



Raising The MySQL Bar

日本オラクル株式会社


山崎 由章 / MySQL Senior Sales Consultant,
Asia Pacific and Japan

April 2014

以下の事項は、弊社の一般的な製品の方向性に関する概要を説明するものです。また、情報提供を唯一の目的とするものであり、いかなる契約にも組み込むことはできません。以下の事項は、マテリアルやコード、機能を提供することをコミットメント(確約)するものではないため、購買決定を行う際の判断材料になさらないで下さい。オラクル製品に関して記載されている機能の開発、リリースおよび時期については、弊社の裁量により決定されます。

OracleとJavaは、Oracle Corporation 及びその子会社、関連会社の米国及びその他の国における登録商標です。文中の社名、商品名等は各社の商標または登録商標である場合があります。

BEST IN CLASS COMPONENTS



MySQL: Next Generation Web Applications
On-Premises, in the Cloud, Distributed Applications

ORACLE

Driving MySQL Innovation: 2010 - 2014

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

MySQL Enterprise Backup 3.7

All GA!

Oracle VM Template for MySQL

Oracle Products Certifications

MySQL Windows Installer

MySQL Enterprise Security

MySQL Enterprise Scalability

MySQL Cluster 7.2

MySQL Cluster Manager 1.3

MySQL Utilities 1.0.6

MySQL Workbench 6.0

All GA!

MySQL Enterprise Backup 3.10

MySQL Enterprise Audit

MySQL Windows Tools

MySQL Database 5.6

MySQL Cluster 7.3

MySQL Workbench 6.1

All GA!

MySQL Fabric 1.4 RC

MySQL Database 5.7 DMR

Available Now!

ORACLE

MySQL 5.6: GA 14 Months Ago – Best GA Ever

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

- Faster query execution 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

Thank You, MySQL Community

Helping to Create a Better MySQL

MySQL 5.6.17 Community Release Notes

[Leave a reply](#)

Thank you to the MySQL Community, on behalf of the MySQL team @ Oracle. Your bug reports, testcases and patches have helped create a better [MySQL 5.6.17](#).

Special Thank you to 徹 赤松, Yoshinori Matsunobu!

Thank You, MySQL Community

Helping to Create a Better MySQL

- Thanks to 徹 赤松 for reporting that SHOW SLAVE STATUS incorrectly reported the master's SSL information. [Bug #70866](#).
- Thanks to Yoshinori Matsunobu for reporting an issue where a slave can't continue replication after the master's crash recovery [Bug #70669](#)
- Thanks to Yoshinori Matsunobu for reporting excessive mutex contention in semi-sync replication. [Bug #70218](#).

Special Thank you to 徹 赤松, Yoshinori Matsunobu!

Welcome, WebScaleSQL!

Thank You, Facebook, Google, LinkedIn & Twitter!

- Investment to improve MySQL for Web Scale
- Great news for the MySQL Community & new developers
- Even easier for Oracle to continue close cooperation

Raising the Bar: **MySQL 5.7**

5.6 GA
Feb 2013

5.7 DMR 1
April 2013

DMR 2
Sep 2013

DMR 3
Dec 2013

...
DMR 4
New features in
labs.mysql.com
Available Now!

MySQL Repositories: Yum

Benefits both End Users and Linux Distributions

- Simple and convenient way to install & update MySQL products
- Supports the following distributions
 - RedHat Enterprise/Oracle Linux
 - Fedora
 - Upcoming: Debian/Ubuntu
- Includes the latest packages
 - MySQL Database
 - MySQL Workbench
 - MySQL Connector/ ODBC
 - MySQL Connector/Python
 - MySQL Utilities

Over 50% of downloads
for Fedora/RedHat already
from the repositories

And now also on
NuGet for Windows Devs
<http://www.nuget.org/profiles/MySQL/>

MySQL 5.7: DMR 4

MySQL 5.7 builds on MySQL 5.6 by improving:

