ORACLE OPEN WORLD
MySQL Central @ OPENWORLD
October 25-29, 2015
San Francisco

#MySQLCentral

The State of the Dolphin: 1.6M QPS (SQL) with MySQL 5.7

Sastry Vedantam
Sastry.vedantam@oracle.com

ORACLE

MySQL

Safe Harbor Statement

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 decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

A Year of Anniversaries!

20 Years: MySQL

10 Years: Oracle stewardship of InnoDB

5 Years: Oracle stewardship of MySQL

Thank You!

ORACLE

The Best MySQL Solutions Ever

Investing & Innovating for You

- ✓ Performance
- ✓ Scalability
- ✓ Manageability
- ✓ Reliability
- ✓ Security
- ✓ Flexibility

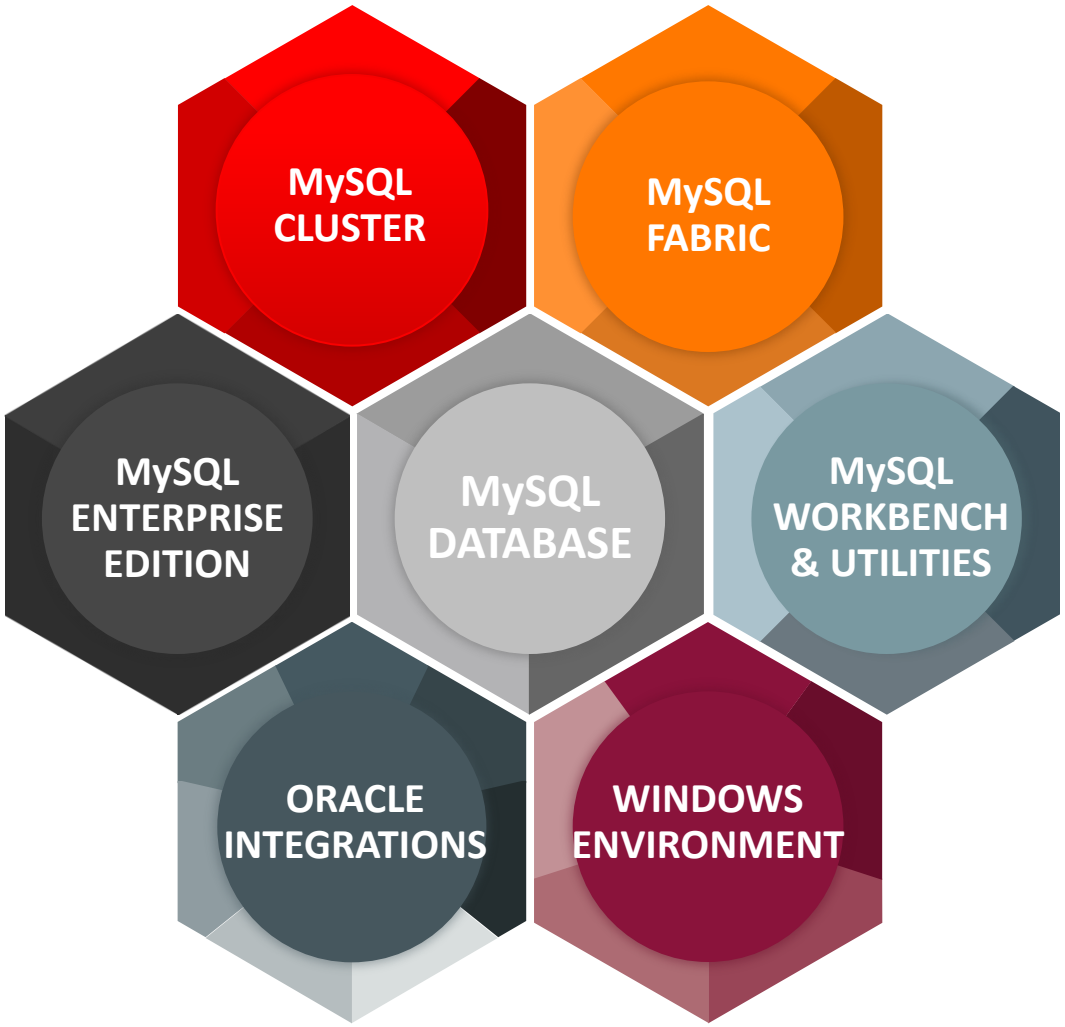

ORACLE

MySQL 5.7 is GA!

Performance & Scalability

- 3 X Faster than MySQL 5.6
- Enhanced InnoDB: faster online & bulk load operations
- Replication Improvements (incl. multi-source, multi-threaded slaves...)
- New Optimizer Cost Model: greater user control & better query performance

Manageability

- Native JSON Support
- Improved Security: safer initialization, setup & management
- Performance Schema Improvements
- MySQL SYS Schema

And many more new features and enhancements. Learn more at: dev.mysql.com

MySQL 5.7 Sysbench Benchmark: SQL Point Selects

3x Faster than MySQL 5.6

1,600,000 QPS

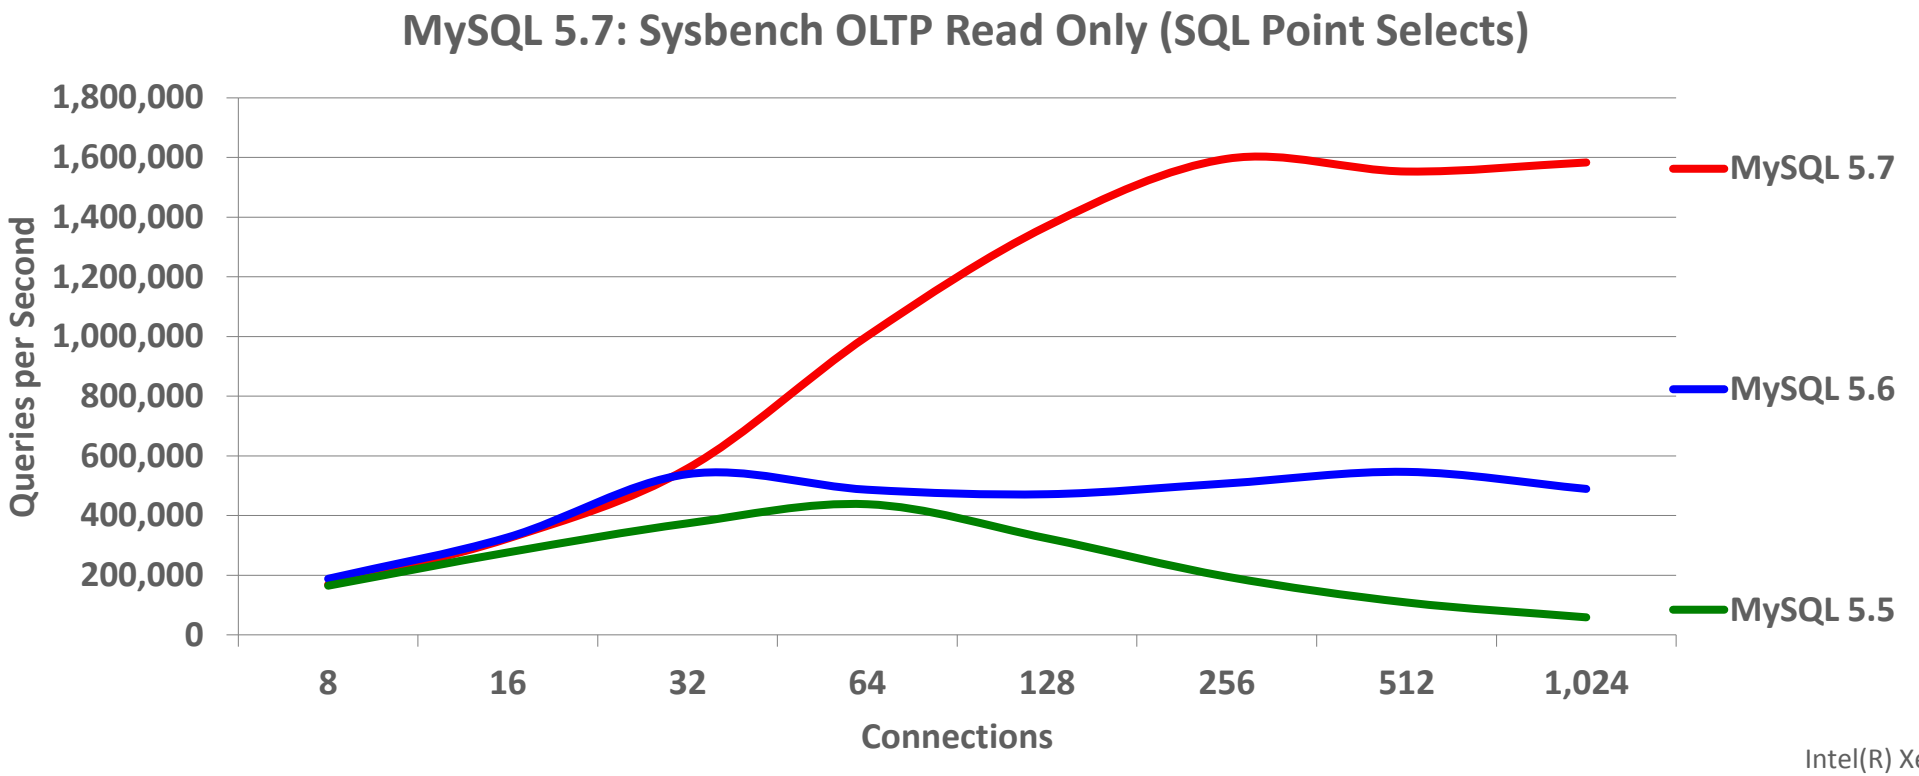

http://dimitrik.free.fr/blog/archives/2015/10/mysql-performance-yes-we-can-do-more-than-16m-qps-sql-on-mysql-57-ga.html

Intel(R) Xeon(R) CPU E7-8890 v3
4 sockets x 18 cores-HT (144 CPU threads)
2.5 Ghz, 512GB RAM
Linux kernel 3.16

