



MySQL Experience Sharing

Mickey Fan
Global System Engineering
Sun Microsystems of California Ltd



Market Opportunity

- \$15 billion RDBMS market (Gartner – 2006, published June 2007)
- Open Source DBMS – 40% growth in next 5 years – passing \$1 billion in revenue by 2012 (Gartner, 10 Jan, 2008)
- 66% of organizations using Open-Source software are using or plan to use open source DBMS in the next 12 months - (Gartner, 10 Jan, 2008)

http://www.sun.com/aboutsun/media/presskits/2008-0116/mysql_factsataglance.pdf




Figure 3: Relevance of Database Systems

http://www.joinvision.com/jv/ext/infow/itfacts/200710/itfacts200710_en.pdf

About MySQL



- Founded in 1995
- 400 employees
- 750 partners
- 50,000 downloads/day
- Customers across every major operating system, hardware vendor, geography, industry, and application type

Leadership, innovation, market acceptance

MySQL — An Open-Source Icon

- MySQL is...
 - > The “M” in LAMP (Linux, Apache, MySQL, and PHP)
 - > What is SAMP?
 - > (Solaris, Apache, MySQL, and PHP)
 - > The industry's most popular and one of the fastest-growing open- source databases
 - > Ubiquitous across Windows, Linux, Solaris™, and Mac OSes
 - > Pervasive across SUN, HP, Dell, IBM, and Fujitsu
 - > Deployed by the most important companies on Earth!



World's most popular open-source database

Multiplatform Offerings



- MySQL will continue to be developed and optimized for Solaris, Linux, Windows, etc.
- Solaris will continue to be optimized for MySQL, Oracle, Sybase, DB2 and PostgreSQL

MySQL — growing with Sun!

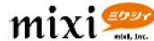
MySQL's Industry-Leading Customers







WIKIPEDIA



Web / Web2.0









Adobe



OEM / ISV's








On Demand, SaaS, Hosting






Telecommunications








Enterprise 2.0


Open-source is powering the Web

Coexistence and Disruption Strategy

- Large percentage of Oracle shops run MySQL
- 52% plan to increase use of open source
- 67% report cost savings as key factor
- 28% report reduced vendor lock-in

FIGURE 15: Open Source or “Express Edition” Databases

(Among Respondents Using Open Source/Express Databases)




MySQL Priorities

Be the Best Online Database




SPEC Benchmarks

- http://www.spec.org/osg/jAppServer2004/results/res2007q2/jAppServer200402_0001.html

 spec	SPECjAppServer®2004 Result Copyright © 2004 Standard Performance Evaluation Corporation		
Sun Java (TM) Systems Application Server 9.0 Platform Edition on SunFire X4100 Cluster with MySQL 5		720.56 SPECjAppServer2004 JOPS@Standard	
Submitter: Sun Microsystems Inc.		SPEC license # 6	Test date: Apr-2007
Software Products Sun Java (TM) Systems Application Server 9.0 Platform Edition Update Release 1 Patch 3 Jakarta Tomcat 5.5.16 Java (TM) 2 Platform Standard Edition Development Kit 5.0 Update 6 32-bit Connector/J 5.0.5 MySQL 5.0.27 64-bit		Software Configurations J2EE Application Servers Emulator SW Config Database SW Config Driver SW Config Primary	Hardware Systems J2EE AppServer HW Database Server HW Load Driver HW System Configuration Diagram
Benchmark Modifications Configuration Bill of Materials Other Info General Notes Full Disclosure Archive			


MySQL Cluster DBT2 Benchmarks

- HW
 - Dual Core Intel Xeon 5150 2.67Ghz
 - 4GB RAM
 - Single Gigabit Ethernet connection
 - Dolphin Express D352
- 8 Data Nodes on 2 Servers
- 16 MySQL Servers on 4 Servers




High Performance

Oracle9i and MySQL top throughput




Throughput is in returned Web pages per second from the application server. Number of users is number of concurrent Web clients driving the load. Response time is the time to complete the six bookstore user action sequences, weighted by frequency of each sequence in the mix. All tests were conducted on an HP NetServer LT 6000r with four 700MHz Xeon CPUs, 2 GB of RAM, a Gigabit Ethernet Intel Corp. Pro/1000 F Server Adapter and 24 9.1GB Ultra3 SCSI hard drives used for database storage.

MySQL OLTP Benchmark



More information on the benchmark results can be found at: <http://www.sun.com/x64/docs/MySQL-sysbench-benchmark.pdf>

MySQL OLTP Benchmark



More information on the benchmark results can be found at: <http://www.sun.com/x64/docs/MySQL-sysbench-benchmark.pdf>

CoolStack & WebStack



memcached

- Open source applications optimized and integrated for the robust Solaris OS
- Binaries that help deliver the best system performance
 - Up to 200% higher performance
- SMF & DTrace Integration
- Reduces time to service
- Free to download, Developer Support