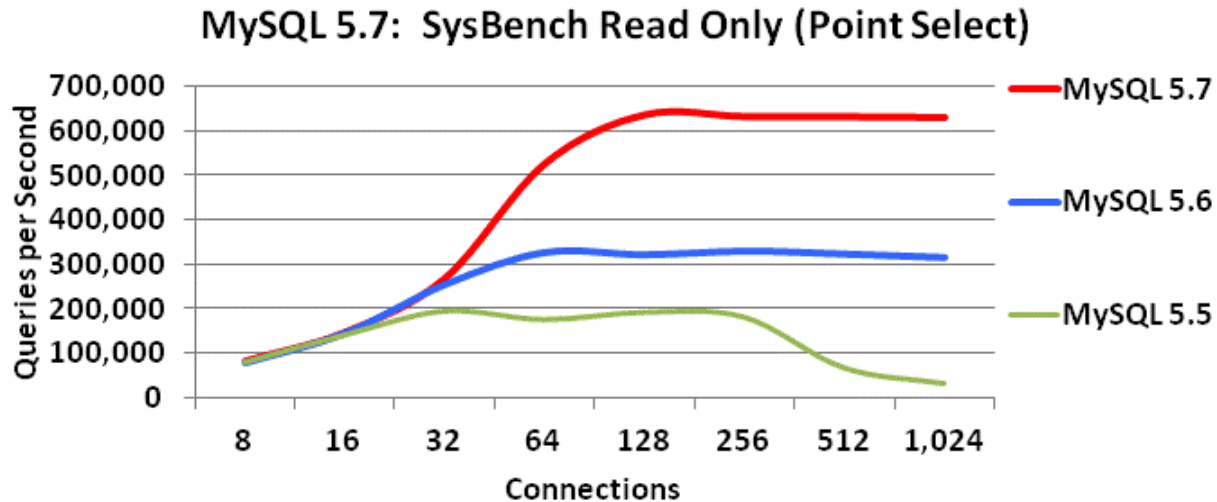
- **InnoDB** for better transactional throughput, availability, IO
- **Replication** for better scalability and availability
- **Utilities** for dev/ops automation
- **Performance Schema** for better performance metrics
- **Optimizer** for better EXPLAINing, query performance, enhanced buffering and partition optimization
- **Connecting** at higher rates, improve session efficiency

Available Now! Get it here: dev.mysql.com/downloads/mysql/

MySQL 5.7 Sysbench Benchmark

Sysbench Point Select

630,000 QPS



Intel(R) Xeon(R) CPU X7560 x86_64
5 sockets x 8 cores-HT (80 CPU threads)
2.27GHz, 256G RAM
Oracle Linux 6.5

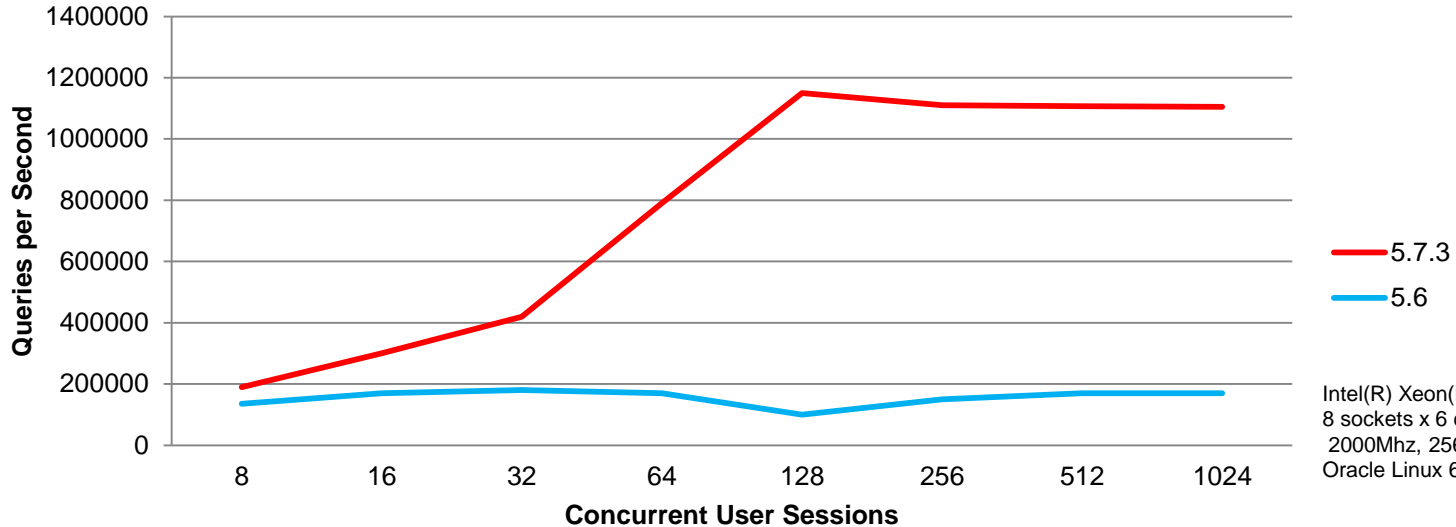
2X Faster than MySQL 5.6
Over 3X Faster than MySQL 5.5