ORACLE

MySQL 5.7: Connections per Second

1.7x Faster than MySQL 5.6

3x Faster than MySQL 5.5

100,000 Connections/Sec

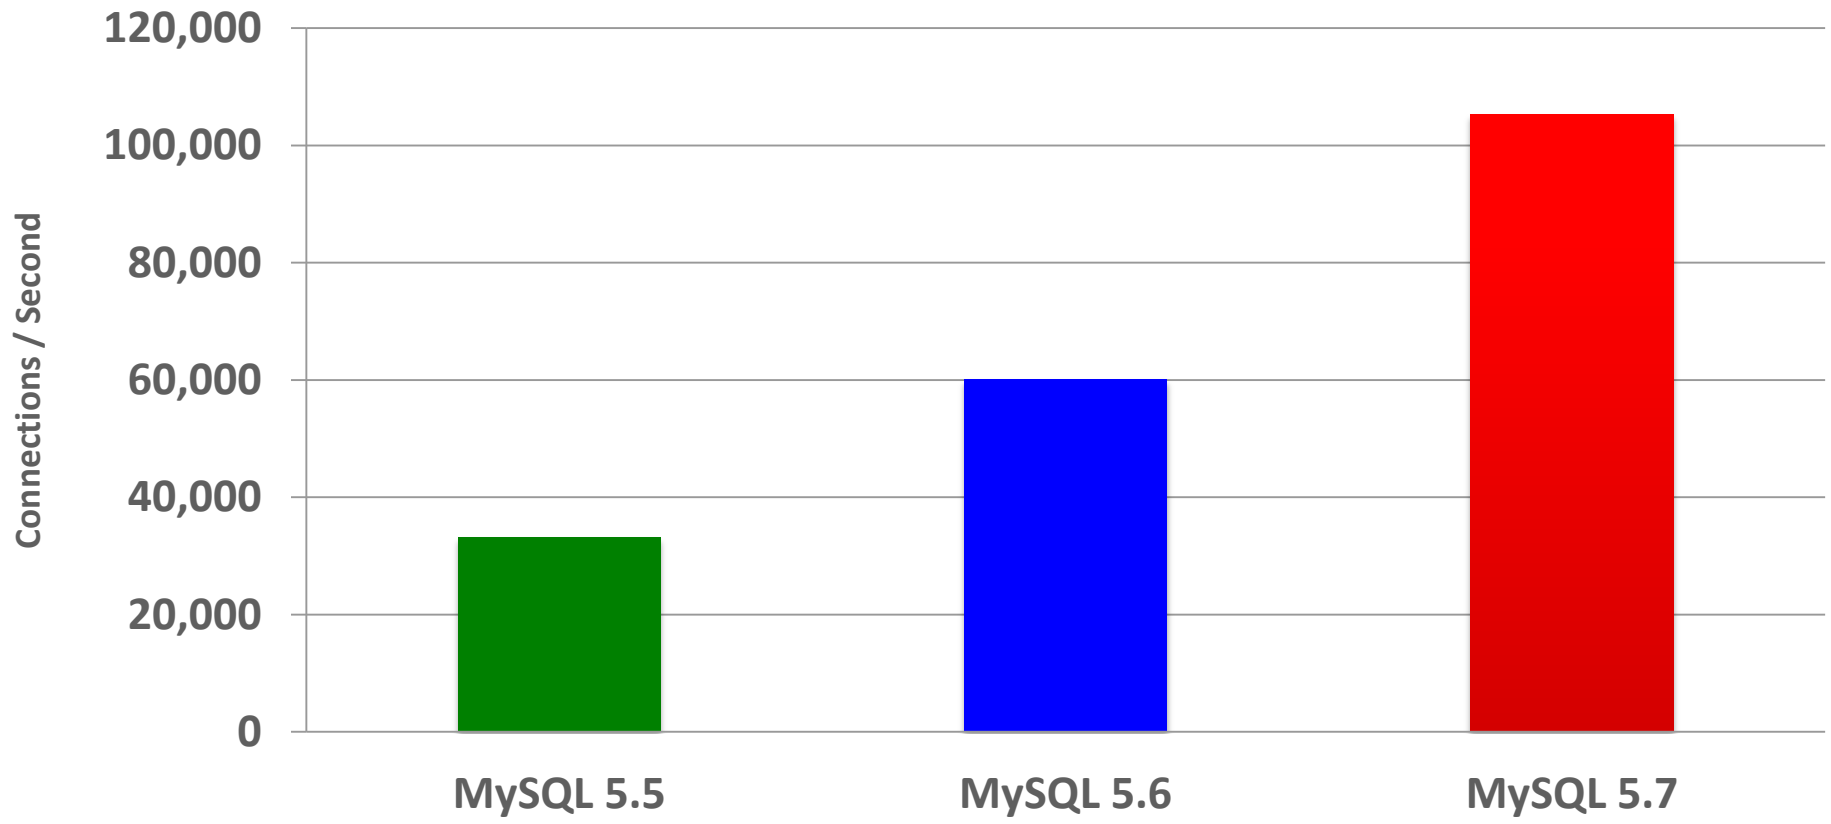

Intel(R) Xeon(R) CPU E7-8890 v3
4 sockets x 18 cores-HT (144 CPU threads)
2.5 Ghz, 512GB RAM
Linux kernel 3.16

ORACLE

MySQL 5.7 Sysbench Benchmark: Mixed OLTP Read Only

3x Faster than MySQL 5.6

~ 1,000,000 QPS

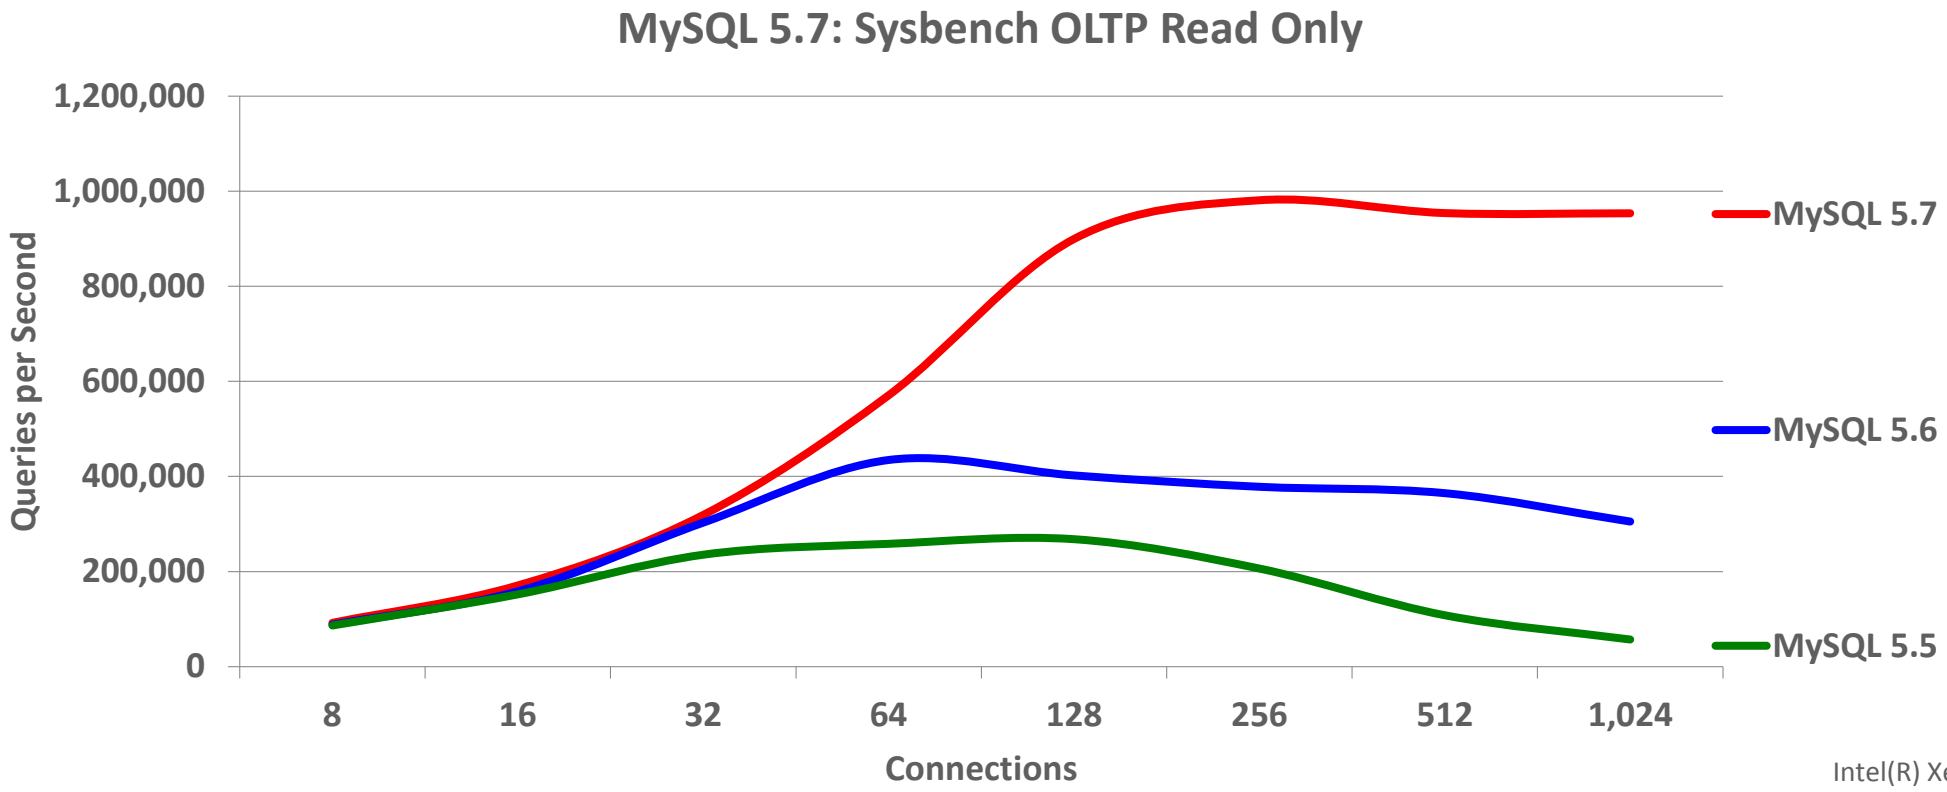

http://dimitrik.free.fr/blog/archives/2015/11/mysql-performance-1m-qps-on-mixed-oltp_ro-on-mysql-57-ga.html

Intel(R) Xeon(R) CPU E7-8890 v3
4 sockets x 18 cores-HT (144 CPU threads)
2.5 Ghz, 512GB RAM
Linux kernel 3.16

MySQL 5.7 Sysbench Benchmark: OLTP Read Write

1.5x Faster than MySQL 5.6

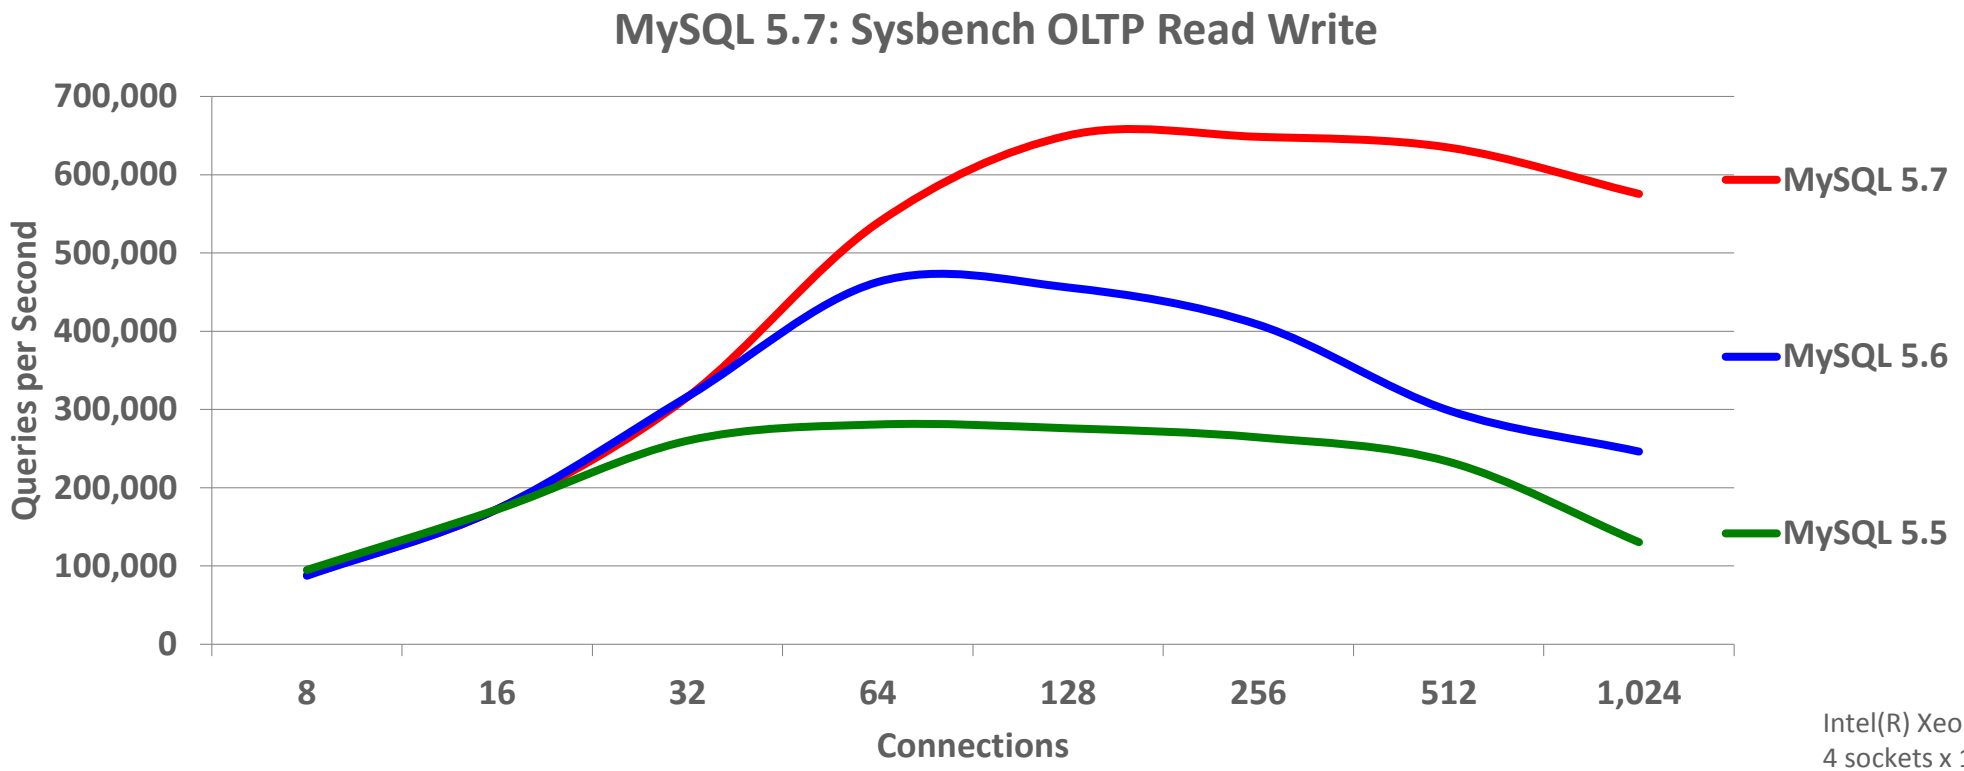

Intel(R) Xeon(R) CPU E7-8890 v3
4 sockets x 18 cores-HT (144 CPU threads)
2.5 Ghz, 512GB RAM
Linux kernel 3.16

MySQL 5.7: Optimizer Improvements

Queries execute faster, while using less CPU and disk space!