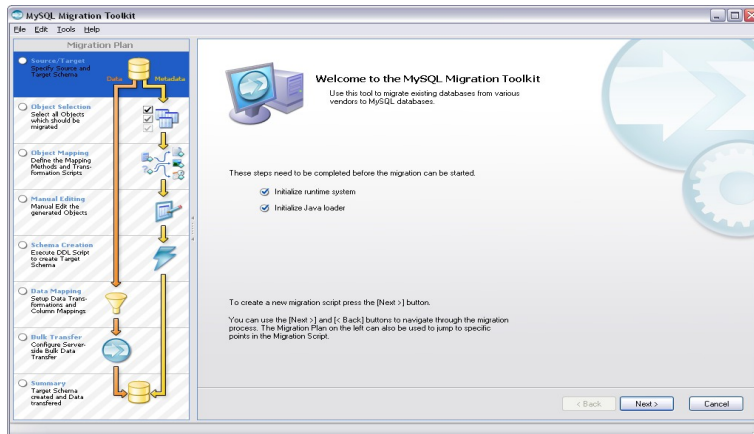
Already Integrated & Optimized For You

MySQL Priorities

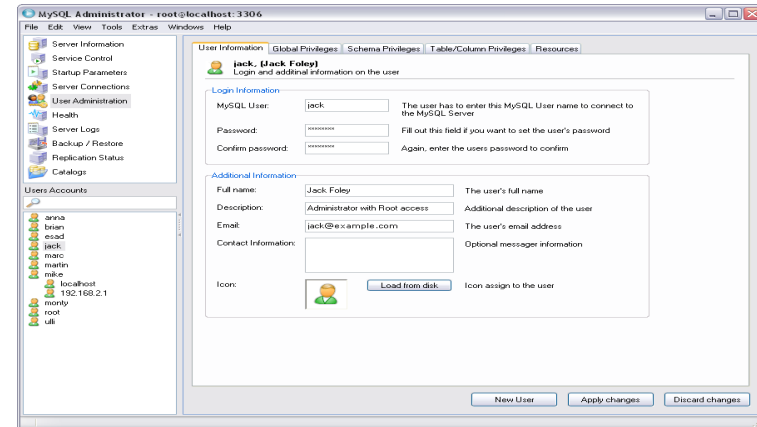
Be the Best Online Database



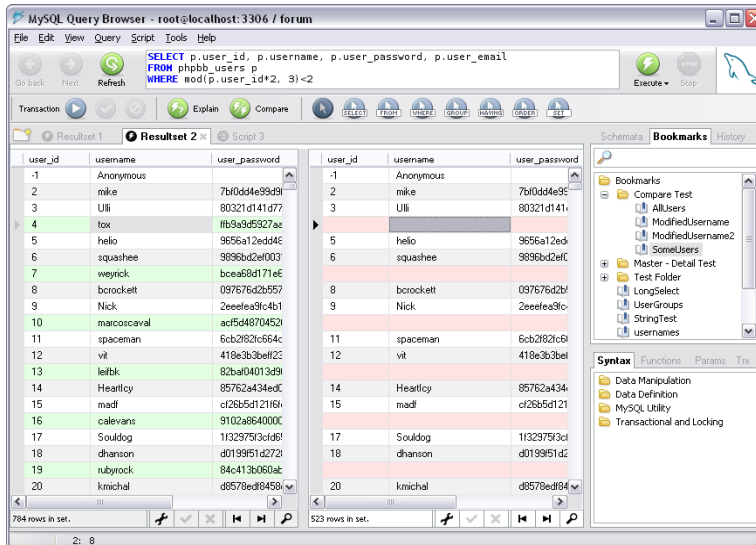
Tools from mySQL