ORACLE

MySQL 5.7: InnoDB Memcached

Thank you Yoshinori (Facebook)

1,150,000 QPS



Intel(R) Xeon(R) CPU X7560 x86_64
8 sockets x 6 cores-HT (96 CPU threads)
2000Mhz, 256G RAM
Oracle Linux 6.2

6x Faster than MySQL 5.6

MySQL 5.7: Optimizer

EXPLAIN on Running Queries

- Problem: a statement in a session is taking a long time to complete
- New Option: run
 - EXPLAIN FOR CONNECTION from another session

```
EXPLAIN [FORMAT=(JSON|TRADITIONAL)] FOR CONNECTION <id>;
```

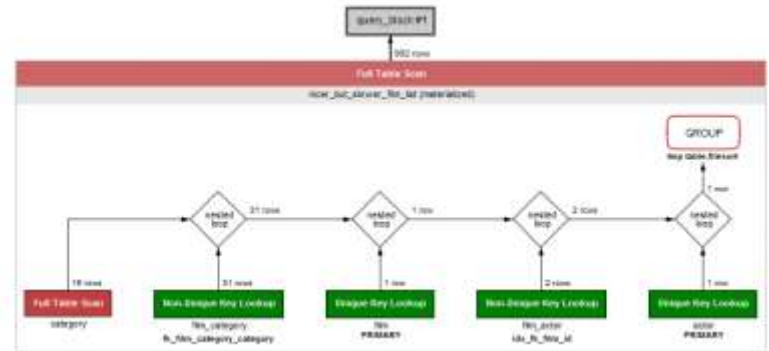
- Discover the root cause of the delay
- Optimize the statement

MySQL 5.7: Optimizer

```
"query_block": {
  "select_id": 1,
  "cost_info": {
    "query_cost": "1.00"
  },
  "table": {
    "table_name": "actor",
    "access_type": "const",
    "possible_keys": [
      "PRIMARY"
    ],
    "key": "PRIMARY",
    "used_key_parts": [
      "actor_id"
    ],
    "key_length": "2",
    "ref": [
      "const"
    ],
  },
  "rows_examined_per_scan": 1,
  "rows_produced_per_join": 1,
  "filtered": 100,
  "cost_info": {
    "read_cost": "0.00",
    "eval_cost": "0.20",
    "prefix_cost": "0.00",
    "data_read_per_join": "280"
  },
  "used_columns": [
    "actor_id",
    "first_name",
    "last_name",
    "last_update"
  ]
}
```

■ Expanded JSON EXPLAIN

- Now includes all available cost info
- Used for Visual Explain In MySQL Workbench 6.1



MySQL 5.7: Optimizer Cost Model

- New cost model API
- Allows storage engines to provide accurate and dynamic cost estimates for key lookups, table scans, range scans, etc...
 - Enables future support of additional factors
 - Whether the data is in RAM, SSD, HDD
- Lays the groundwork for making costs configurable
 - Based on your hardware performance characteristics
- Improves records per key estimates
- Cost values included in JSON Explain output

