



# Developer and DBA Guide to What's New in MySQL 5.6

# Safe Harbor Statement DTCC2013

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decision. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

# Program Agenda

DTCC2013

- Oracle's Investment in MySQL
- DBA/Developer Guide to MySQL 5.6
- Development Priorities
- Questions?

# MySQL 5.6 is GA!

DTCC2013



- The Best MySQL Release Ever
- All New Features and enhancements Available to all in the MySQL Community Edition
- Designed to Meet Your Most Demanding Web, Cloud and Embedded Application Requirements

# World's Most Popular Open Source Database

DTCC2013

- Over 15 million estimated installations
- Used by 9 of top 10 web sites in the world
- Embedded by 8 of the top 10 ISVs
- #1 database in the Cloud
- Integrated w/Hadoop in 80% of Big Data platforms
- Facebook: 182K fans, +35% YoY Growth
- Twitter: 29K followers, +67% YoY Growth
- Numerous Awards: Linux Journal, CRN, PHP Architect...

# Oracle Significantly Invests in MySQL


DTCC2013

- Complete Solutions
- Best of Breed at Every Level
- On Premise and in the Cloud
- **MySQL: Web, Mobile & Embedded**



# Industry Leaders Rely on MySQL

DTCC2013



UNMATCHED INVESTMENT

DTCC2013  
InnoDB

PERFORMANCE IMPROVEMENTS

REPLICATION

EMBEDDED

CLOUD

WINDOWS

HUNDREDS OF EXPERTS

# ORACLE DRIVES MySQL INNOVATION

STRATEGIC

WORLD-CLASS SUPPORT

MySQL CLUSTER

WEB

NoSQL

LINUX

MySQL ENTERPRISE EDITION

LARGEST MySQL ENGINEERING & SUPPORT ORGANIZATION



# DRIVING MySQL INNOVATION

MySQL Enterprise Monitor 2.2  
MySQL Cluster 7.1  
MySQL Cluster Manager 1.0  
MySQL Workbench 5.2  
MySQL Database 5.5  
MySQL Enterprise Backup 3.5  
MySQL Enterprise Monitor 2.3  
MySQL Cluster Manager 1.1

**All GA!**

**2010**

MySQL Enterprise Backup 3.7  
Oracle VM Template for MySQL Enterprise Edition  
MySQL Enterprise Oracle Certifications  
MySQL Windows Installer  
MySQL Enterprise Security  
MySQL Enterprise Scalability

**All GA!**

MySQL Database 5.6 DMR\*  
MySQL Cluster 7.2 DMR

**MySQL Labs!**

**("early and often")**

**2011**

DTCC2013  
MySQL Cluster 7.2  
MySQL Cluster Manager 1.4  
MySQL Utilities 1.0.6  
MySQL Migration Wizard  
MySQL Enterprise Backup 3.8  
MySQL Enterprise Audit  
MySQL Database 5.6

**All GA!**

MySQL Cluster 7.3 DMR  
**Available Now!**

**A BETTER  
MySQL**

**2012-13**

**\*Development Milestone Release**

# Oracle's Investment in MySQL Community DTCC2013

**Available to download and use under the GPL**

- MySQL Database (Community Edition)
- MySQL Cluster
- MySQL Workbench
- MySQL Migration Wizard
- MySQL Utilities (in Python)
- MySQL Connectors
- MySQL Proxy
- Forums



[mysql.com/downloads/](https://mysql.com/downloads/)

# MySQL 5.6: Best Release Ever!

DTCC2013

## IMPROVED PERFORMANCE AND SCALABILITY

- Scales to 48 CPU Threads
- Up to 230% performance gain over MySQL 5.5

## IMPROVED INNODB

- Better transactional throughput and availability

## IMPROVED OPTIMIZER

- Better query exec times and diagnostics for query tuning and debugging

## IMPROVED REPLICATION

- Higher performance, availability and data integrity

## IMPROVED PERFORMANCE SCHEMA


- Better Instrumentation, User/Application level statistics and monitoring

## New! NoSQL ACCESS TO INNODB

- Fast, Key Value access with full ACID compliance, better developer agility

# MySQL 5.6: Scalability


DTCC2013



- Users can fully utilize latest generations of hardware and OS
- Scales as data volumes and users grow