- Optimizer and Parser refactoring
  - Readability, maintainability and stability
  - Separate parsing, optimizing, execution stages
  - Easier feature additions, with lessened risk
- New hint framework
  - Easier to manage
  - With support for additional new hints
- Improved JSON EXPLAIN
- EXPLAIN for running thread

- New Cost based Optimizer
  - Easier to extend
  - Configurable and tunable
    - mysql.server_cost and mysql.engine_cost tables
    - API for where data resides: on disk or in cache
- InnoDB for internal temp tables
- Better ONLY_FULL_GROUP_BY mode
- Many specific new optimizations
- Generated Columns & Functional Indexes

MySQL 5.7: Query Rewrite Plugin

- New pre and post parse query rewrite APIs
  - Users can write their own plug-ins
- Provides a post-parse query plugin
  - Rewrite problematic queries without the need to make application changes
  - Add hints
  - Modify join order
  - Many more …
- Improve problematic queries from ORMs, third party apps, etc
- Eliminates many legacy use cases for proxies

MySQL 5.7: Optimizer - Cost Info in JSON EXPLAIN

- Expanded JSON EXPLAIN
  - Now includes all available cost info
  - Used for Visual Explain In MySQL Workbench

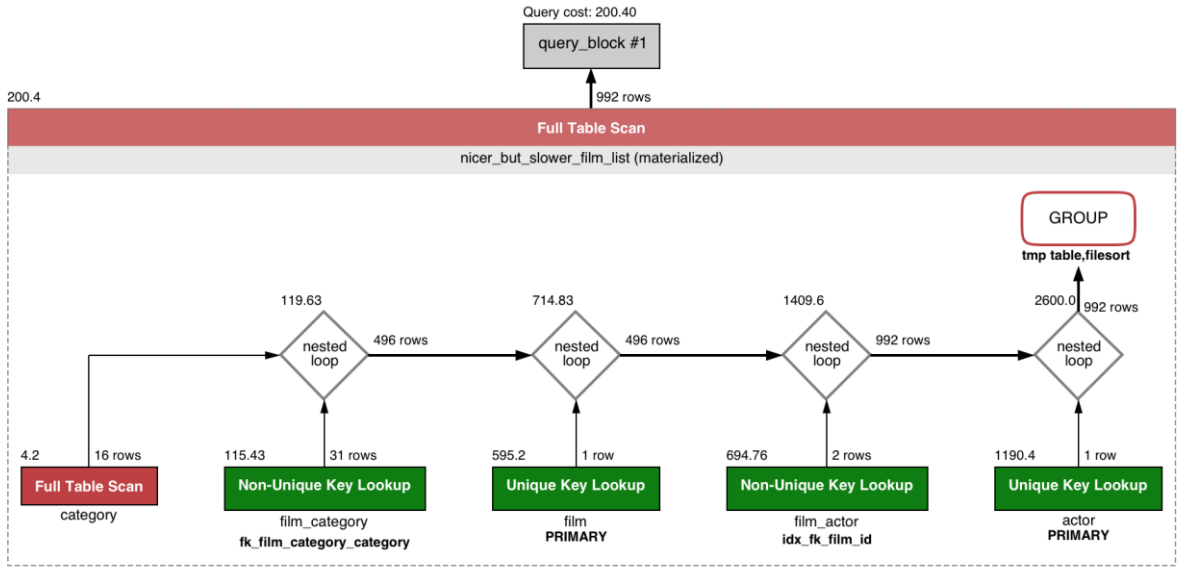


```
{
  "query_block": {
    "select_id": 1,
    "cost_info": {
      "query_cost": "200.40"
    },
    "table": {
      "table_name": "nicer_but_slower_film_list",
      "access_type": "ALL",
      "rows_examined_per_scan": 992,
      "rows_produced_per_join": 992,
      "filtered": 100,
      "cost_info": {
        "read_cost": "2.00",
        "eval_cost": "198.40",
        "prefix_cost": "200.40",
        "data_read_per_join": "852K"
      },
      "used_columns": [
        "FID",
        "title",
        "description",
        "category",
        "price",
        "length",
        "rating",
        "actors"
      ],
...
```

MySQL 5.7: JSON

- Native JSON data type
  - Native internal binary format for efficient processing & storage
- Built-in JSON functions
  - Allowing you to efficiently store, search, update, and manipulate Documents
- JSON Comparator
  - Allows for easy integration of Document data within your SQL queries
- Indexing of Documents using Generated Columns
  - InnoDB supports indexes on both stored and virtual Generated Columns
  - New expression analyzer automatically uses the best “functional” index available
- New inline syntax for easy SQL integration

MySQL 5.7: JSON and Text Datatype Comparison

Unindexed traversal of 206K documents

```
# With feature column as JSON type
SELECT DISTINCT
 feature->"$.type" as json_extract
FROM features;
+--------------+
| json_extract |
+--------------+
| "Feature"    |
+--------------+
1 row in set (1.25 sec)
```

```
# With feature column as TEXT type
SELECT DISTINCT
 feature->"$.type" as json_extract
FROM features;
+--------------+
| json_extract |
+--------------+
| "Feature"    |
+--------------+
1 row in set (12.85 sec)
```

Explanation: Binary format of JSON type is very efficient at searching. Storing as TEXT performs over 10x worse at traversal.

MySQL 5.7: Functional Indexes with JSON

From table scan on 206K documents to index scan on 206K materialized values

```
ALTER TABLE features ADD feature_type VARCHAR(30) AS (JSON_UNQUOTE(feature->'$.type'));
Query OK, 0 rows affected (0.01 sec)
Records: 0  Duplicates: 0  Warnings: 0
```

Meta data change only (FAST).
Does not need to touch table.

```
ALTER TABLE features ADD INDEX (feature_type);
Query OK, 0 rows affected (0.73 sec)
Records: 0  Duplicates: 0  Warnings: 0
```

Creates index only, does not touch row data.

```
SELECT DISTINCT feature_type FROM features;
+--------------+
| feature_type |
+--------------+
| "Feature"    |
+--------------+
1 row in set (0.06 sec)
```

Down from 1.25 sec to 0.06 sec

MySQL 5.7: Performance Schema

Memory Instrumentation

- Aggregates 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

- Stored Procedures
- Stored Functions
- Prepared Statements
- Transactions

Additional Information

- Replication slave status
- MDL lock instrumentation
- Status and variables per thread
- Server stage tracking
- Track long running SQL
- Improved configuration and ease-of-use
- All while reducing total footprint and overhead

MySQL 5.7: SYS Schema

Helper objects for DBAs, Developers and Operations staff

- Helps simplify DBA / Ops tasks
  - Monitor server health, user, host statistics
  - Spot, diagnose, and tune performance issues
- Easy to understand views with insights into
  - IO hot spots, Locking, Costly SQL statements
  - Schema, table and index statistics
- SYS is similar to
  - Oracle V$ catalog views
  - Microsoft SQL DMVs (Dynamic Mgmnt Views)

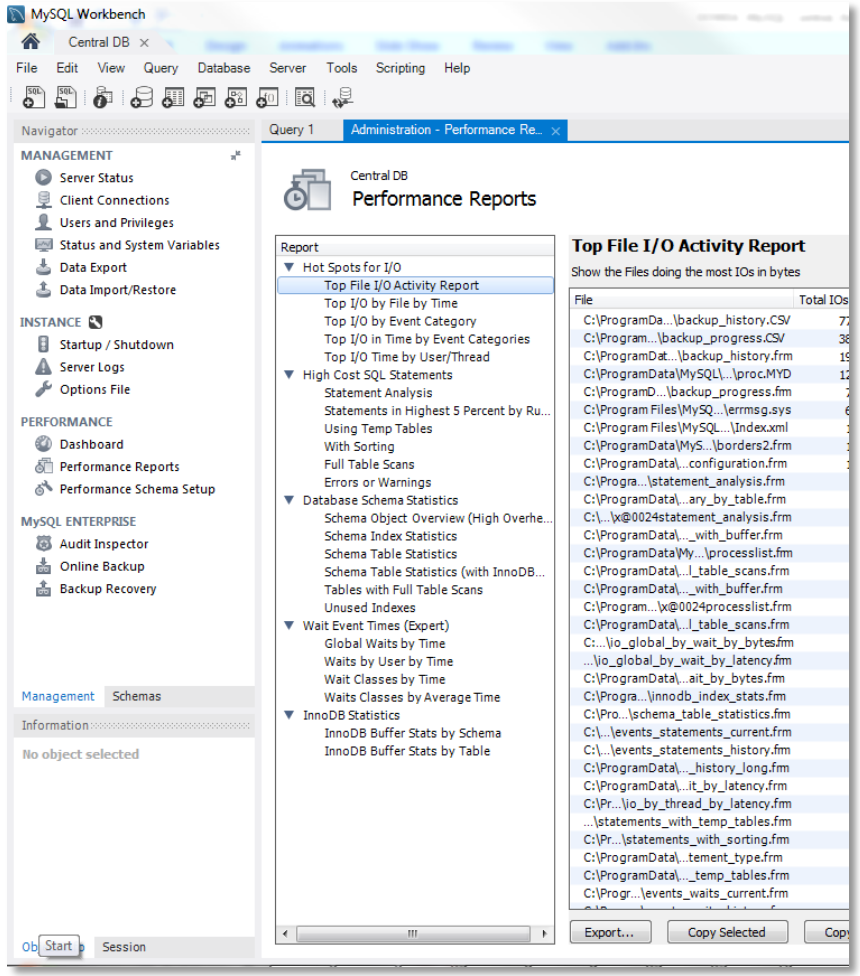

MySQL 5.7: GIS Improvements

- Replaced custom code with Boost.Geometry
  - For spatial calculations
  - For spatial analysis
  - Enabling full OGC compliance
  - We're also Boost.Geometry contributors!
- InnoDB R-tree based spatial indexes
  - Full ACID, MVCC, & transactional support
  - Index records contain minimum bounding box
- GeoHash
- GeoJSON
- Helper functions such as `ST_Distance_Sphere()` and `ST_MakeEnvelope()`

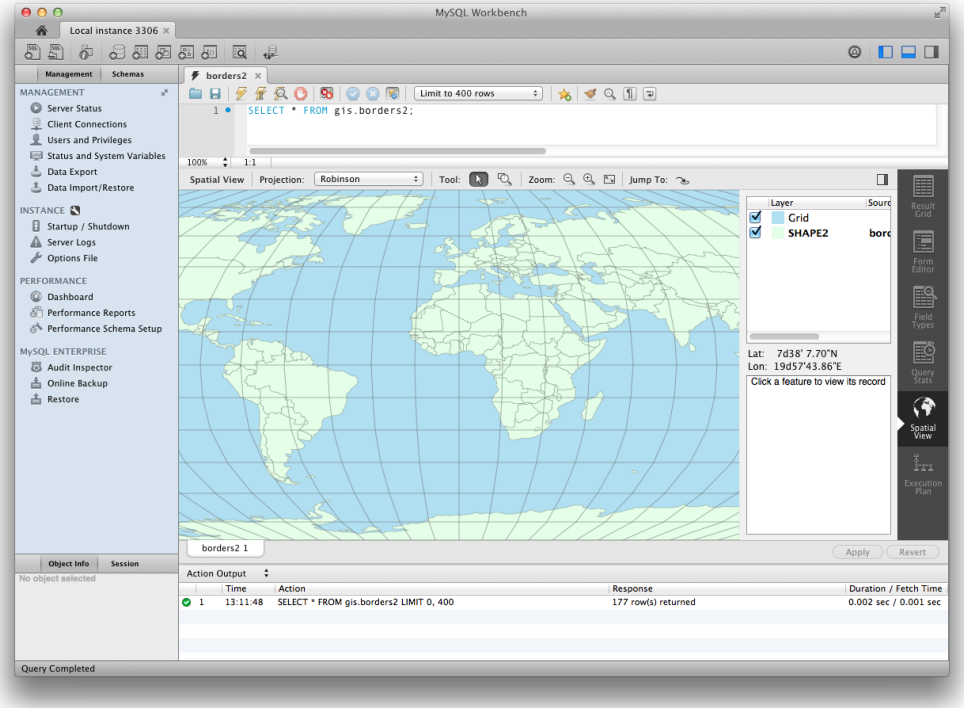

MySQL 5.7: InnoDB Improvements

- Native Partitioning
  - Eliminates previous limitations
  - Eliminates resource usage problems
  - Transportable tablespace support
- Native Full-Text Search
  - Including full CJK support!
- Native Spatial Indexes
- Transparent page compression
- Support for 32K and 64K pages
  - Use with transparent page compression for very high compression ratios
- General TABLESPACE support
  - Store multiple tables in user defined shared tablespaces
- Support for MySQL Group Replication
  - High priority transactions
- Improved support for cache preloading
  - Load your hottest data loaded at startup
- Configurable fill-factor
  - Allows for improvements in storage footprint
- Improved bulk-data load performance
- Resize the InnoDB Buffer Pool online