MySQL 5.7: Parser Refactoring

- Improves readability, maintainability, and stability
 - Cleanly separate the parsing, optimizing, and execution stages
 - Allows for easier feature additions, with lessened risk
- Transforms SQL Parser into a true bottom-up parser
 - Better support for larger syntax rules in the future
 - Fewer intermediate nodes in the parse tree
 - Fewer node allocations
 - Faster grammar

MySQL 5.7: InnoDB

Many Improvements

- Improved InnoDB Online Alter Table
 - Online Rename Index, Online Change Varchar
- Enhanced FusionIO Integration
 - Doublewrite buffer automatically disabled when DirectFS detected
- Parallel “Dirty Page” Flushing
 - Higher throughput, performance, and scalability
- Partitions – support for Transportable Tablespaces (TTS)
 - TTS support for individual partitions

MySQL 5.7: InnoDB Compression

labs.mysql.com

Thank you, Fusion-io

- Transparent Page Level Compression
 - Happens transparently in background threads
 - Managed entirely within the IO layer
 - Uses sparse file and "hole punching" in OS kernels and File Systems
- Reduces IO
 - Improves performance
 - Reduces write cycles, thus increasing SSD lifespan
- Applies to all tables, including the system tablespace and UNDO logs

MySQL 5.7: InnoDB Temp Tables

- New Separate tablespace for temporary tables
 - Improved CREATE/DROP performance
- DDL changes are transient
 - Eliminates some disk IO
- Optimize DML operations
 - No REDO logging, no change buffering, less locking

MySQL 5.7: Server-side Statement Timeouts

Thank you Davi Arnaut

- Server Side statement timeouts
 - Global for server, per session, or for individual SELECT statements
- ```
SELECT MAX_STATEMENT_TIME = 109 * FROM my_table;
```
- Based on contribution from Davi Arnaut
    - Expanded to Windows and Solaris, restricted by removing USER option

# MySQL 5.7: Improved MDL locking

- Removes bottlenecks around DML access to a single table
  - 10% increased throughput in OLTP\_RO/POINT\_SELECT sysbench tests on higher core counts
  - Optimized for typical DML heavy workloads
- Implemented fast-path for DML locks
- Implemented lock-free DML lock acquisition
- Implemented a lock-free hash
  - Now uses MurmurHash library

# MySQL 5.7: Security

- AES 256 Encryption
  - Default in MySQL 5.7
- Password Rotation Policies
  - Can be set globally, and at the user level
- Deployment: Enable Default Secure Unattended Install
  - Random password set on install
  - Remove anonymous accounts
  - Deployment without test account, schema, demo files

- InnoDB Spatial Index Support
  - Optimized R-tree indexes
- Boost.Geometry integration
  - Improved accuracy
  - Support for additional types, combinations, and functions
  - Improved performance
  - Full Open Geospatial Consortium compliance
  - Enhancements to Boost.Geometry pushed upstream



# MySQL 5.7: Performance Schema



## Memory Usage

- Instruments for Memory
  - Added for over 200 Memory Types
- Aggregates memory usage statistics by
  - Type of memory used (caches, internal buffers, ...)
  - Thread/account/user/host indirectly performing the memory operation
- Attributes include
  - Memory used (bytes), Operation counts, High/Low Water Marks

## Statement Instrumentation

- Extended the statement instrumentation
  - Stored Procedures and Stored Functions
  - Prepared Statements
  - Transactions

## Additional Data

- Additional instrumentation and metrics
  - Replication slave status
  - MDL lock instrumentation

# MySQL SYS Schema

## Simplified Views on Performance Schema & Information Schema

- Helps simplify DBA tasks
  - Monitor health, growth rates, and other key metrics
  - Spot, diagnose, and tune performance problems
- Provides easy to understand insights into
  - IO hot spots
  - Costly SQL statements
  - Dynamic table, index, and schema statistics
  - Wait time analysis
  - Locking
  - InnoDB statistics