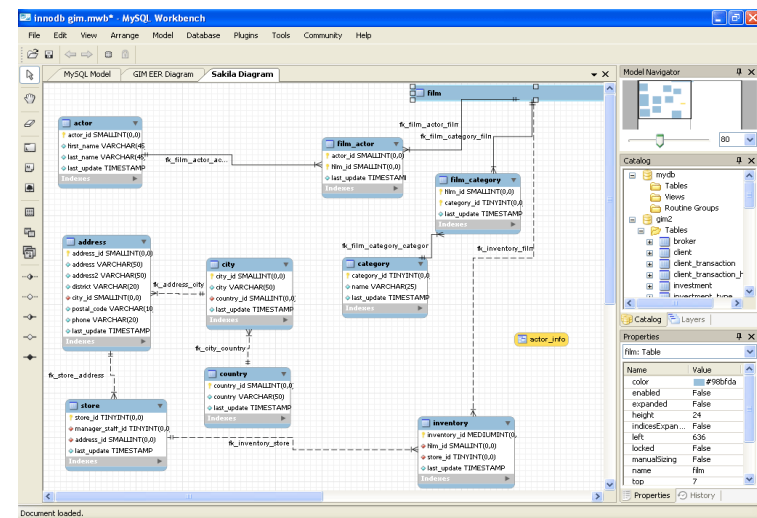
MySQL Migration Toolkit



MySQL Administrator



MySQL Query Browser



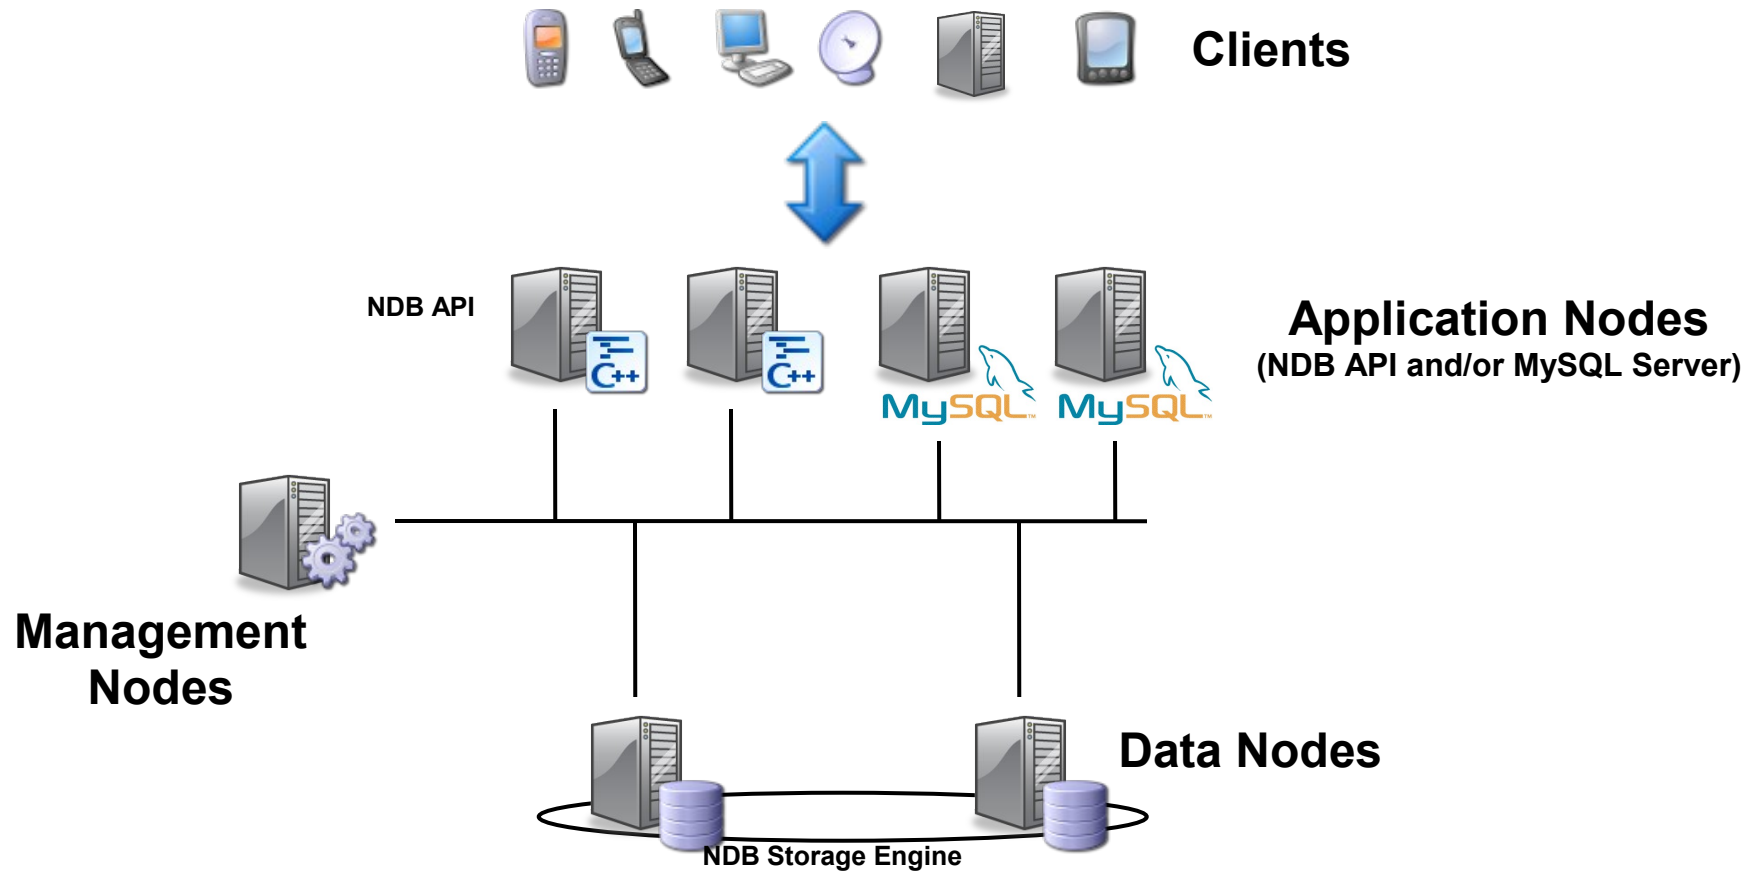
MySQL Workbench

MySQL Priorities

Be the Best Online Database