MySQL 5.7: InnoDB Compression

Thank you, SanDisk Fusion-io

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

MySQL 5.7: Security Improvements

- AES 256 Encryption now the default
- Password rotation policies
  - Can be set globally, and at the user level
- Deployment: enable secure unattended install by default
  - Random password set on install
  - Remove anonymous accounts
  - Deployment without test account, schema, demo files
- Easier instance initialization and setup: mysqld –initialize
- New detection and support for systemd

- SSL
  - Enabled by default
  - Auto-detection of existing keys and certs
  - Auto generation of keys and certs when needed
  - New helper utility: `mysql_ssl_rsa_setup`
  - New `--require_secure_transport` option to prevent insecure communications
  - Added SSL support to binary log clients
- Extended Proxy User Support
  - Added Built-in Authentication Plugins support for Proxy Users
  - Allows multiple users to share a single set of managed privileges

MySQL 5.7: Replication Improvements

- GTID enhancements
  - On-line, phased deployment of GTIDs
  - Binary logging on slave now optional
- Enhanced Semi-synchronous replication
  - Write guaranteed to be received by slave before being observed by clients of the master
  - Option to wait on Acks from multiple slaves
- Multi-Source Replication
  - Consolidate updates from multiple Masters into one Slave
- Dynamic slave filters

- **8-10x** Faster slave throughput
  - Often removes slave as a bottleneck; keep pace with master with 8+ slave threads
  - Option to preserve Commit order
  - Automatic slave transaction retries

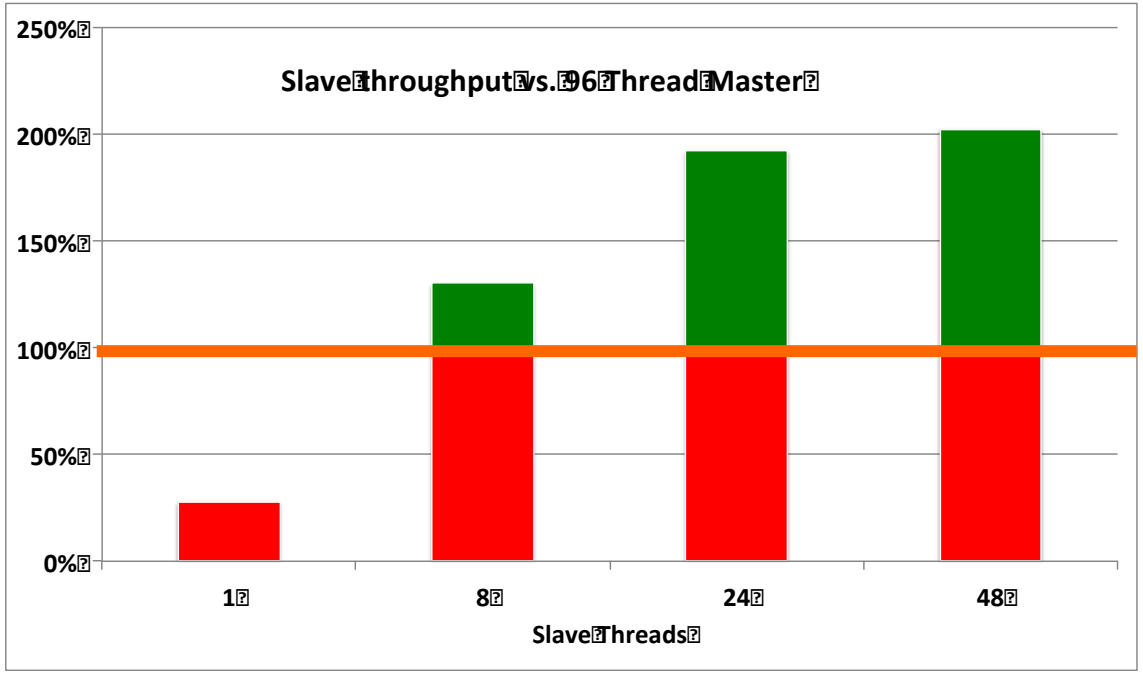

MySQL 5.7: Replication Improvements

- Multi-Source Replication
  - Consolidate updates from multiple Masters into one Slave
    - Consolidated view of all shards
    - More flexible topologies
    - Centralized point for backups
  - Compatible with Semi-Sync Replication & enhanced MTS
- Performance Schema tables for monitoring slave
- Online Operations: Dynamic Replication Filters, switch master

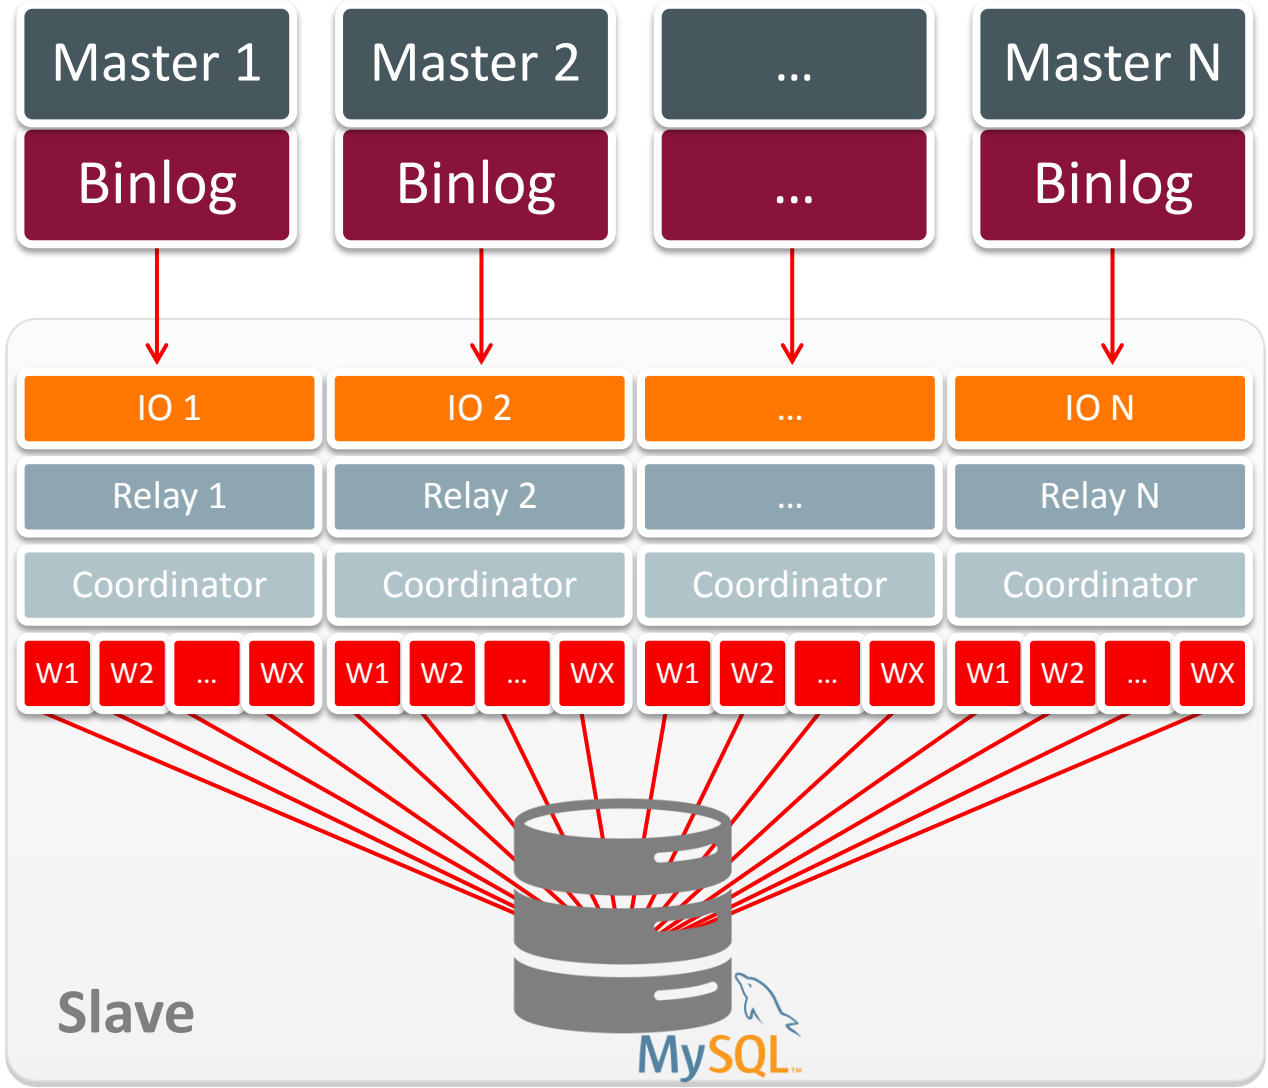

MySQL Router

- Connection and Transaction routing
- Transparently improve your MySQL apps
  - Transparent MySQL Fabric support
    - Transparent HA
    - Transparent Sharding
  - Transparent support for MySQL Group Replication clusters
  - Transparent support for custom clusters and HA setups
- Easily extendable using plugin APIs
- Many new plugins to come – Aggregation, Binary Log, Load Balancing, …