# MySQL 5.6: Scalability


DTCC2013



- Users can fully utilize latest generations of hardware and OS
- Scales as data volumes and users grow

# MySQL 5.6 SysBench Benchmarks DTCC 2013


SysBench (Read Write): MySQL 5.6 vs. 5.5 (Linux)



**Up to 151% Performance Gain**

# MySQL 5.6 SysBench Benchmarks DTCC2013

SysBench (Read Only): MySQL 5.6 vs. 5.5 (Linux)



MySQL 5.6

MySQL 5.5

Oracle Linux 6  
Intel(R) Xeon(R) E7540 x86\_64  
MySQL leveraging:  
- 60 of 96 available CPU threads  
- 2 GHz, 512GB RAM

**Up to 234% Performance Gain**

# MySQL 5.6: InnoDB

DTCC2013

## BETTER RESOURCE UTILIZATION

- Removal of legacy bottlenecks
- Improved threading/concurrency
- Optimized for Read-Only Workloads
- SSD Optimizations

## BETTER AVAILABILITY, FASTER SCALING

- Online DDL Operations
- Transportable Tablespaces
- Dump, Restore/Warm Buffer Pool

## DEVELOPER AGILITY

- Full Text Search
- NoSQL, Key-value access to InnoDB



# MySQL 5.6: InnoDB

DTCC2013

## Better Performance, Scalability

- Several internal improvements (e.g. Split kernel mutex, efficient Buffer Pool Flushing, more...)
- Optimized for Read Only workloads
- Persistent Optimizer Statistics
  - Increased plan stability, accurate statistics
  - Better user control, automatic/manual
- SSD Optimizations
  - 4, 8k page sizes
  - .ibd files outside of MySQL data dir
  - separate tablespaces for undo log



# MySQL 5.6: InnoDB

DTCC2013

## Optimized for Read Only Workloads

- Ideal for highly concurrent, read intensive web apps
- Enables developer control of read only transactional overhead

```
SET autocommit = 1;  
SELECT c FROM sbtest WHERE id=N;
```

On by default

```
SET autocommit = 0;  
START TRANSACTION READ ONLY;  
SELECT c FROM sbtest WHERE id=N;  
COMMIT;
```

Developer controlled

Less complexity, more apps can be standardized on InnoDB

# MySQL 5.6: InnoDB

## Online DDL Operations

DTCC2013

- CREATE INDEX
  - DROP INDEX
  - Change AUTO\_INCREMENT value for a column
  - ADD/DROP FOREIGN KEY
  - Rename COLUMN
  - Change ROW FORMAT, KEY\_BLOCK\_SIZE for a table
  - Change COLUMN NULL, NOT\_NULL
  - Add, drop, reorder COLUMN
- Adds flexible schemas, online changes, no downtime
  - No need to consider NoSQL options

# MySQL 5.6: InnoDB

DTCC2013

## Dump and Restore/Warm Buffer Pool

- Shortens warm up times after restart (from hours to minutes)
- At shutdown/startup or manually at any time
- Small footprint on disk - stores tablespace and page IDs only

Dump current state of Buffer Pool at shutdown:

```
mysql> SET innodb_buffer_pool_dump_at_shutdown=ON;
```

Load Buffer Pool from dump at startup:

```
mysql> SET innodb_buffer_pool_load_at_startup=ON;
```

- Better Elasticity - Quickly spin up new instances to meet demand
- Great for Cloud, Hosted, SaaS, On-premise deployments

# MySQL 5.6: InnoDB

## Full Text Search

- FULLTEXT indexes on InnoDB tables
- Keys on text-based content
- Speeds up searches for words, phrases
- Fully transactional, fast look up
- Natural language/Boolean modes, proximity search, relevance ranking