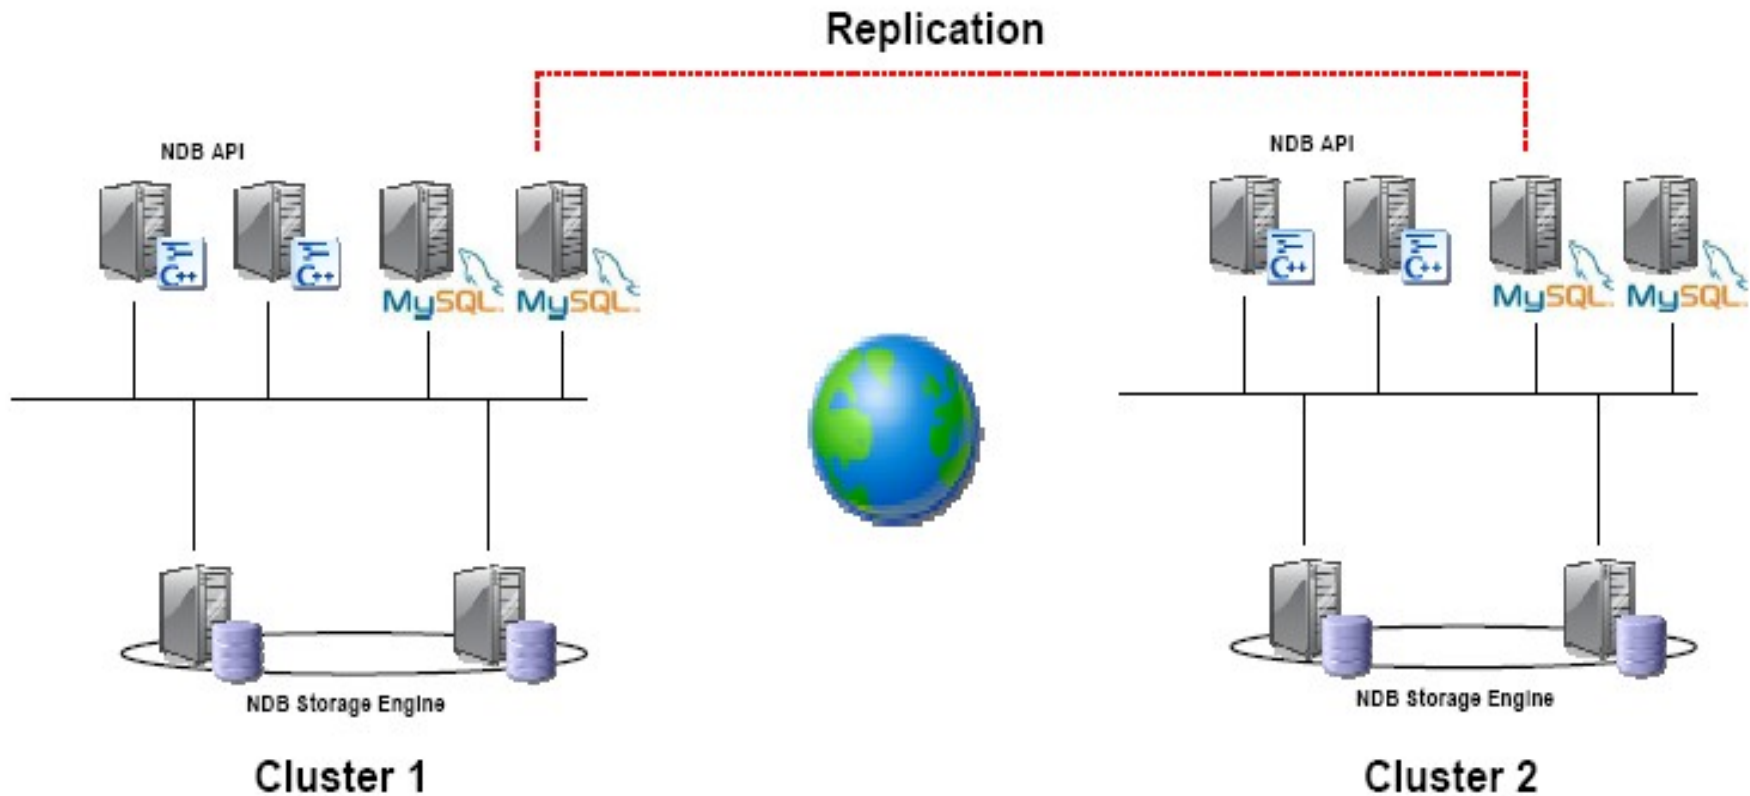


Basic Cluster Architecture




Inside MySQL Replication

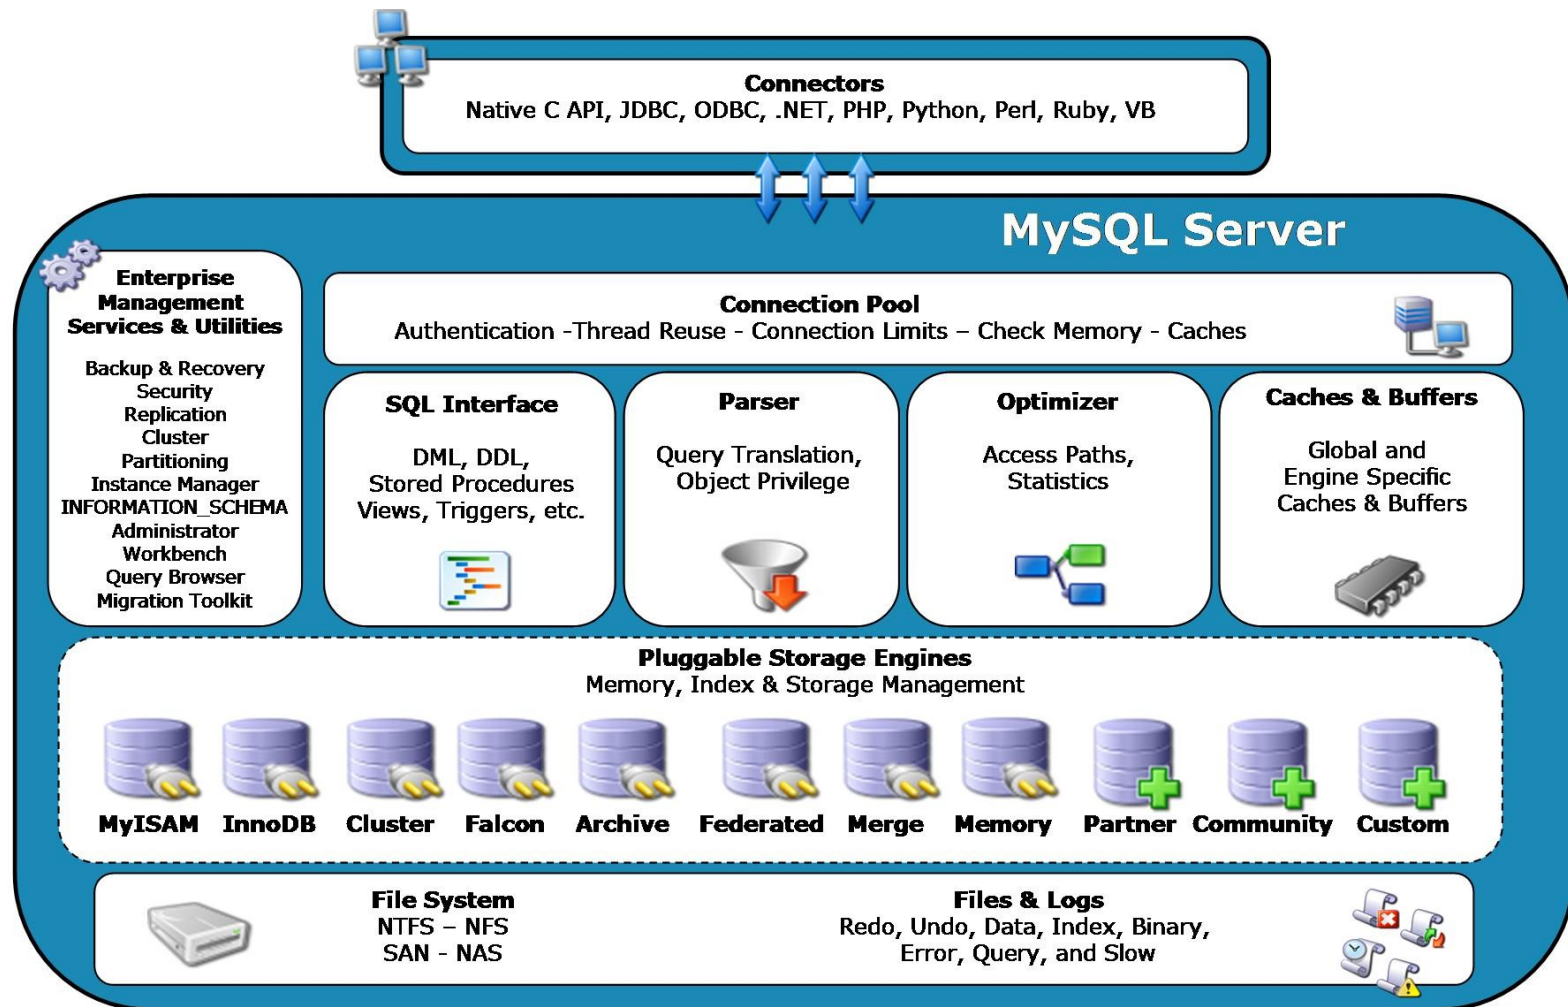
- Asynchronous MySQL Replication between Clusters
- Enables a higher degree of availability and continuity