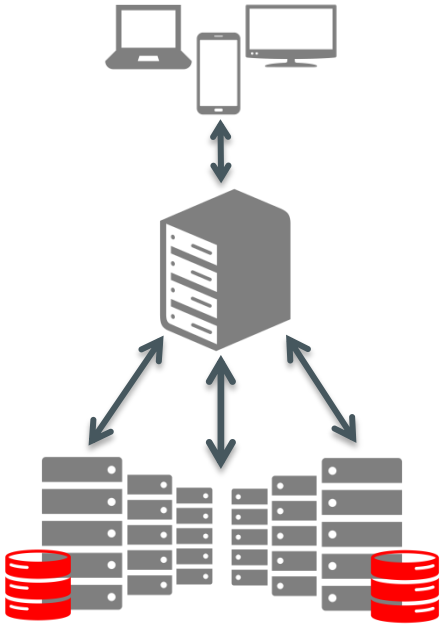

MySQL Fabric

High Availability + Sharding-Based Scale-out

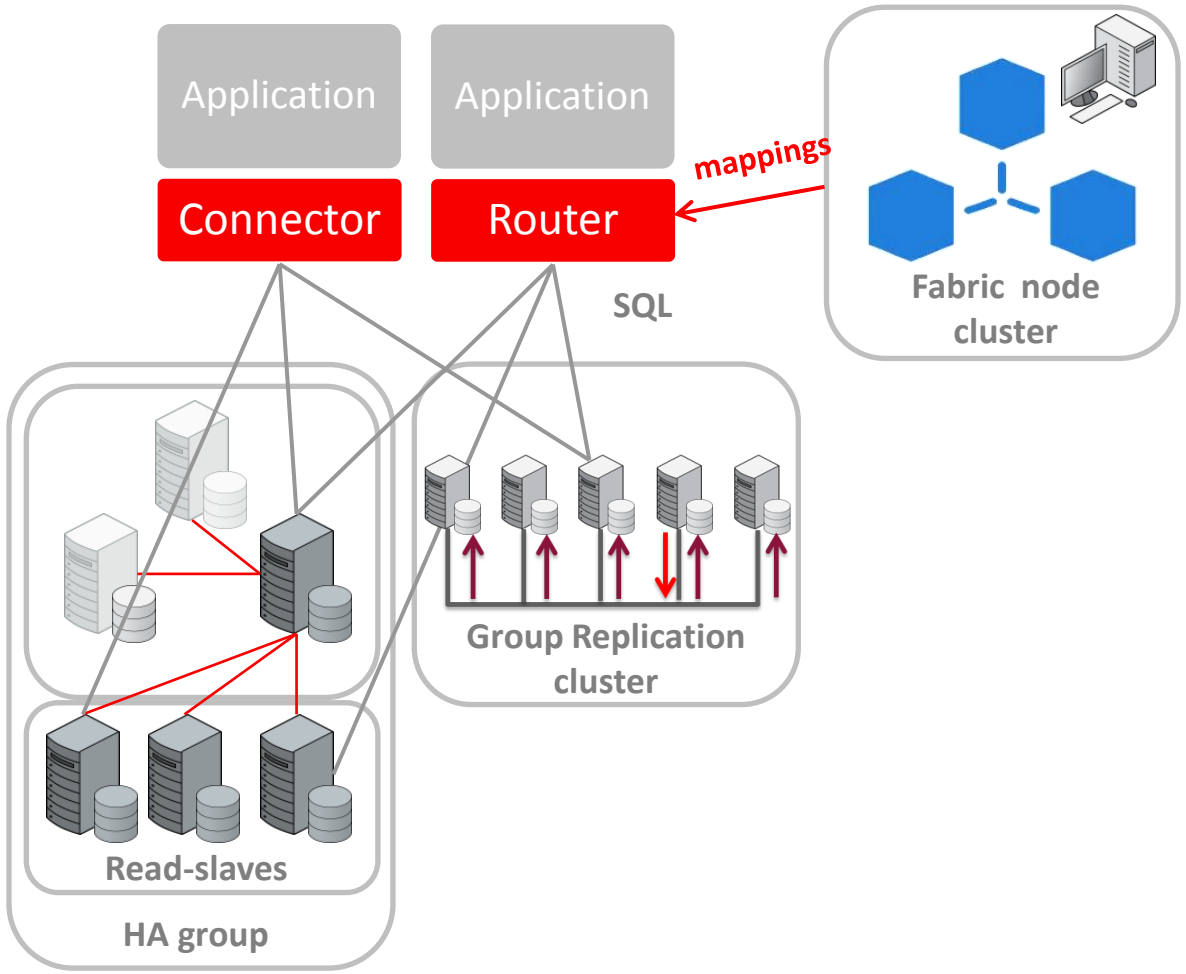

- High Availability
  - Server monitoring with auto-promotion and transparent application failover
- Scale-out through sharding
  - Application provides shard key
  - Tools for shard management
  - Global updates & tables
- Server provisioning using OpenStack
  - Support for Nova and Neutron APIs
- New in 1.6 Release (labs.mysql.com)
  - No single point of failure (SPOF) for HA
  - MySQL Router for connections, in addition to Fabric-aware connectors

The Future of MySQL Scaling (HA + Sharding)

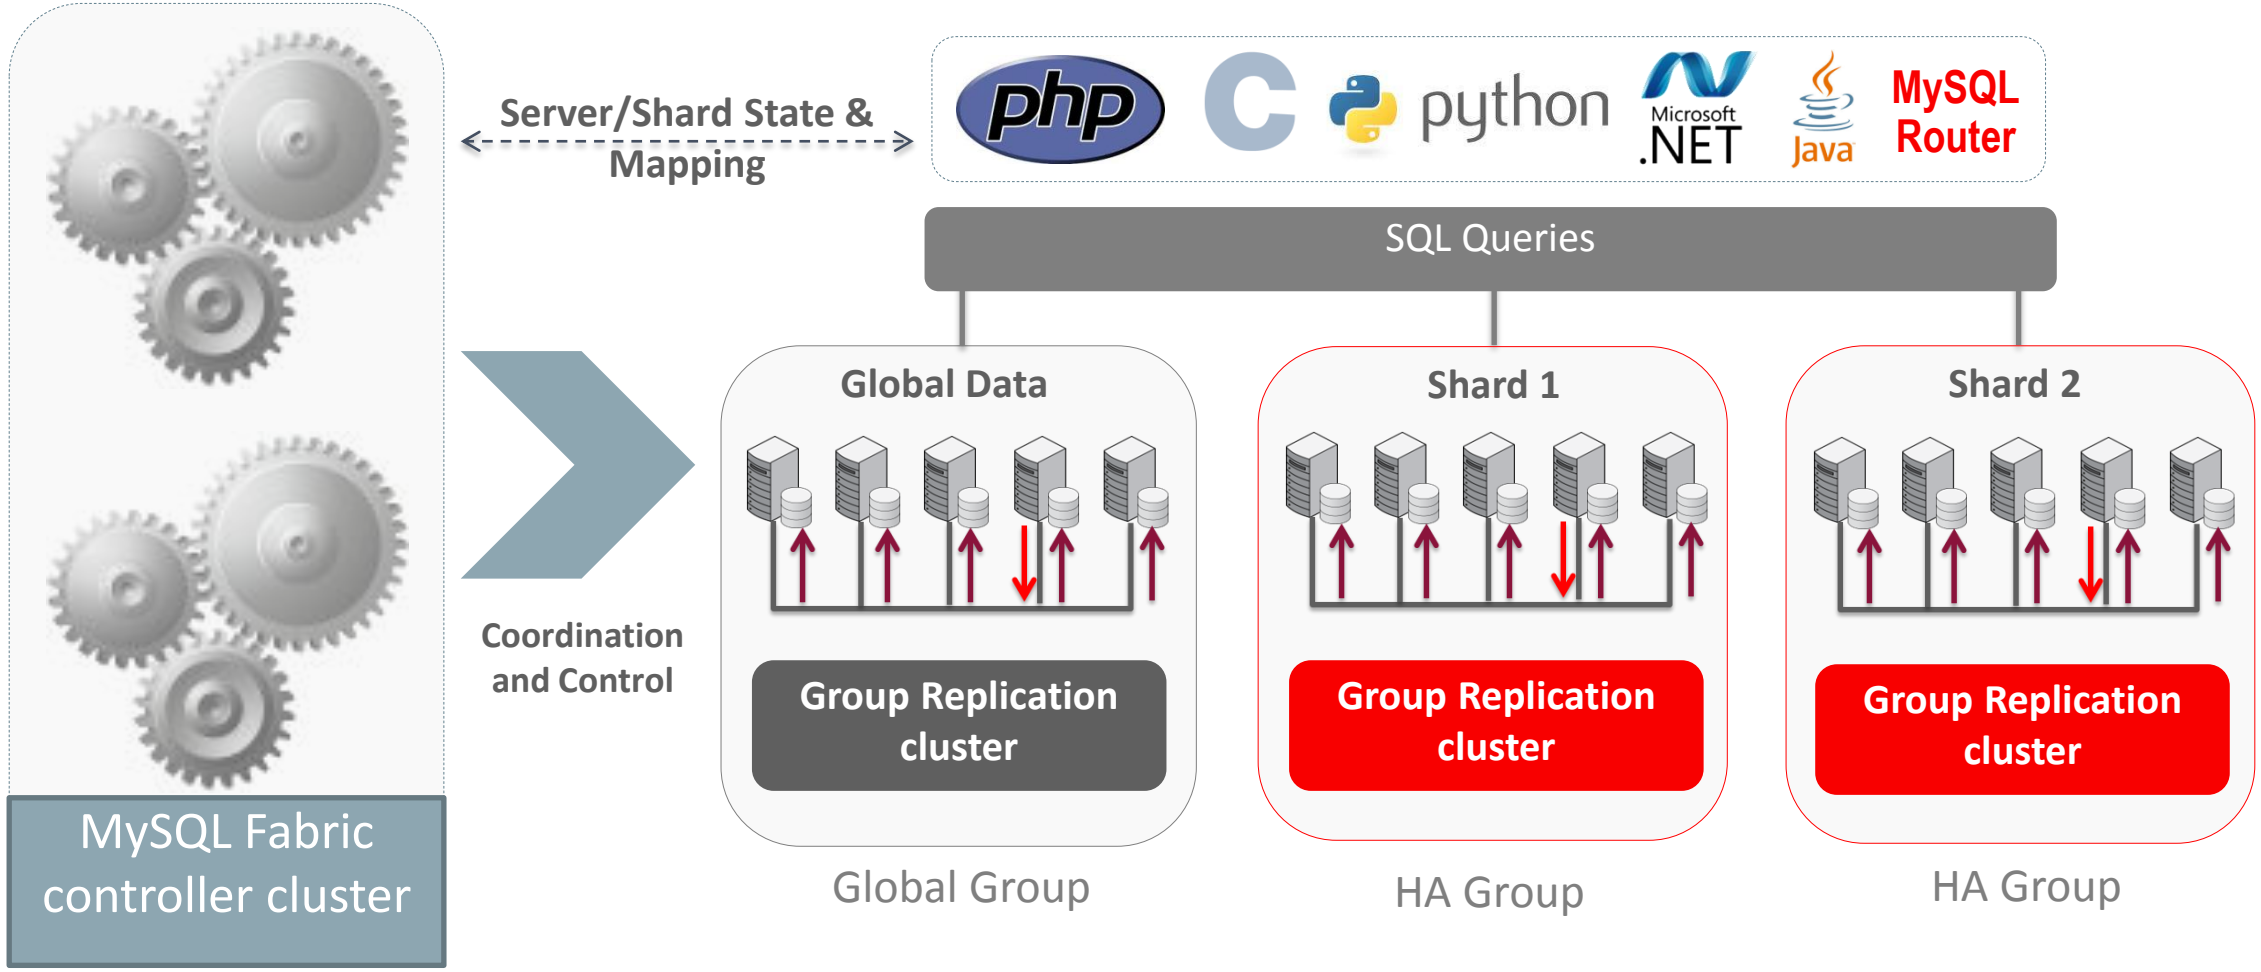

MySQL Group Replication

labs.mysql.com

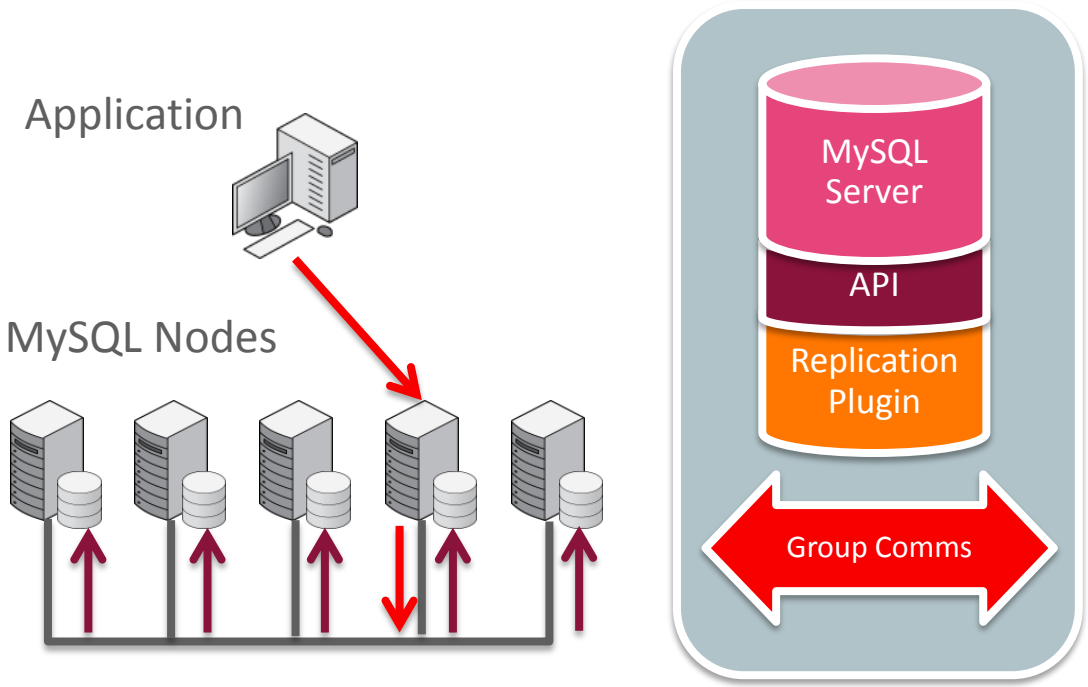

- Active/Active Update Anywhere
  - Conflict detection and resolution (transaction rollback)
  - Optimistic State Machine Replication
- Automatic group membership management and failure detection
  - No need for server fail-over
  - Elastic scale out/in
  - No single point of failure
  - Automatic reconfiguration
- Well integrated
  - InnoDB
  - GTID-based replication
  - PERFORMANCE_SCHEMA