DTCC2013

```
create table quotes
( id int unsigned
  auto_increment primary
  key
  , author varchar(64)
  , quote varchar(4000)
  , source varchar(64)
  , fulltext(quote)
) engine=innodb;
```

```
select author as "Apple" from quotes
  where match(quote) against ('apple' in natural language mode);
```

Less complexity, more apps can be standardized on InnoDB

# MySQL 5.6: InnoDB

DTCC2013

## Transportable Tablespaces

- Enables export/import of tables between running MySQL instances

Export:

```
CREATE TABLE t(c1 INT) engine=InnoDB;  
FLUSH TABLE t FOR EXPORT; -- quiesce the table and create the meta data file  
$innodb_data_home_dir/test/t.cfg  
UNLOCK TABLES;
```

Import:

```
CREATE TABLE t(c1 INT) engine=InnoDB; -- if it doesn't already exist  
ALTER TABLE t DISCARD TABLESPACE;  
-- The user must stop all updates on the tables, prior to the IMPORT  
ALTER TABLE t IMPORT TABLESPACE;
```

- Better Elasticity - Quickly spin up new instances to meet demand
- Great for Cloud, Hosted, SaaS, On-premise deployments

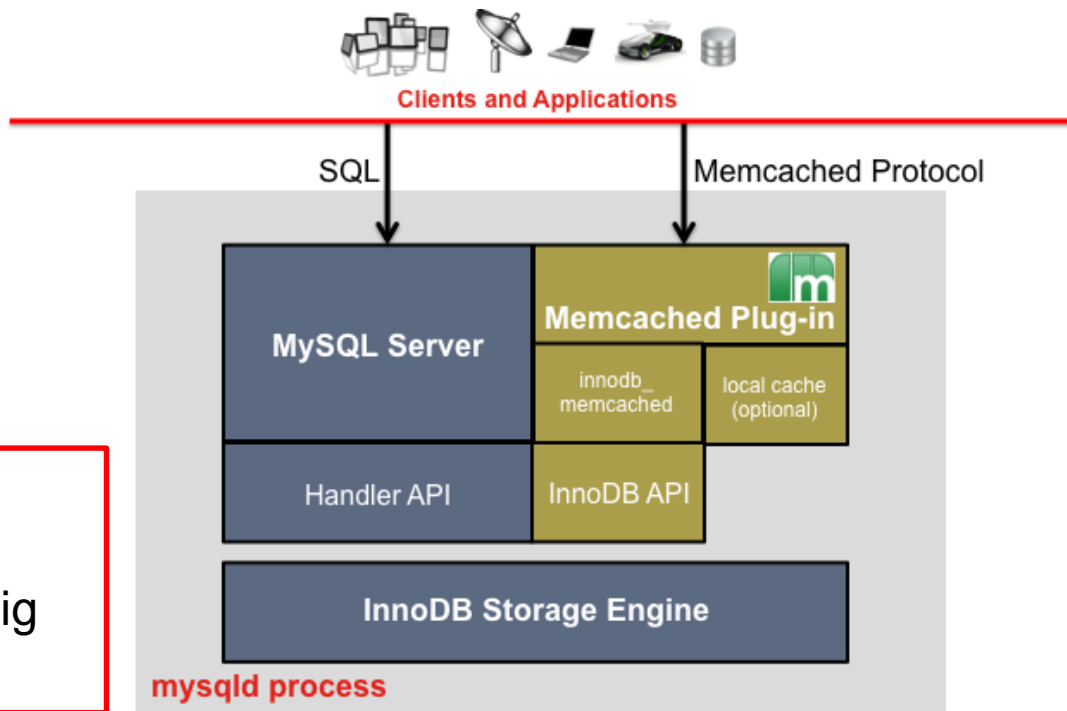
# MySQL 5.6: InnoDB

DTCC2013

## NoSQL Key Value Access to InnoDB

Same app can leverage:


- Key-value access to InnoDB via familiar Memcached API
  - SQL for rich queries, JOINS, FKs, etc.
  - Fully transactional
- Up to 9x performance boost for updates
  - Great for fast data ingestion in Big Data pipeline



# Performance

DTCC2013

## MySQL 5.6: NoSQL Benchmarking



**Up to 9x Higher “SET / INSERT” Throughput**

[blogs.oracle.com/mysqlinnodb/entry/new\\_enhancements\\_for\\_innodb\\_memcached](https://blogs.oracle.com/mysqlinnodb/entry/new_enhancements_for_innodb_memcached)



## BETTER QUERY PERFORMANCE, EXECUTION TIMES


- Subquery Optimizations
- File Sort Optimizations for most common web use cases
- Index Condition Pushdown
- Batched Key Access and Multi-Range Read

## BETTER DIAGNOSTICS

- EXPLAIN for INSERT, UPDATE, DELETE
- Structured EXPLAIN output (JSON)
- Optimizer Traces

# MySQL 5.6: Optimizer

DTCC2013



- Subquery Optimizations
- File sort optimizations for most web use cases
  - 4x better execution time – 40s to 10s
- Index Condition Pushdown
  - 160x better execution time – 15s to 90ms
- Batched Key Access and Multi Range Read
  - 280x better execution time – 2800s to 10s

- Better complex query execution times ever growing data sets (Big Data!)
- MEM + Query Analyzer key to utilizing full benefits of 5.6 Optimizer
- MySQL Consultative Support provides guidance on configuration

# MySQL 5.6: Optimizer

## Subquery Optimizations

DTCC2013

```
SELECT title FROM film WHERE film_id IN  
  (SELECT film_id FROM film_actor  
   GROUP BY film_id HAVING count(*) > 12);
```

- For Developers
  - No more re-writing legacy application or packaged subqueries into joins
- Performance
  - DBT 3 Query #13 benchmark:
  - Execution time drops from **DAYS** to **seconds**

## Postpone Materialization of views/subqueries in FROM

```
EXPLAIN SELECT * FROM (SELECT * FROM a_big_table);  
SELECT ... FROM derived_table AS dt  
          join table AS t WHERE dt.fld = t.dlf
```

- Late materialization
  - Allows fast EXPLAINS for views/subqueries
  - Avoid materialization when possible, faster bail out
  - A key can be generated for derived tables
- 240x better execution time (drops from ~8 min to ~2 sec)

## File Sort Optimizations with Small Limit

```
CREATE TABLE products(  
  productid int auto_increment PRIMARY KEY,  
  productname varchar(200)  
);
```

Web use case – list top 100 products sorted by name

```
SELECT * FROM products ORDER BY productname LIMIT 100;
```

- Use `sort_buffer_size` to sort all in memory, avoid creating intermediate sorted files
- Produce ordered result set using a single table scan
- Example above: 20 million rows, using default `sort_buffer_size`
  - 4x better execution time (drops from 40s to 10s)

# MySQL 5.6: Optimizer

DTCC2013

## Index Condition Pushdown (ICP)

```
CREATE TABLE person (  
    personid INTEGER PRIMARY KEY,  
    firstname CHAR(20),  
    lastname CHAR(20),  
    postalcode INTEGER,  
    age INTEGER,  
    address CHAR(50),  
    KEY k1 (postalcode,age)  
) ENGINE=InnoDB;
```


- With ICP Disabled
  - 15 s (buffer pool 128 Mb)
  - 1.4 s (buffer pool 1.5 Gb)
- With ICP Enabled
  - Execution time drops to 90 ms for both
- Controlled by `optimizer_switch` `index_condition_pushdown` flag

```
SELECT lastname, firstname FROM person  
WHERE postalcode BETWEEN 5000 AND 5500 AND age BETWEEN 21 AND 22;
```