High Availability Solutions



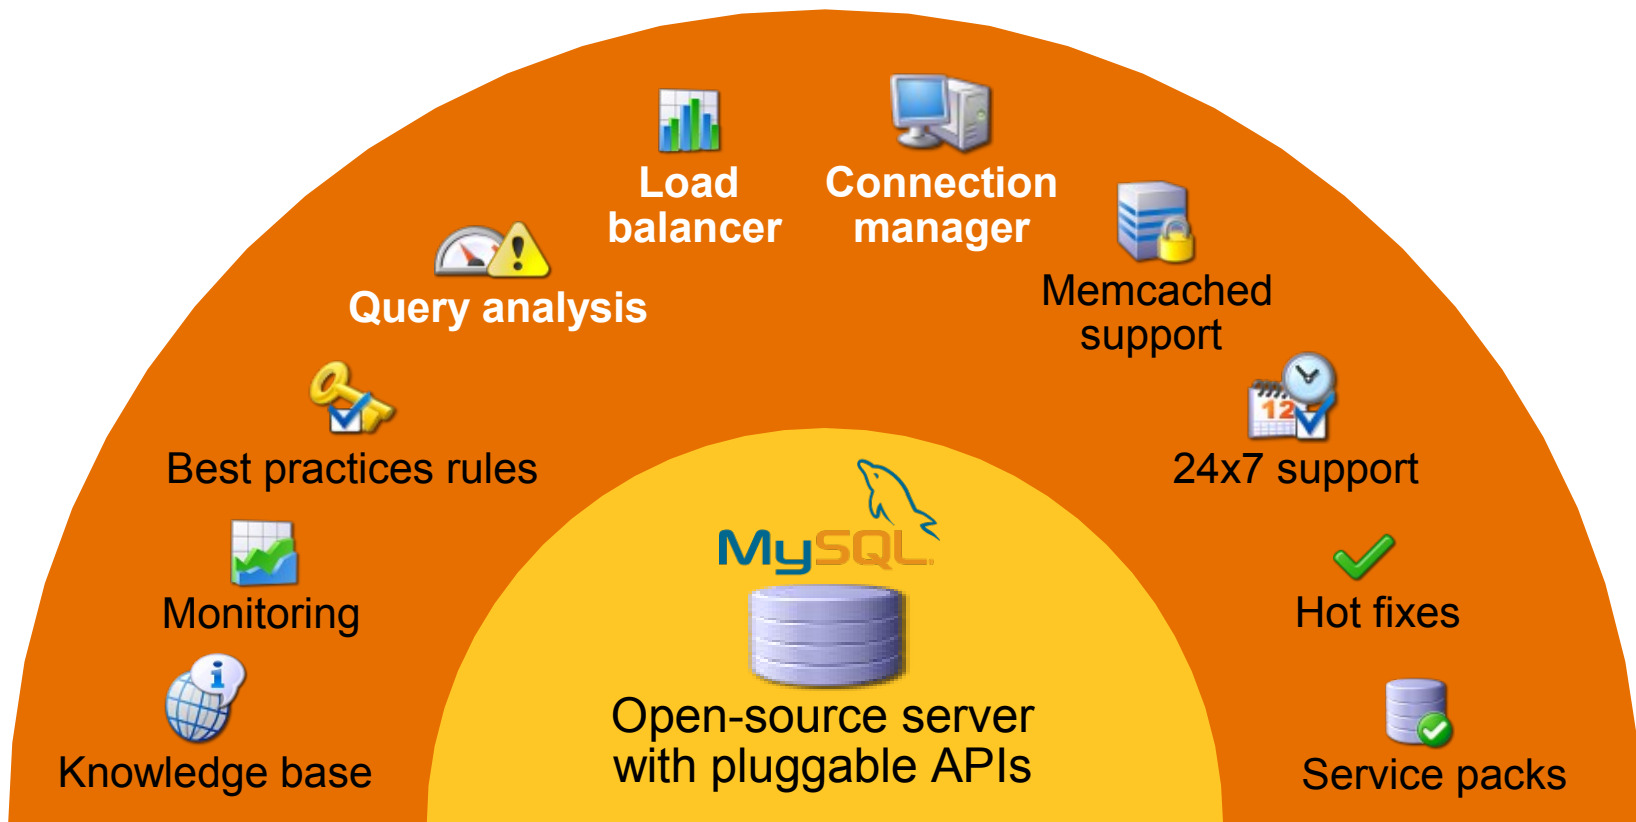
Pluggable Storage Engine Architecture



Choose the right storage engine for the job!