MySQL Enterprise Edition

What's New

MySQL Enterprise Edition

Advanced Features

- Scalability
- High Availability
- Security
- Audit

Management Tools

- Monitoring
- Backup
- Development
- Administration
- Migration

Support

- Technical Support
- Consultative Support
- Oracle Certifications

MySQL Security Overview

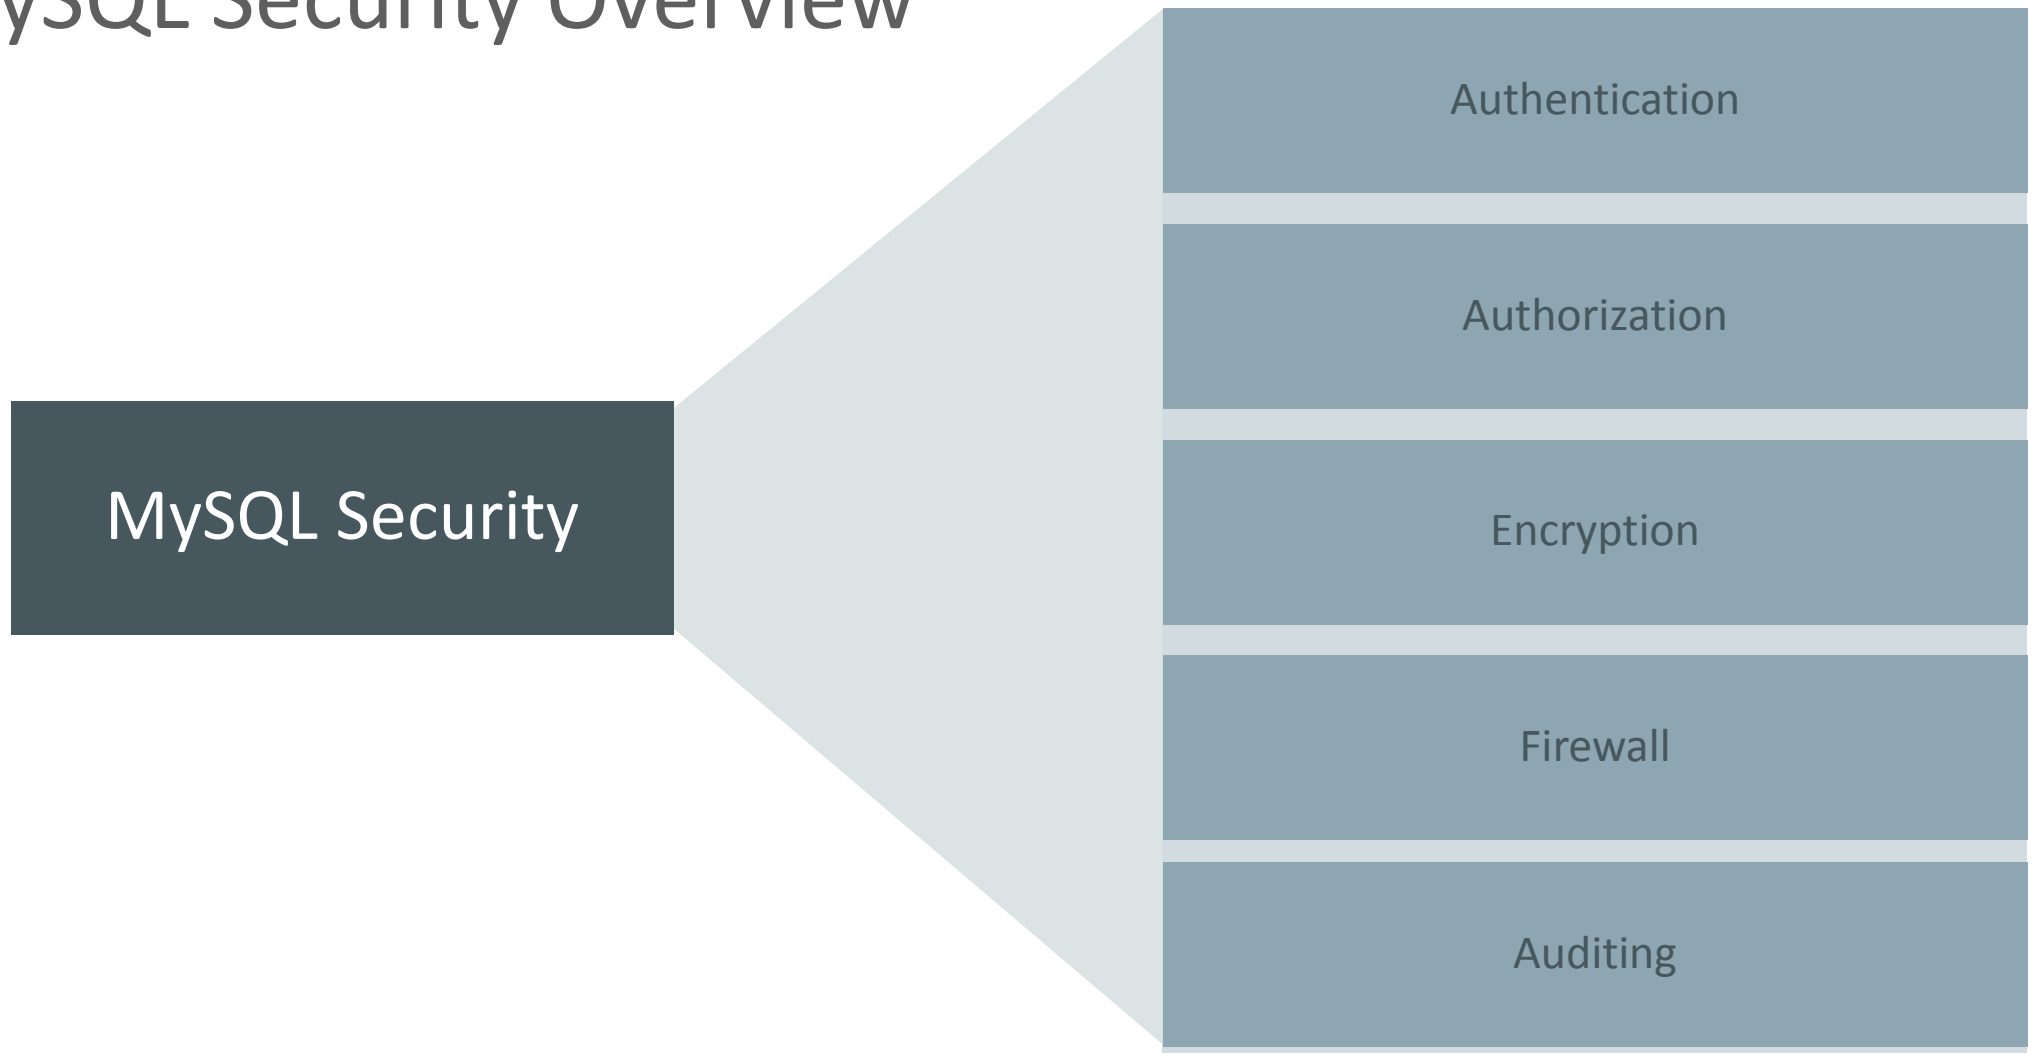

MySQL Enterprise Authentication

Integrates MySQL with existing security infrastructures

- Integrate with Centralized Authentication Infrastructure
  - Centralized Account Management
  - Password Policy Management
  - Groups & Roles
- PAM (Pluggable Authentication Modules)
  - Standard interface (Unix, LDAP, Kerberos, others)
  - Windows
    - Access native Windows service - Use to Authenticate users using Windows Active Directory or to a native host

MySQL Enterprise Authentication: PAM

- Standard Interface
  - LDAP
  - Unix/Linux
- Proxy Users

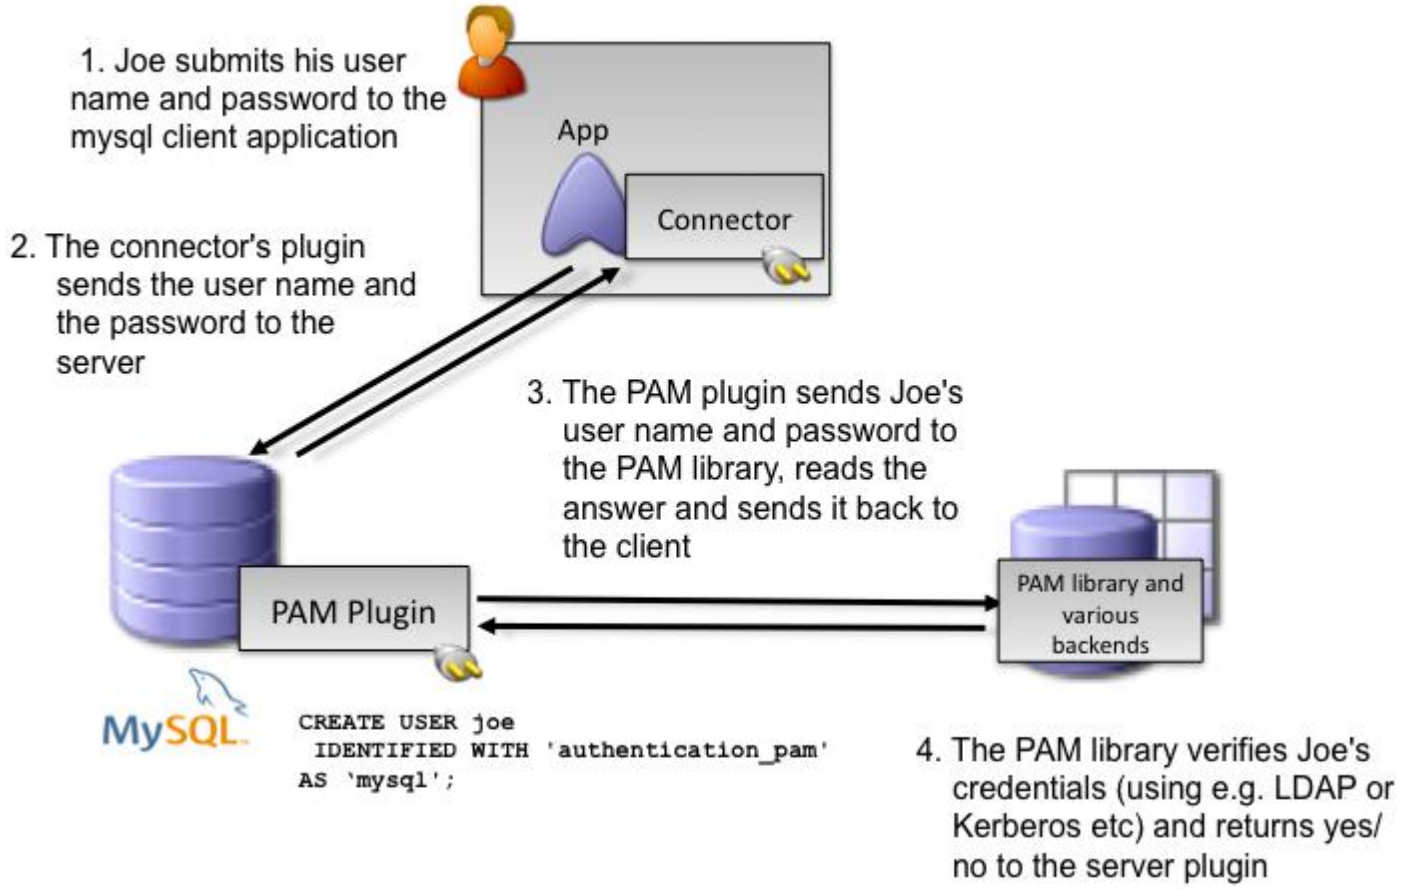

Prevent:

SQL INJECTION ATTACKS

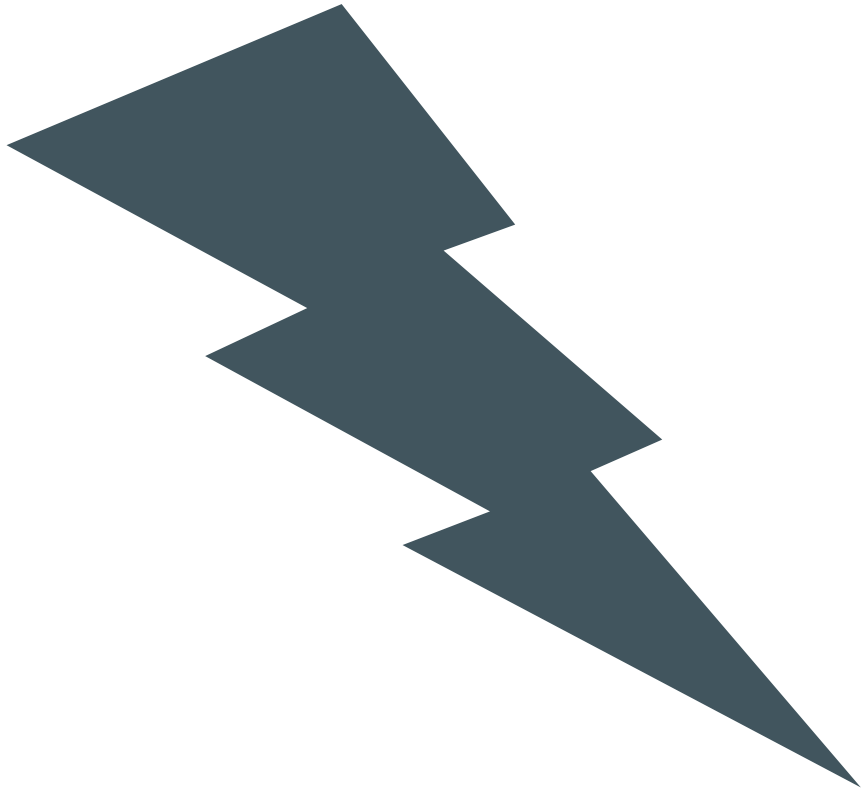

MySQL Enterprise Firewall

- Real Time Protection
  - Queries analyzed and matched against White List
- Blocks SQL Injection Attacks
  - Positive Security Model
- Block Suspicious Traffic
  - Out of Policy Transactions detected & blocked
- Learns White List
  - Automated creation of approved list of SQL command patterns on a per user basis
- Transparent
  - No changes to application required

