE

RightNow Relies on InnoDB



online customer relationship
management solutions

“RightNow is a longtime MySQL user. We converted to InnoDB a few years ago, which allowed us to scale our application much higher than was previously possible. We are pleased with the overall performance and reliability that InnoDB provides.”

**Ryan Huddleston,
DBA Manager**

- 25 TB of production data
- 20 million page turns per day
- 4 billion queries per week

INNOBASE

Flickr, Google and Wikipedia Rely on InnoDB



photo-sharing service

“InnoDB is the MySQL storage engine of choice for Flickr, due to its utmost stability and speed. With InnoDB, we are able to achieve billions of transactions per day at a fraction of the cost of other database technologies. InnoDB is able to meet the demands of our highly mixed traffic while offering the ability to backup data without interruption of service.”

Dathan Vance Pattishall

Google

“Google runs critical business systems with InnoDB and MySQL. The systems require 24×7 operation with minimal downtime. The systems support large OLTP and reporting workloads. We are very happy with the scalability, reliability and manageability of this software.”



WIKIPEDIA
The Free Encyclopedia

“The InnoDB-based MySQL solution is critical in building efficient high-performance database environments, especially when scale-out strategies are applied.”

Domas Mituzas,
DBA & Site Engineer

INNOBASE

Choosing a Storage Engine

	InnoDB	MyISAM	Other
Transactions?	Yes	No	?
Row-level locking?	Yes	No	?
Multi-version concurrency?	Yes	No	?
Serializable transactions? (safe for SQL-based replication)	Yes	Yes (table locks)	?
Referential integrity?	Yes	No	?
Reliable disk writes? (double-write)	Yes	Partial	?
Efficient disk i/o? (async and direct i/o)	Yes	No	?
Flexible storage options? (one/multiple tables/file)	Yes	No	?

Choosing a Storage Engine

	InnoDB	MyISAM	Other
Primary key data clustering?	Yes	No	?
Queries resolved in index?	Yes	Yes	?
Consistent online (“hot”) backup?	Yes (separate product)	No	?
Compression?	coming soon	Partial (readonly)	?
Distributed by MySQL?	Yes	Yes	No
Supported by MySQL?	Yes	Yes	No
User forums, bug database at MySQL.COM?	Yes	Yes	No
Industry-proven?	Yes	Yes	No