MySQL Enterprise Value Proposition

- Reduce DBMS TCO by 90% with MySQL Enterprise subscription
- Key differentiator: MySQL Enterprise Monitor reduces downtime.



MySQL 5

Performance

5

- Stored Procedures
- Cluster query push down
- Query optimisations
- Archive Engine
- InnoDB storage

5.1

- Table/Index Partitioning
- Full-Text Index Improvements
- Archive Engine Improvements
- Faster ALTER TABLE
- Faster ADD/DROP Index
- Parallel Data Import
- New Load Testing Utility
- MyISAM Memory Option
- New Process/SQL Diagnostics

Reliability

5

- SQL Mode
- Triggers
- Views
- Precision Math
- Distributed Transactions
- Increased object support

5.1

- Row-Based Replication
- Disk-based Cluster
- Replication for Cluster
- Federated Engine Txn

Ease of Use

5

- Instance Manager
- Information Schema
- Cursors
- Enhanced GUI Tools
- Migration Toolkit

5.1

- XML/XPath Support
- Task Scheduler
- Storage Engine Plug-in API

MySQL 5.1 - Events

- Events can execute stored procedure at a specific time or in a interval
 - > Similar to Unix "cron job" or the Windows Task Scheduler
 - > Either on-shot or repeating events
 - > Repeating events can have end-time which also might remove the event

CREATE EVENT


- One Shot events
 - > AT <timestamp>
- Repeating events
 - > EVERY <INTERVAL> [STARTS <timestamp>] [ENDS <timestamp>]
- Self-Removing events
 - > ON COMPLETION [NOT] PRESERVES

```
CREATE EVENT ev_name ON SCHEDULE
    AT CURRENT_TIMESTAMP + INTERVAL 1 HOUR
    (or) EVERY 1 WEEK STARTS '00:01:00' ENDS ...
    DO CALL <sp>
```

Partitioning

- Reasons to use partitioning
 - > Performance
 - > Ease of Management
 - VLDB
- Distribution of rows is controlled by a partitioning expression
- The MySQL Cluster was using Partitioning since the first release

Horizontal Partitioning



Partition Expressions

- KEY splits by PRIMARY KEY
- HASH hashes a field with MD5 or PASSWORD
- RANGE creates a partition for a range of values
- LIST creates a partition for each set of values

Partition Pruning

- READs from a partition only take those partitions into account which pass the WHERE clause
- If you only have to scan 1 of 4 partitions the query is executed 4 times faster
 - > EXPLAIN PARTITION
SELECT * from CUST_ORDER where
YEAR(OORDER_DATE)=1996

Partitioned by Date – Year(Order Date)

Partition 1996

Order Id	Order Date	Descn	Qty

Partition 1997

Order Id	Order Date	Descn	Qty

Partition 1998

Order Id	Order Date	Descn	Qty

Partition 1999

Order Id	Order Date	Descn	Qty

Managing Partitions

- ALTER TABLE – to change the layout of partitions
 - > ADD/DROP PARTITION
 - > REORGANIZE – to change the partition rule
 - > COALESCE – to reduce the number of partition for HASH & KEY partitions
 - > Syntax
 - ADD PARTITION (partition_definition)
 - DROP PARTITION partition_names
 - COALESCE PARTITION number
 - REORGANIZE PARTITION partition_names INTO (partition_definitions)
 - ANALYZE PARTITION partition_names
 - CHECK PARTITION partition_names
 - OPTIMIZE PARTITION partition_names
 - REBUILD PARTITION partition_names
 - REPAIR PARTITION partition_names
 - REMOVE PARTITIONING

Row-based Replication (MySQL 5.1)

- Statement-based replication
 - > Replicate statement doing changes
 - > Requires up-to-date slave
 - > Requires determinism
- Row-based replication
 - > Replicate actual row changes
 - > Does not require up-to-date slave
 - > Can handle any statement

Logging modes

- SET BINLOG_FORMAT=mode
 - > Session and global variable
 - > Mode is one of STATEMENT, ROW, or MIXED
 - > STATEMENT: statements are logged in statement format
 - > ROW: statements are logged in row format
 - > MIXED (default)
- Statements are logged in statement format by default
- Statements are logged in row format in some cases (such as DDL)

MIXED mode

- Safe statements are usually logged in statement format
- Unsafe statements are logged in row format
- Heuristic decision on what is unsafe, currently:
 - > Statement containing UUID() or calls to UDFs
 - > Statements updating >1 table with auto-increment columns
 - > INSERT DELAYED statements
 - problems with RAND() and user-defined variables

Example: multi-table update

- UPDATE t1,t2 SET t1.b = ..., t2.b = ...

```
mysql> show binlog events from 480;
```

Log_name	Pos	Event_type	...	Info
...	480	Table_map	...	table_id: 16 (test.t1)
...	520	Table_map	...	table_id: 17 (test.t2)
...	560	Update_rows	...	table_id: 16
...	625	Update_rows	...	table_id: 17 flags: STMT_END_F

4 rows in set (0.00 sec)

Example: CREATE-SELECT

- `CREATE t3 SELECT * FROM t1`

```
mysql> show binlog events from 690;
```

Log_name	Pos	Event_type	...	Info
...	480	Table_map	...	use `test`; CREATE TABLE `t3` (a INT(11) DEFAULT NULL, b INT(11) DEFAULT NULL)
...	520	Table_map	...	table_id: 18 (test.t3)
...	625	Write_rows	...	table_id: 18 flags: STMT_END_F