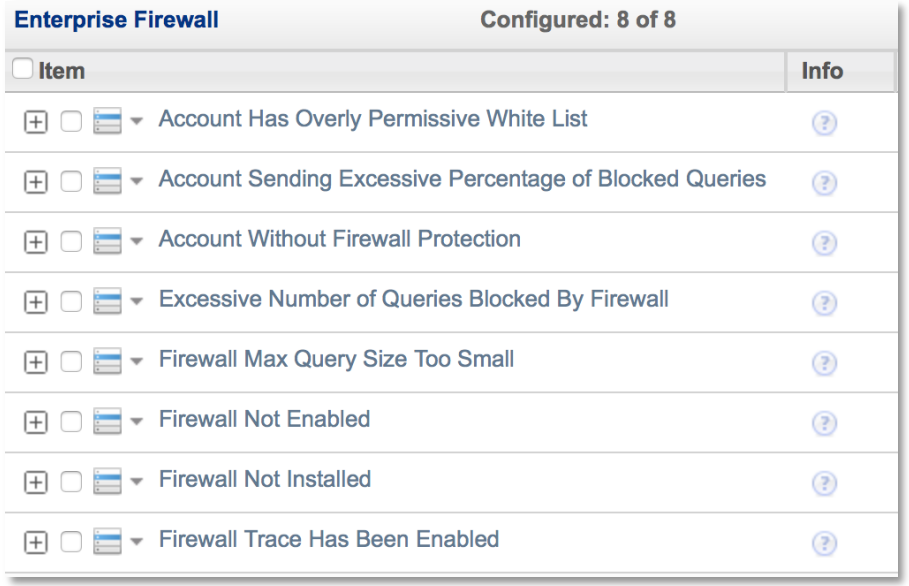

MySQL Enterprise Firewall monitoring

MySQL Enterprise Firewall

- SQL Injection Protection with Positive Security Model

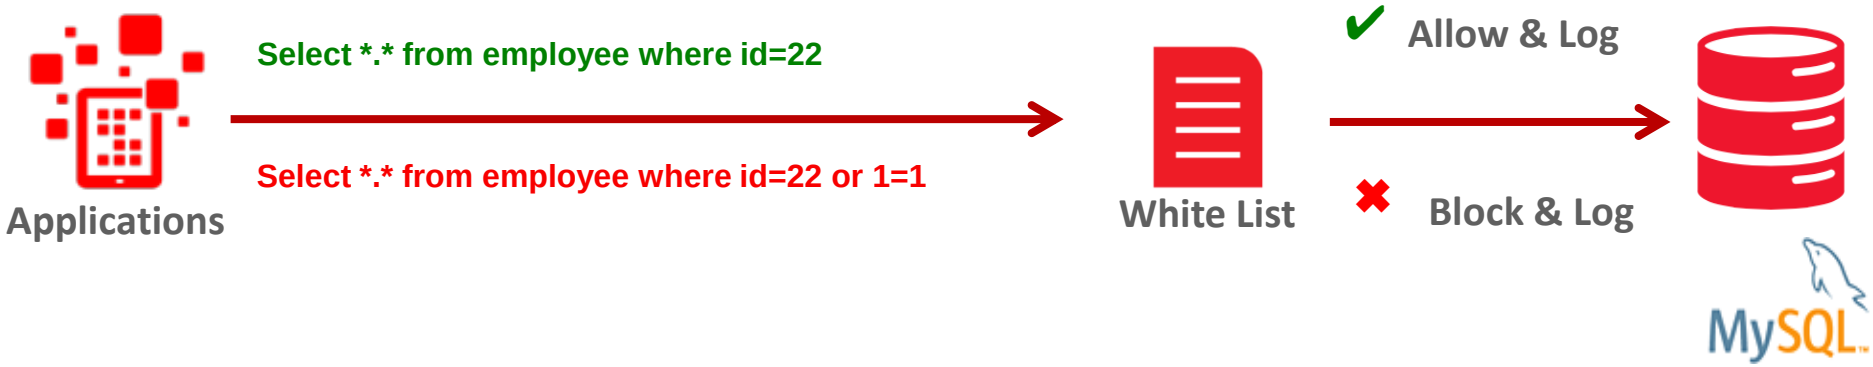

- Out of policy database transactions detected and blocked
- Logging & Analysis

Keep:

DATA SECURE AT ALL TIMES

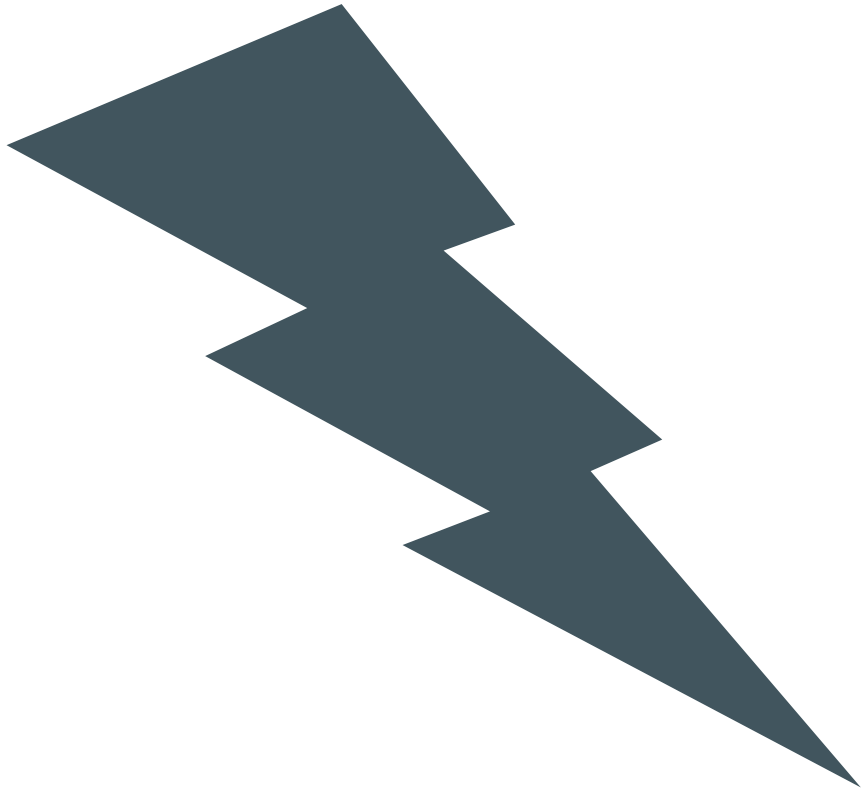

MySQL Enterprise Encryption

- MySQL encryption functions
  - Symmetric encryption AES256 (All Editions)
  - Public-key / asymmetric cryptography – RSA
- Key management functions
  - Generate public and private keys
  - Key exchange methods: DH
- Sign and verify data functions
  - Cryptographic hashing for digital signing, verification, & validation – RSA,DSA