# MySQL SYS Schema

- Driven from positive feedback on ps\_helper
  - Over 80 views, versioned, self updating, and server version aware
  - Supporting MySQL 5.5, 5.6, and 5.7
- SYS is similar to:
  - Oracle V\$ catalog views
  - Microsoft SQL Server DMVs (Dynamic Management Views)
  - IBM DB2 SYSIBM catalog
- Available with Workbench 6.1 or via GitHub
  - Workbench includes integrated “one click” reports

# MySQL Workbench 6.1

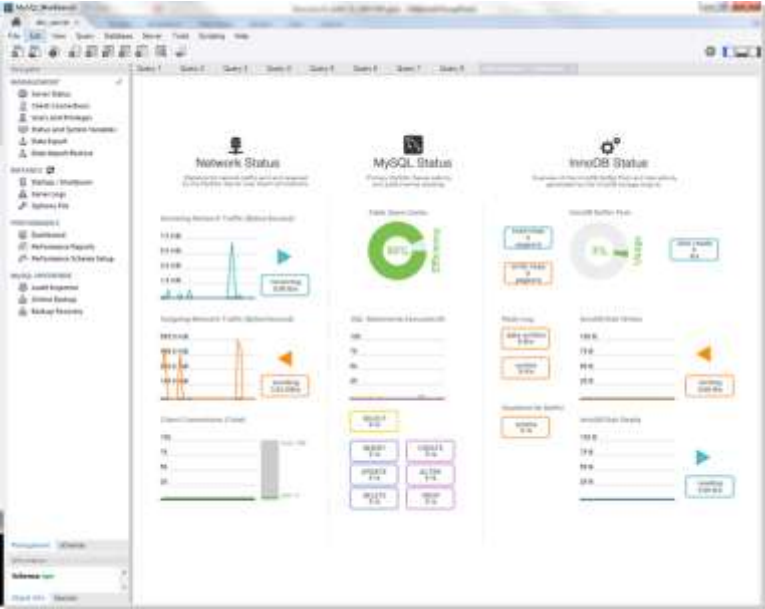


- Performance Dashboard
  - View real-time server performance stats on a graphical dashboard
  - Performance Schema Reports based on new "SYS" schema
- Visual Explain 2.0
  - Much improved visual explain output, with simpler and clearer layout
- Enhanced Enterprise Backup Support
  - Automatic setup of requirements and support for existing setups
  - Enhanced support for management of new and existing backup accounts



# MySQL Workbench 6.1

## Performance and Status Dashboards



Network, Server, InnoDB

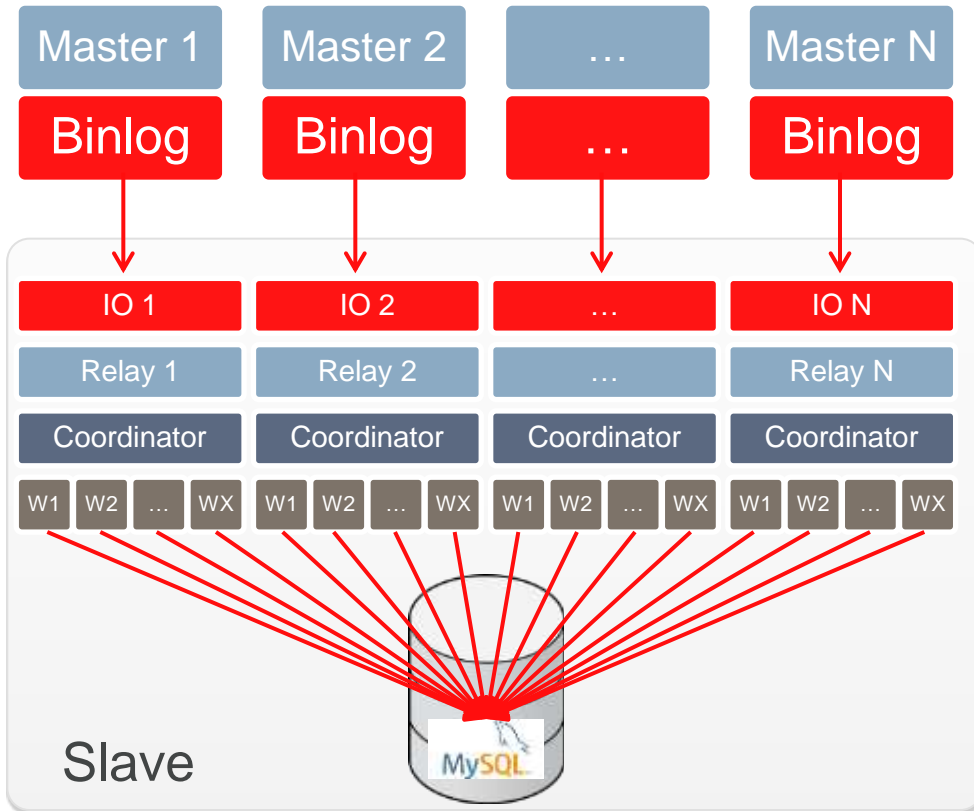
Analyze hotspots, costly SQL statements, wait times, locks, InnoDB stats, and more