# MySQL 5.6: Optimizer

DTCC2013

## Batched Key Access (BKA) and Multi-Range Read (MRR)



# Customer Requirements for Replication

DTCC2013



## PERFORMANCE

- Improve read consistency from slaves
- Reduce risk of data loss if master fails
- Minimize overhead

## FAILOVER & RECOVERY

- Self-Healing
- Improve availability

## DATA INTEGRITY

- Correct, consistent & accessible

## DEV/OPS AGILITY

- Enhance responsiveness to business
- Reduce TCO



# MySQL 5.6: Best Replication Features Ever

DTCC2013



## PERFORMANCE

- Multi-Threaded Slaves
- Binary Log Group Commit
- Optimized Row-Based Replication

## FAILOVER & RECOVERY

- Global Transaction Identifiers
- Replication Failover & Admin Utilities
- Crash Safe Slaves

## DATA INTEGRITY

- Replication Event Checksums

## DEV/OPS AGILITY

- Time Delayed Replication
- Remote Binlog Backup
- Informational Log Events


# MySQL 5.6: Replication

## Global Transaction Ids


DTCC2013

- Simple to track & compare replication across the cluster
  - Unique identifier for each transaction written to the Binlog
- Automatically identify the most up-to-date slave for failover
- Deploy n-tier replication hierarchies

Eliminates the need for *complex*  
3<sup>rd</sup> party solutions




# MySQL 5.6: Replication HA Utilities (Python) DTCC2013



- Enabling self-healing replication topologies
- Automated failover & recovery
  - `mysqlfailover` Utility
- Switchover & administration
  - `mysqlrpladmin` Utility



**Monitoring** ←  **HA Utilities**

- Delivers HA within the core MySQL distribution
  - Eliminates the need for 3<sup>rd</sup> party solutions
  - Allows extensibility to support variety of HA mechanisms


# MySQL 5.6: Replication

## Multi-Threaded Slaves

DTCC2013

- Increases slave throughput, reducing lag
- Applies events to different databases in parallel using concurrent SQL threads
- **5x performance gain**

Multi-Threaded Slave Performance



- SysBench, running across 10 x schemas
- Oracle Linux 6.1, Oracle Sun Fire x4150 m2 Server


- Great for systems which isolate application data using databases
  - Cloud, SaaS, Hosting, other multi-tenant deployments

# MySQL 5.6: Replication

## Binary Log Group Commit

DTCC2013

- Increases replication throughput by increasing performance of the master
- Commits multiple transactions as a group to Binlog on disk
- Finer grained locking; reducing lock wait times




- Better transactional throughput, less slave lag when coupled with MTS
- MySQL Consultative Support provides guidance on configuration

# Binary Log Group Commit Performance

DTCC 2013

**Binlog=1**

MySQL 5.6 vs. 5.5 - Read Write (Linux)



**180% Performance Gain**

# MySQL 5.6: Replication


## Optimized Row Base Replication

DTCC2013

Prior to MySQL 5.6 or `binlog-row-image = full`



With MySQL 5.6, `binlog-row-image = minimal`



- New option: `binlog-row-image=minimal`
- Increases throughput for master and slave
  - Reduces Binlog size, memory & network bandwidth
- Only replicates elements of the Row image that have changed

# MySQL 5.6: Performance Schema DTCC2013

## *New Instrumentation*

- Statements/Stages
- Table and Index I/O
- Table locks
- Users/Hosts/Accounts
- Network I/O

```
[mysqld]  
performance_schema=on
```

## *New Features*

- Show contents of Host cache
- New Summary tables
- Easier configuration
  - Start up defaults in my.cnf
  - Auto tune
- Reduced overhead
- On by default

Provides user/session level stats on resource usage for Cloud-based consumption/reporting/charge back