3 rows in set (0.00 sec)

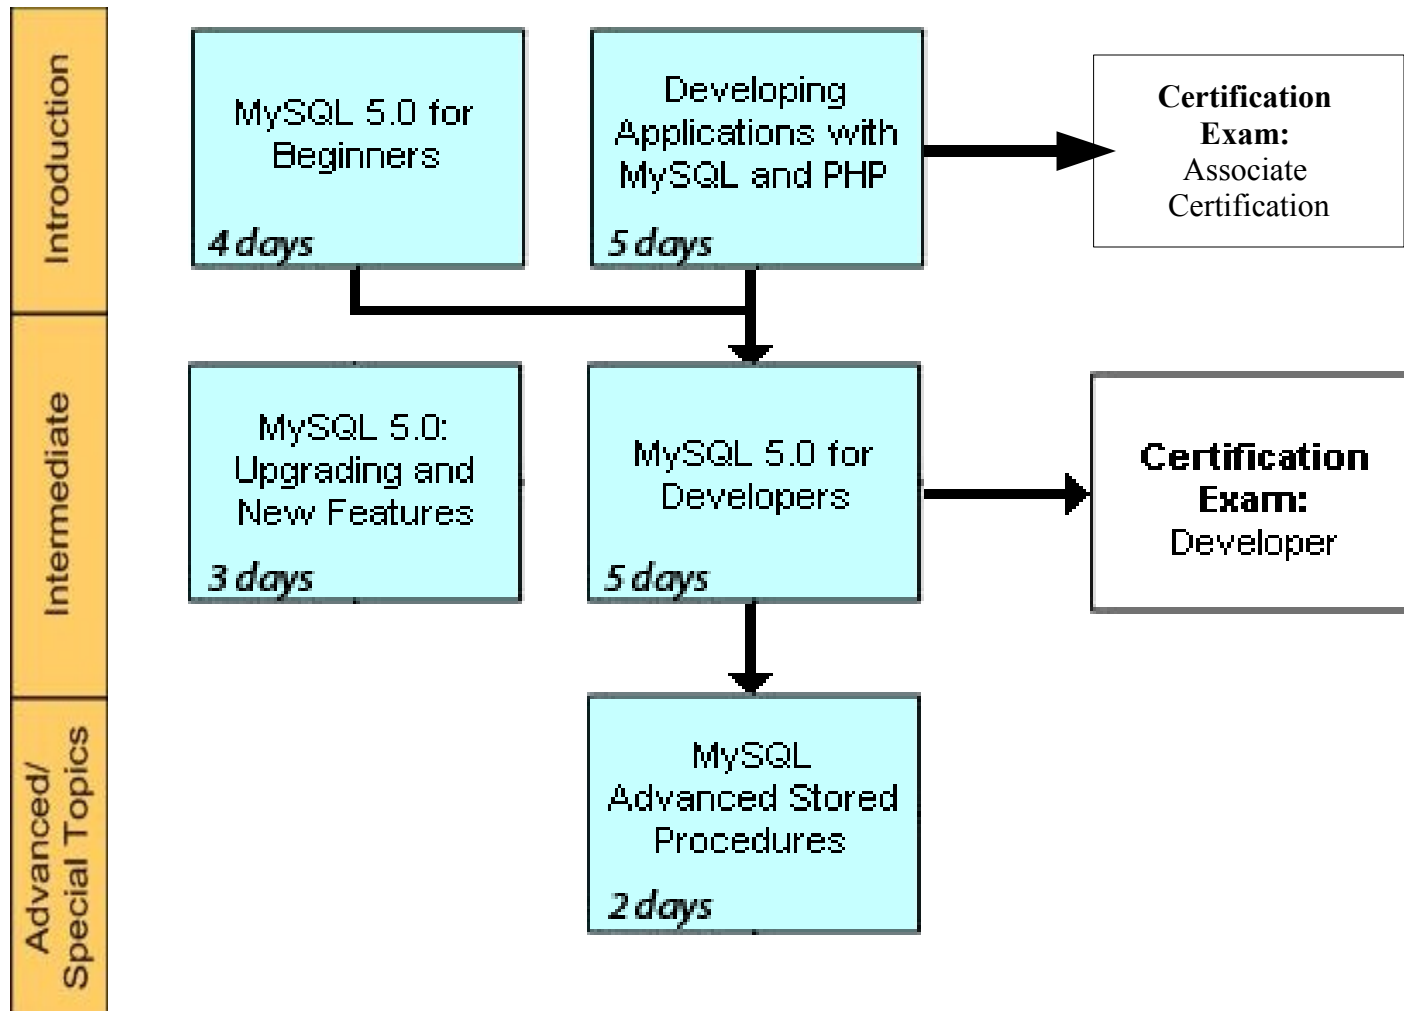
Row-based vs Statement-based Replication

- Advantages of Row-based Replication (RBR)
 - > Can replicate non-deterministic statements (e.g. UDFs, LOAD_FILE(), UUID(), USER(), FOUND_ROWS())
 - > Makes it possible to replicate between MySQL Clusters (having multiple MySQL servers or using NDB API)
 - > Less execution time on slave
 - > Simple conflict detection (that is currently being extended)
- Advantages of Statement-based Replication (SBR)
 - > Proven technology (since MySQL 3.23)
 - > Sometimes produces smaller log files
 - > Binary log can be used for auditing