# MySQL 5.7.4: Replication

## Better Performance, Improved Usability and Enhanced HA

- Higher throughput
  - Slave : applies transactions in parallel even within same database
  - Master: better synchronization between replication-user sessions
- DMR4 improves master performance for Semi-synchronous Replication
- Lossless Replication through enhanced Semi-sync
- Performance Schema tables for monitoring slave
- Dynamic Replication Filters: change filters on the fly

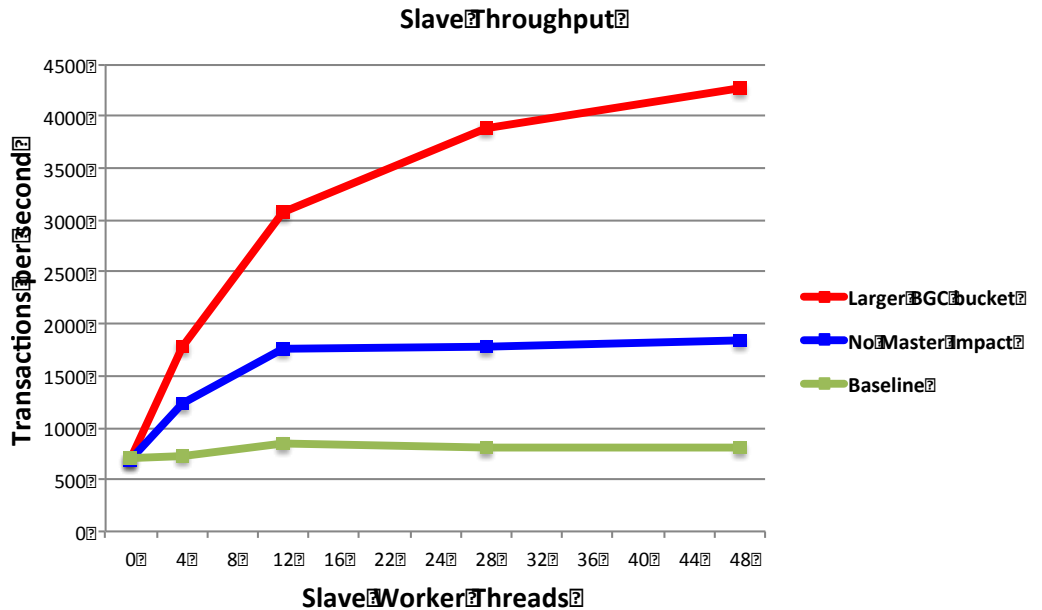
# Multi-Source Replication



- Consolidate updates from multiple Masters into one Slave
  - Consolidated view of all shards
  - More flexible topologies
  - Centralized point for backups, easier reporting
- Compatible with Semi-Synchronous Replication & enhanced MTS
- Master-specific slave filters planned for GA
- Application must keep data sets disjoint between sources

# MySQL Multi-Threaded Slave

Slave keeps pace with the master



- The evolution...
  - 2010: Inter-schema MTS (labs)
  - 2013: Inter-schema MTS (5.6 GA)
  - 2013: **Intra**-schema MTS (5.7.2 DMR)
  - Today: Tune size of group commit buckets - trading commit latency on master (labs)
- 6X slave throughput