# MySQL 5.6: In Summary

DTCC2013

## IMPROVED PERFORMANCE AND SCALABILITY

- Scales to 48 CPU Threads
- Up to 230% performance gain over MySQL 5.5

## IMPROVED INNODB

- Better transactional throughput and availability

## IMPROVED OPTIMIZER

- Better query exec times and diagnostics for query tuning and debugging

## IMPROVED REPLICATION

- Higher performance, availability and data integrity

## IMPROVED PERFORMANCE SCHEMA

- Better Instrumentation, User/Application level statistics and monitoring

## New! NoSQL ACCESS TO INNODB

- Fast, Key Value access with full ACID compliance, better developer agility

# MySQL Database Development Priorities

DT002013

- Optimized for Web, Cloud-based, Embedded use cases
- Simplified, Pluggable architecture
  - Maintainability, more extensible
  - More NoSQL options (HTTP, JSON, JavaScript, etc.)
- Refactoring
  - Data Dictionary in InnoDB
  - Optimizer/Parser/Protocol
- InnoDB
  - Optimized for SSD
  - GIS
- Easy HA, Replication and Sharding

# Learn More

DTCC2013

- [mysql.com](http://mysql.com)
  - MySQL Products, Editions, Training, Consulting
  - TCO calculator
  - Customer use cases and success stories
- [dev.mysql.com](http://dev.mysql.com)
  - Downloads, Documentation
  - Forums
  - PlanetMySQL
- [eDelivery.oracle.com](http://eDelivery.oracle.com)
  - Download and evaluate all MySQL products



# New MySQL 5.6 Training

Learn about the world's most popular open-source database



DTCC2013

## Learn MySQL From Oracle

- Expert-led training to help you install, configure, and administer MySQL 5.6.
- Extensive hands-on practices guide you through each concept
- Explore real-world problems and discover best practices as you work with the tools and techniques used by professional MySQL database administrators
- Content developed in collaboration with product engineering.
- Available in traditional or virtual classroom as well as self-study formats.
- Custom training solutions to match your organization's specific business needs
- Backed by Oracle University's **100% Satisfaction Program**

## MySQL for Database Administrators **NEW**

Schedule/Purchase	Training Formats 	Price	Duration	Course Materials	Language
<a href="#">▶ View Schedule</a>	Classroom Training	US\$ 2,250	5 Days	English	
<a href="#">▶ View Schedule</a>	Live Virtual Class	US\$ 2,250	5 Days	English	Multiple

### What You Will Learn

### Audience

### Course Topics

### Course Objectives

### Related Training

The MySQL for Database Administrators course is for DBAs and other database professionals who want to install and configure the MySQL Server, set up replication and security, perform database backups and performance tuning, and protect MySQL databases.

#### Learn To:

- ▶ Describe the MySQL Architecture.
- ▶ Install and upgrade MySQL.
- ▶ Configure MySQL server options at runtime.
- ▶ Evaluate data types and character sets for performance issues.
- ▶ Understand the use of the InnoDB storage engine with MySQL.
- ▶ Perform backup and restore operations.

#### Setting up the Server and Databases

This course teaches you how to install, configure, maintain and tune databases for a wide variety of workloads and environments. You will install and configure MySQL and set appropriate configuration options. You will also connect to the server, using both graphical and command-line clients, while creating and populating databases.

#### Maintaining and Securing Data

Expert Oracle instructors will also help you examine the structure of tables, their columns and data types and the different ways to check and repair tables. You will learn to program MySQL stored routines, events and triggers to process data in complex and automated ways. Explore how to partition tables, work with locking and transactions as well as the key features of InnoDB and other storage engines. You will develop an understanding of how data is processed, stored, and retrieved.

[oracle.com/education/mysql](http://oracle.com/education/mysql)

# 网上讲座直播— MySQL

DTCC2013

通过 MySQL 获得更好的成效，就从您的计算机开始！

获取世界上最流行的开源数据库的最新信息

本次免费网上讲座由现场演示和互动问答两部分组成。

主题：MySQL网上讲座

- 日期：2013 年 4 月 23 日
- 时间：10:00AM -11:30AM



**立即注册！** <http://valuemarketing.com.cn/oracle/20130403/oracle.html>



# Developer and DBA Guide to What's New in MySQL 5.6

Questions?