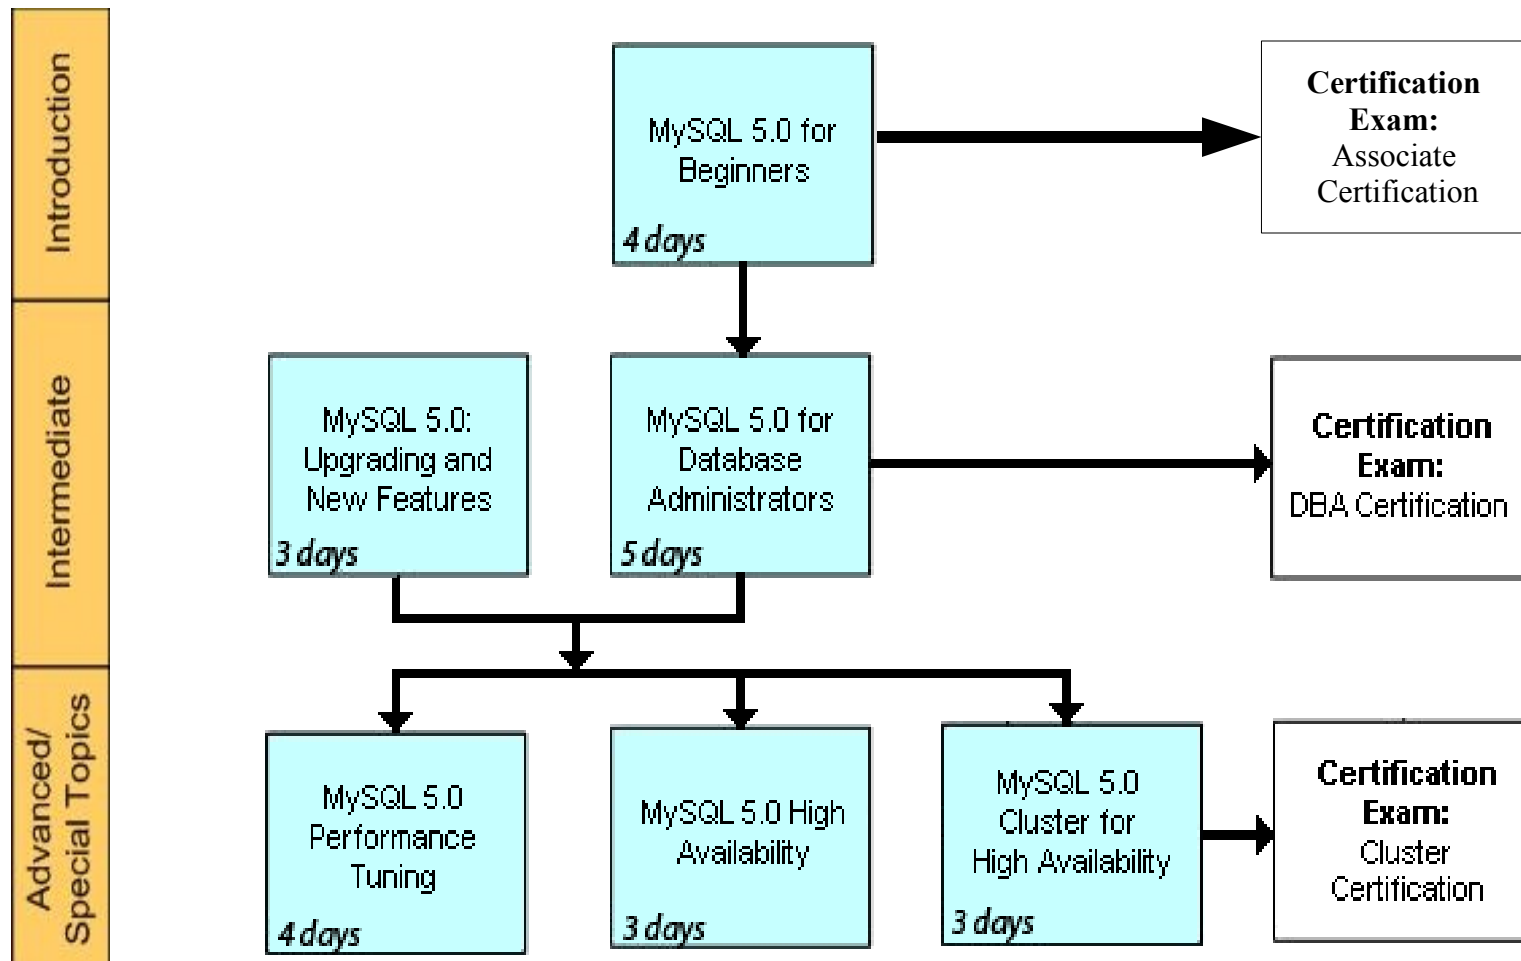
MySQL Certification and Training

- MySQL Certified Associate
- MySQL Certified Database Administrator
- MySQL Certified Developer
- MySQL Certified Cluster Database Administrator

MySQL Certified Developers



MySQL Certified DBAs



Making Users / Making Customers

Free to All

- **MySQL Community Server**
- **MySQL GUI Management Tools**
- **MySQL Connectors (JDBC, ODBC, etc.)**
- **Documentation**
- **Forums**

Customer

- **Subscription:**
 - **MySQL Enterprise**
 - **MySQL Unlimited**
- **License (OEM):**
 - **Embedded Server**
 - **Support**
- **MySQL Cluster Carrier-Grade**
- **Training**
- **Consulting**

Key Takeaway

- Free
- Widely Extensible and Adopted
- Enterprise and Developers Productivity Features – Join our hands-on workshop to know more!!!
- Download and start using it

URL for MySQL

www.mysql.org





MySQL Experience Sharing

mickey.fan@sun.com

Sun Microsystems



Appendix - Demo

- MySQL Workbench
- MySQL Query Workbench
- MySQL Migration Assistant
- MySQL DBA Administrator
- MySQL Enterprise Monitoring
- MySQL Partitioning