# MySQL Utilities 1.4 RC

## Powerful DevOps Management tools for MySQL

- Automate common Dev/Ops tasks
  - Replication: provisioning, testing, monitoring and failover
  - Database comparisons: consistency checking
  - Database administration: users, connections, tables
  - Auditing
- Python scripts
  - Now standalone or launched from MySQL Workbench
  - Extensible to include custom scripting; Python library for extensibility
- New: Round-Robin Multi-Source Replication & slave synchronization check

# MySQL Fabric 1.4 RC (MySQL Utilities 1.4.2)

## High Availability + Sharding-Based Scale-out





- High Availability:
  - Server monitoring with auto-promotion and transparent application failover
- Fabric-aware connectors
  - Python, Java, PHP
- No Proxy -> no extra latency
- Optionally scale-out through sharding
  - Application provides shard key
  - Range or Hash
  - Tools for re-sharding
  - Global updates & tables

# MySQL Cluster 7.4

labs.mysql.com

Better performance and operational simplicity



- Performance gain over 7.3
  - 47% (Read-Only)
  - 38% (Read-Write)

- Faster node restarts
  - Recovering nodes rejoin the cluster faster

ORACLE

# Oracle MySQL HA & Scaling Solutions

|                          | MySQL Replication | MySQL Fabric    | Oracle VM Template | Solaris Cluster | Windows Cluster | DRBD           | MySQL Cluster |
|--------------------------|-------------------|-----------------|--------------------|-----------------|-----------------|----------------|---------------|
| Clustering Mode          | Master + Slaves   | Master + Slaves | Active/Passive     | Active/Passive  | Active/Passive  | Active/Passive | Multi-Master  |
| App Auto-Failover        | ✗                 | ✓               | ✓                  | ✓               | ✓               | ✓              | ✓             |
| Data Layer Auto-Failover | ✗                 | ✓               | ✓                  | ✓               | ✓               | ✓              | ✓             |
| Zero Data Loss           | MySQL 5.7         | ✓               | ✓                  | ✓               | ✓               | ✓              | ✓             |
| Failover Time            | N/A               | Secs            | Secs +             | Secs +          | Secs +          | Secs +         | < 1 Sec       |
| Scale-out                | Reads             | ✓               | ✗                  | ✗               | ✗               | ✗              | ✓             |
| Cross-shard operations   | N/A               | ✗               | N/A                | N/A             | N/A             | N/A            | ✓             |
| Transparent routing      | ✗                 | For HA          | ✓                  | ✓               | ✓               | ✓              | ✓             |
| Shared Nothing           | ✓                 | ✓               | ✗                  | ✗               | ✗               | ✓              | ✓             |
| Storage Engine           | InnoDB            | InnoDB          | InnoDB             | InnoDB          | InnoDB          | InnoDB         | NDB           |
| Single Vendor Support    | ✓                 | ✓               | ✓                  | ✓               | ✗               | ✓              | ✓             |
| Platform Support         | All               | All             | Linux              | Solaris         | Windows         | Linux          | All           |

# MySQL Enterprise Backup 3.10

Speed, Compression, Flexibility, Efficiency, Security



- New compression options
- New Advanced Encryption Standard (AES) Security
  - Compress and/or secure in 1 streamlined and fast/efficient step
- Zero space / single step / streaming backup & recovery
  - In 3.9 full and incremental
  - Enhanced in 3.10 for tables
- New filtering options for partial backups



ORACLE

# MySQL Innovations for You

- Oracle Delivers & Raises the Bar. Past 12 Months:
  - **GA:** MySQL Cluster 7.3, MySQL Workbench 6.0 & 6.1, MySQL Enterprise Monitor 3.0, MySQL Utilities, MySQL Connectors...and more
  - **DMRs & Labs:** MySQL 5.7, MySQL Fabric, Multi-source Replication, Hadoop Applier for MySQL, GIS, MySQL Cluster 7.4...
- Oracle Makes MySQL Better for Next Generation Web, Cloud and Big Data Applications
- Try the Products and Give us Feedback!

ORACLE®