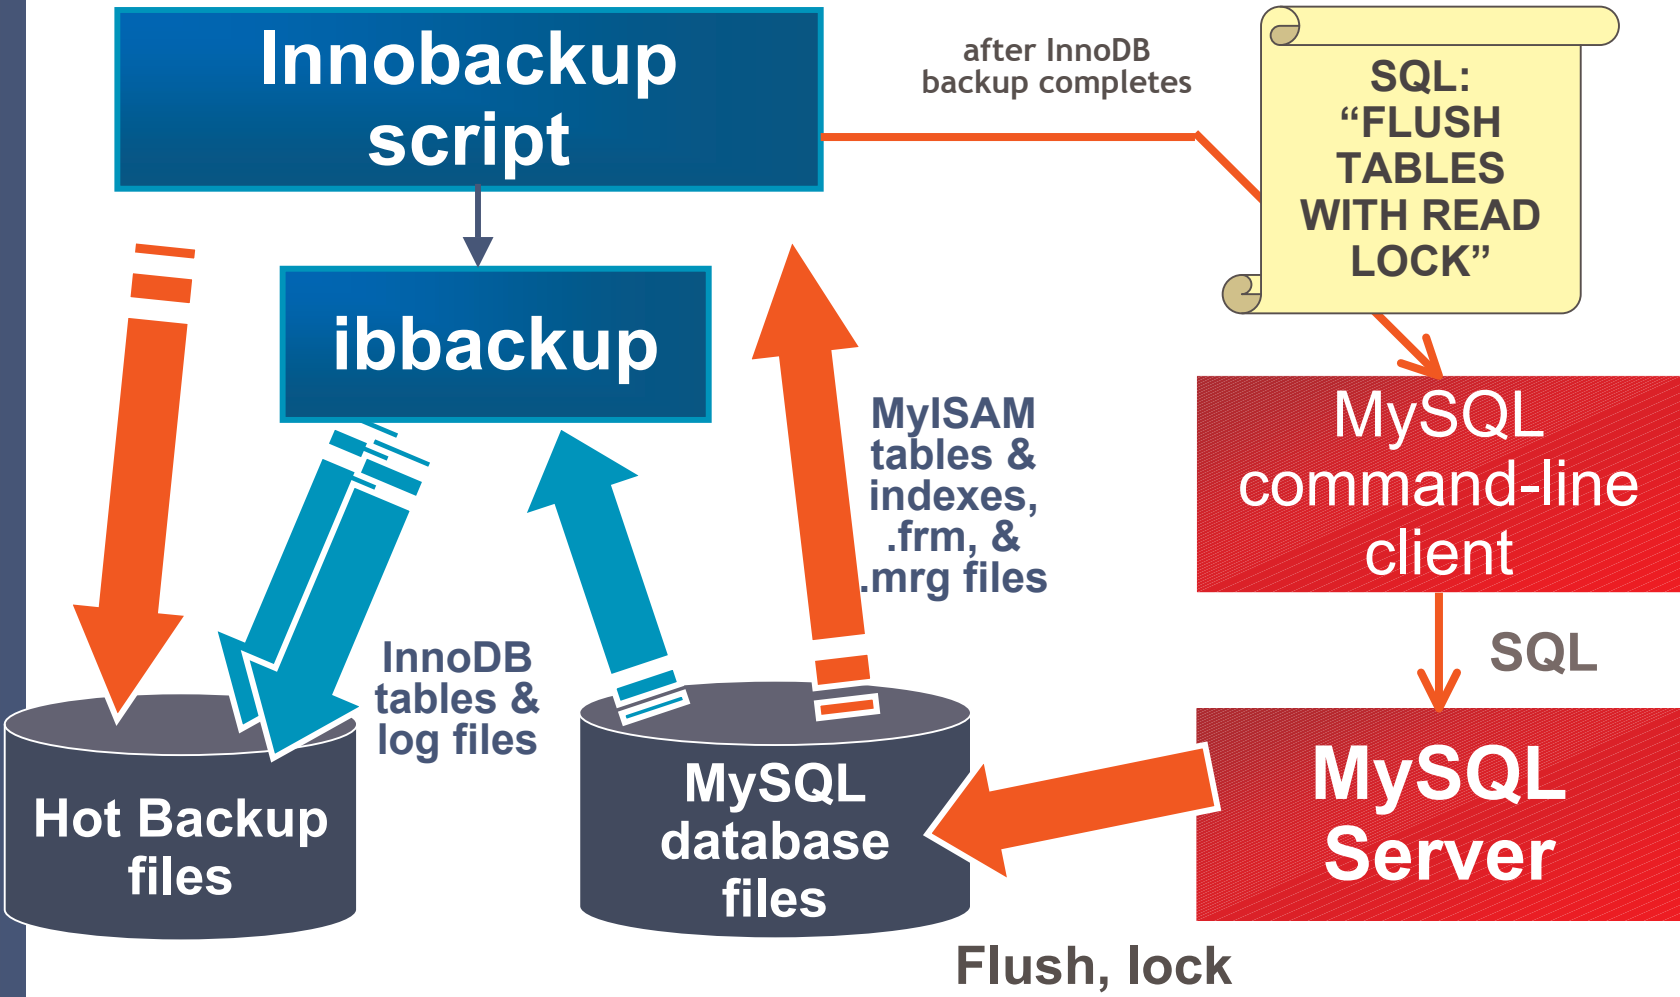


What is InnoDB Hot Backup?

- 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

MyISAM Backup with Innobackup



INNOBASE



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

Backup Tool Comparison

	InnoDB Hot Backup	MySQL dump	MySQL HotCopy	Zmanda RM for MySQL	Linux LVM	Veritas
Allows updates	Yes (InnoDB tables only)	No	No	No	Yes	Yes
Backs up .frm files	Yes	Yes	Yes	Yes	Yes	Yes
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



Summary

- InnoDB is the leading transactional storage engine for MySQL
- InnoDB delivers high performance and protects data integrity for MySQL applications
 - transactions, row-level locking, multi-version concurrency control, referential integrity and crash recovery
- The Innobase and MySQL relationship is strong, and the future of InnoDB is sound
- New features in InnoDB for MySQL V5.0 and V5.1 improve performance and convenience; more new developments underway!
- InnoDB Hot Backup permits on-line backup of running MySQL databases

INNOBASE



For More Information ...

2007 MySQL User Conference Sessions

- *InnoDB: Status, Architecture, and New Features*
- *InnoDB Birds of a Feather*
- *Clash of the DB Egos*
- *InnoDB: Efficient, Reliable Transactional Storage for MySQL*
- *InnoDB Performance Optimization*
 - Heikki Tuuri and Peter Zaitzev (Percona)
- *InnoDB Performance Potential in High-end Environments*
 - Yasufumi Kinoshita (NTT COMWARE CORPORATION)
- *The innotop Workshop*
 - Baron Schwartz, Rimm-Kaufmann Group

Please visit www.innodb.com

INNOBASE



INNOBASE

developer of

INNODB[®]

and

INNODB[®] Hot Backup

ORACLE[®]

Innbase is an Oracle company