All within MySQL

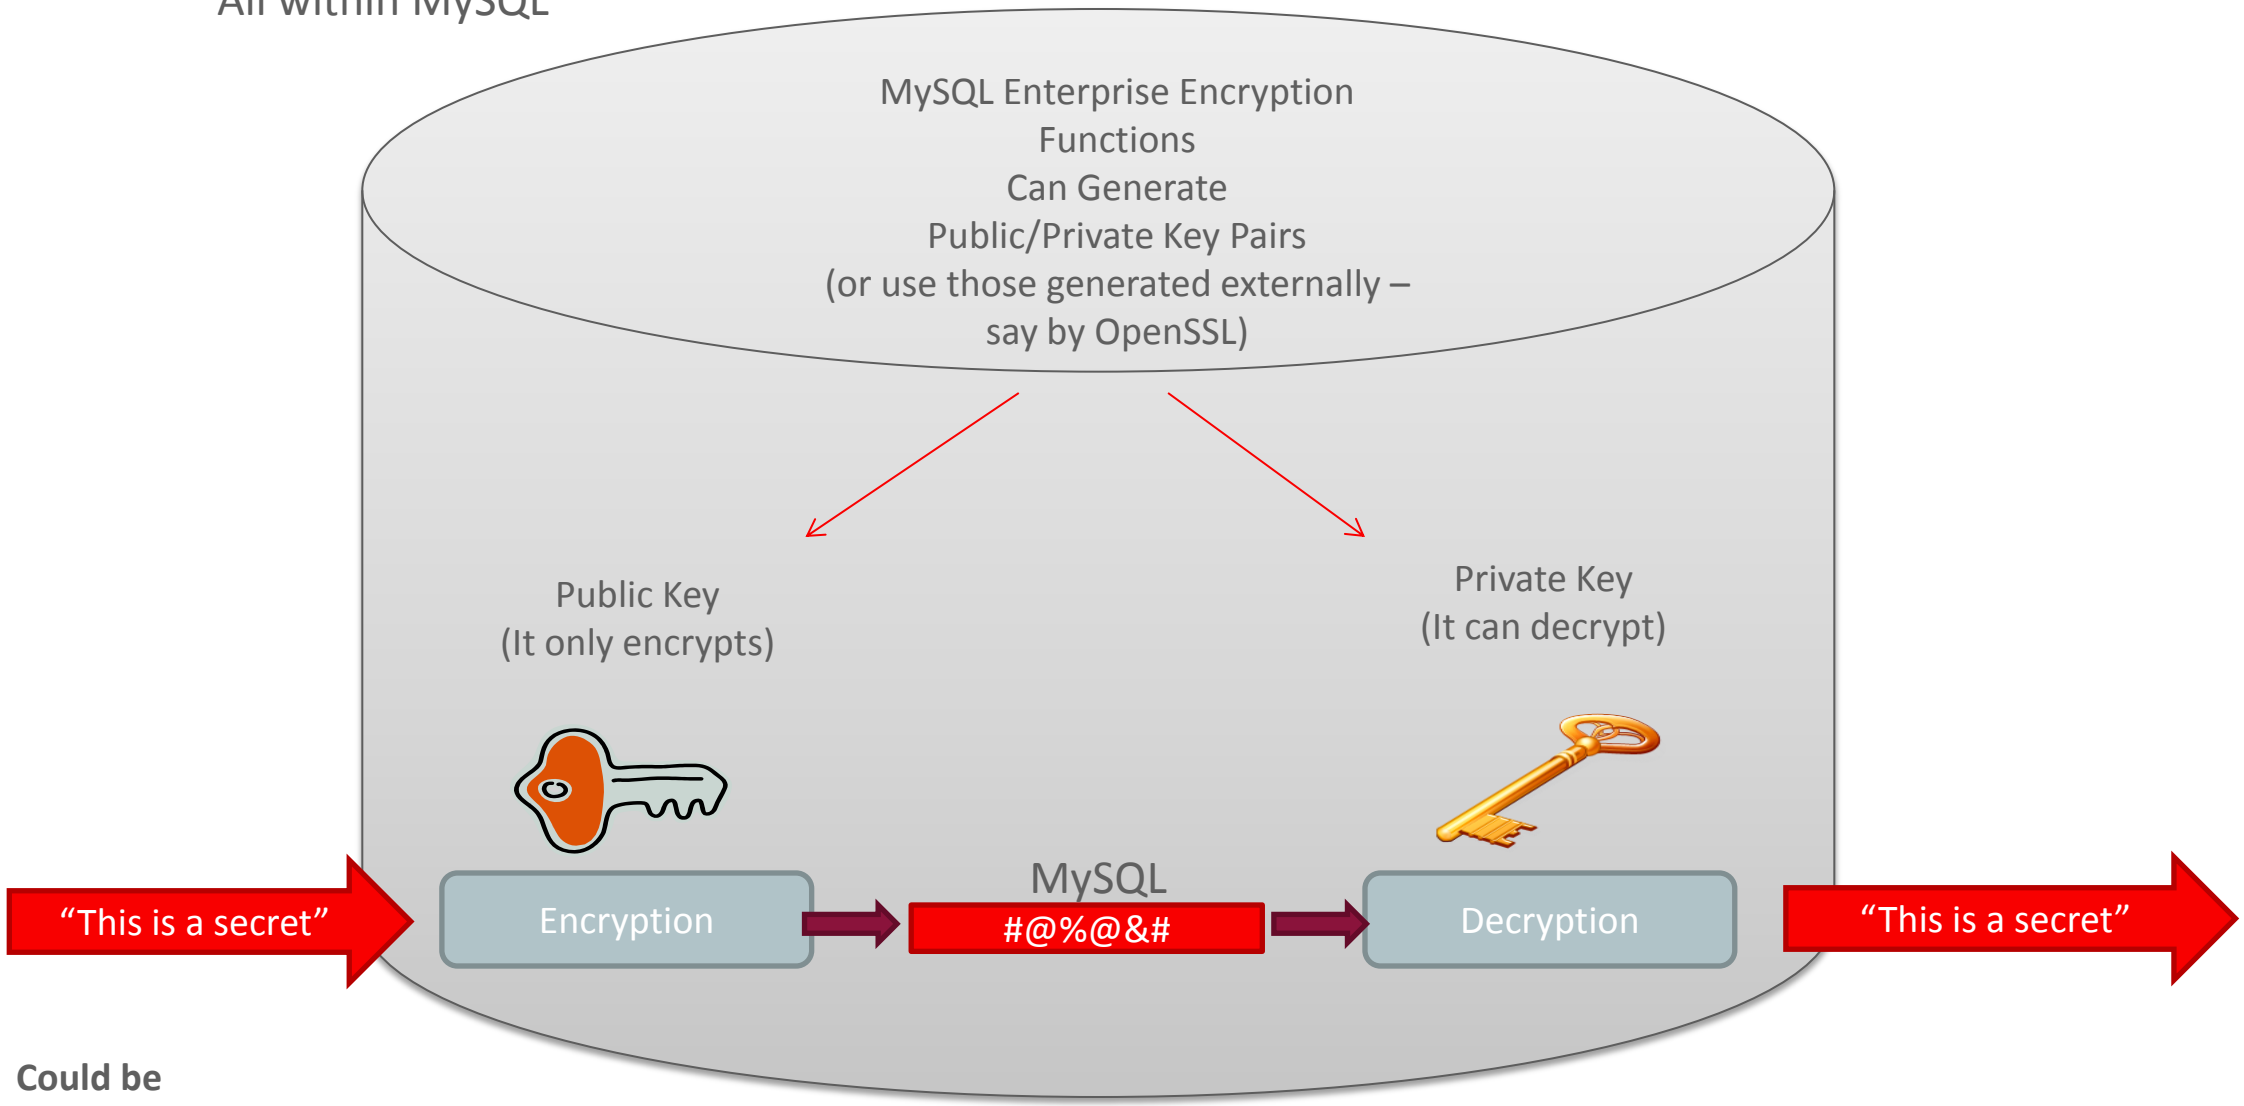

Could be
From Client App
Within MySQL (function call)

App encrypts
MySQL Stores Decrypts

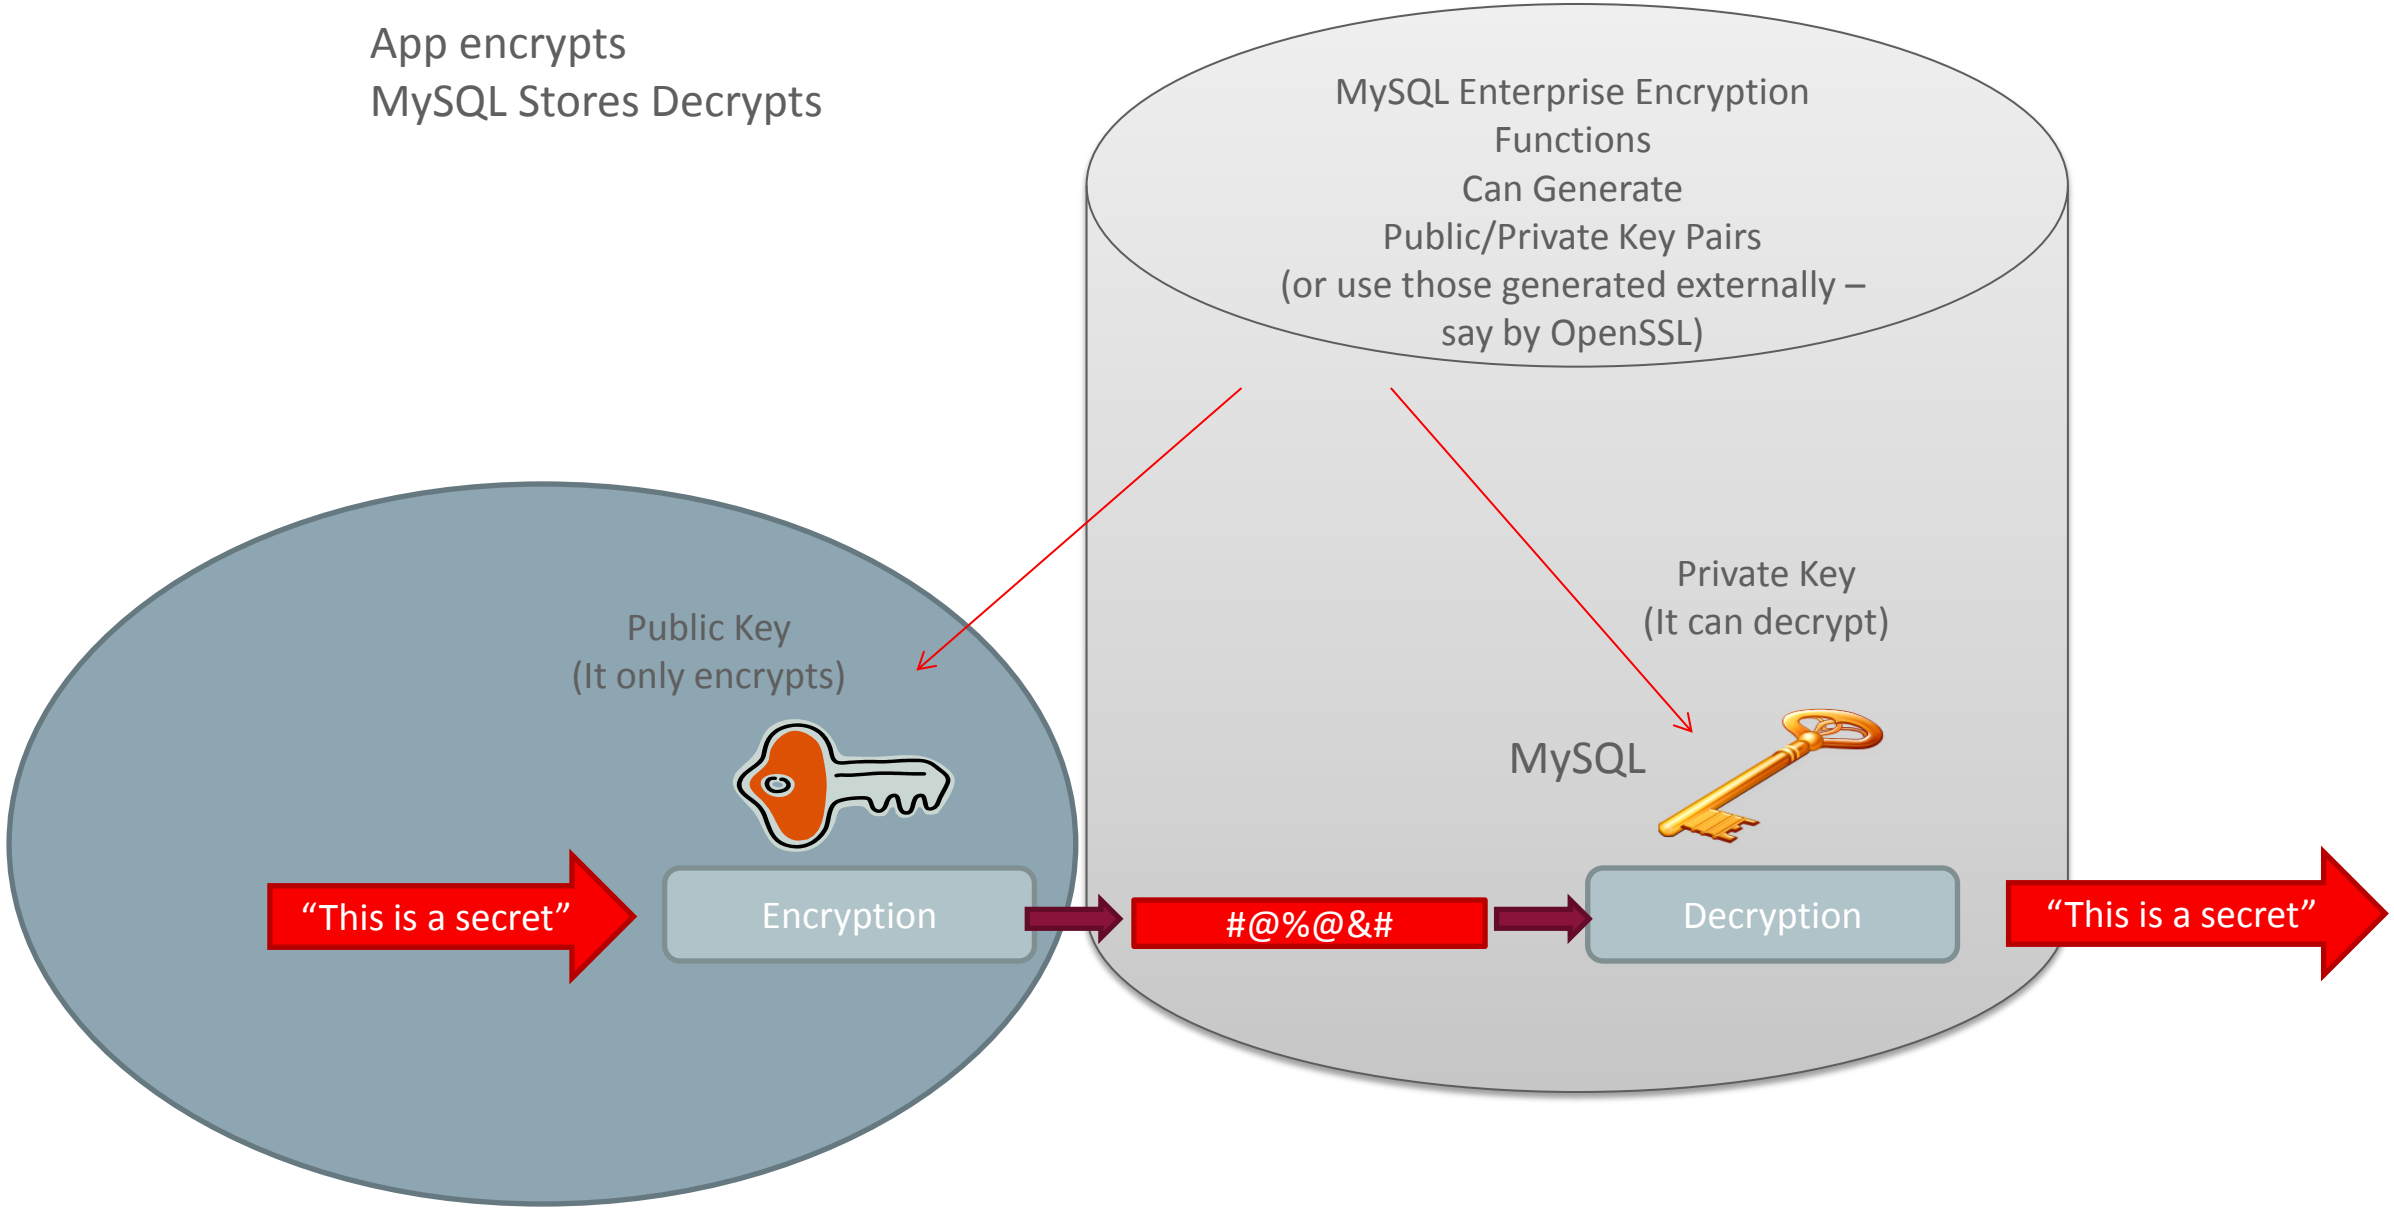

App encrypts
MySQL Stores
App Decrypts

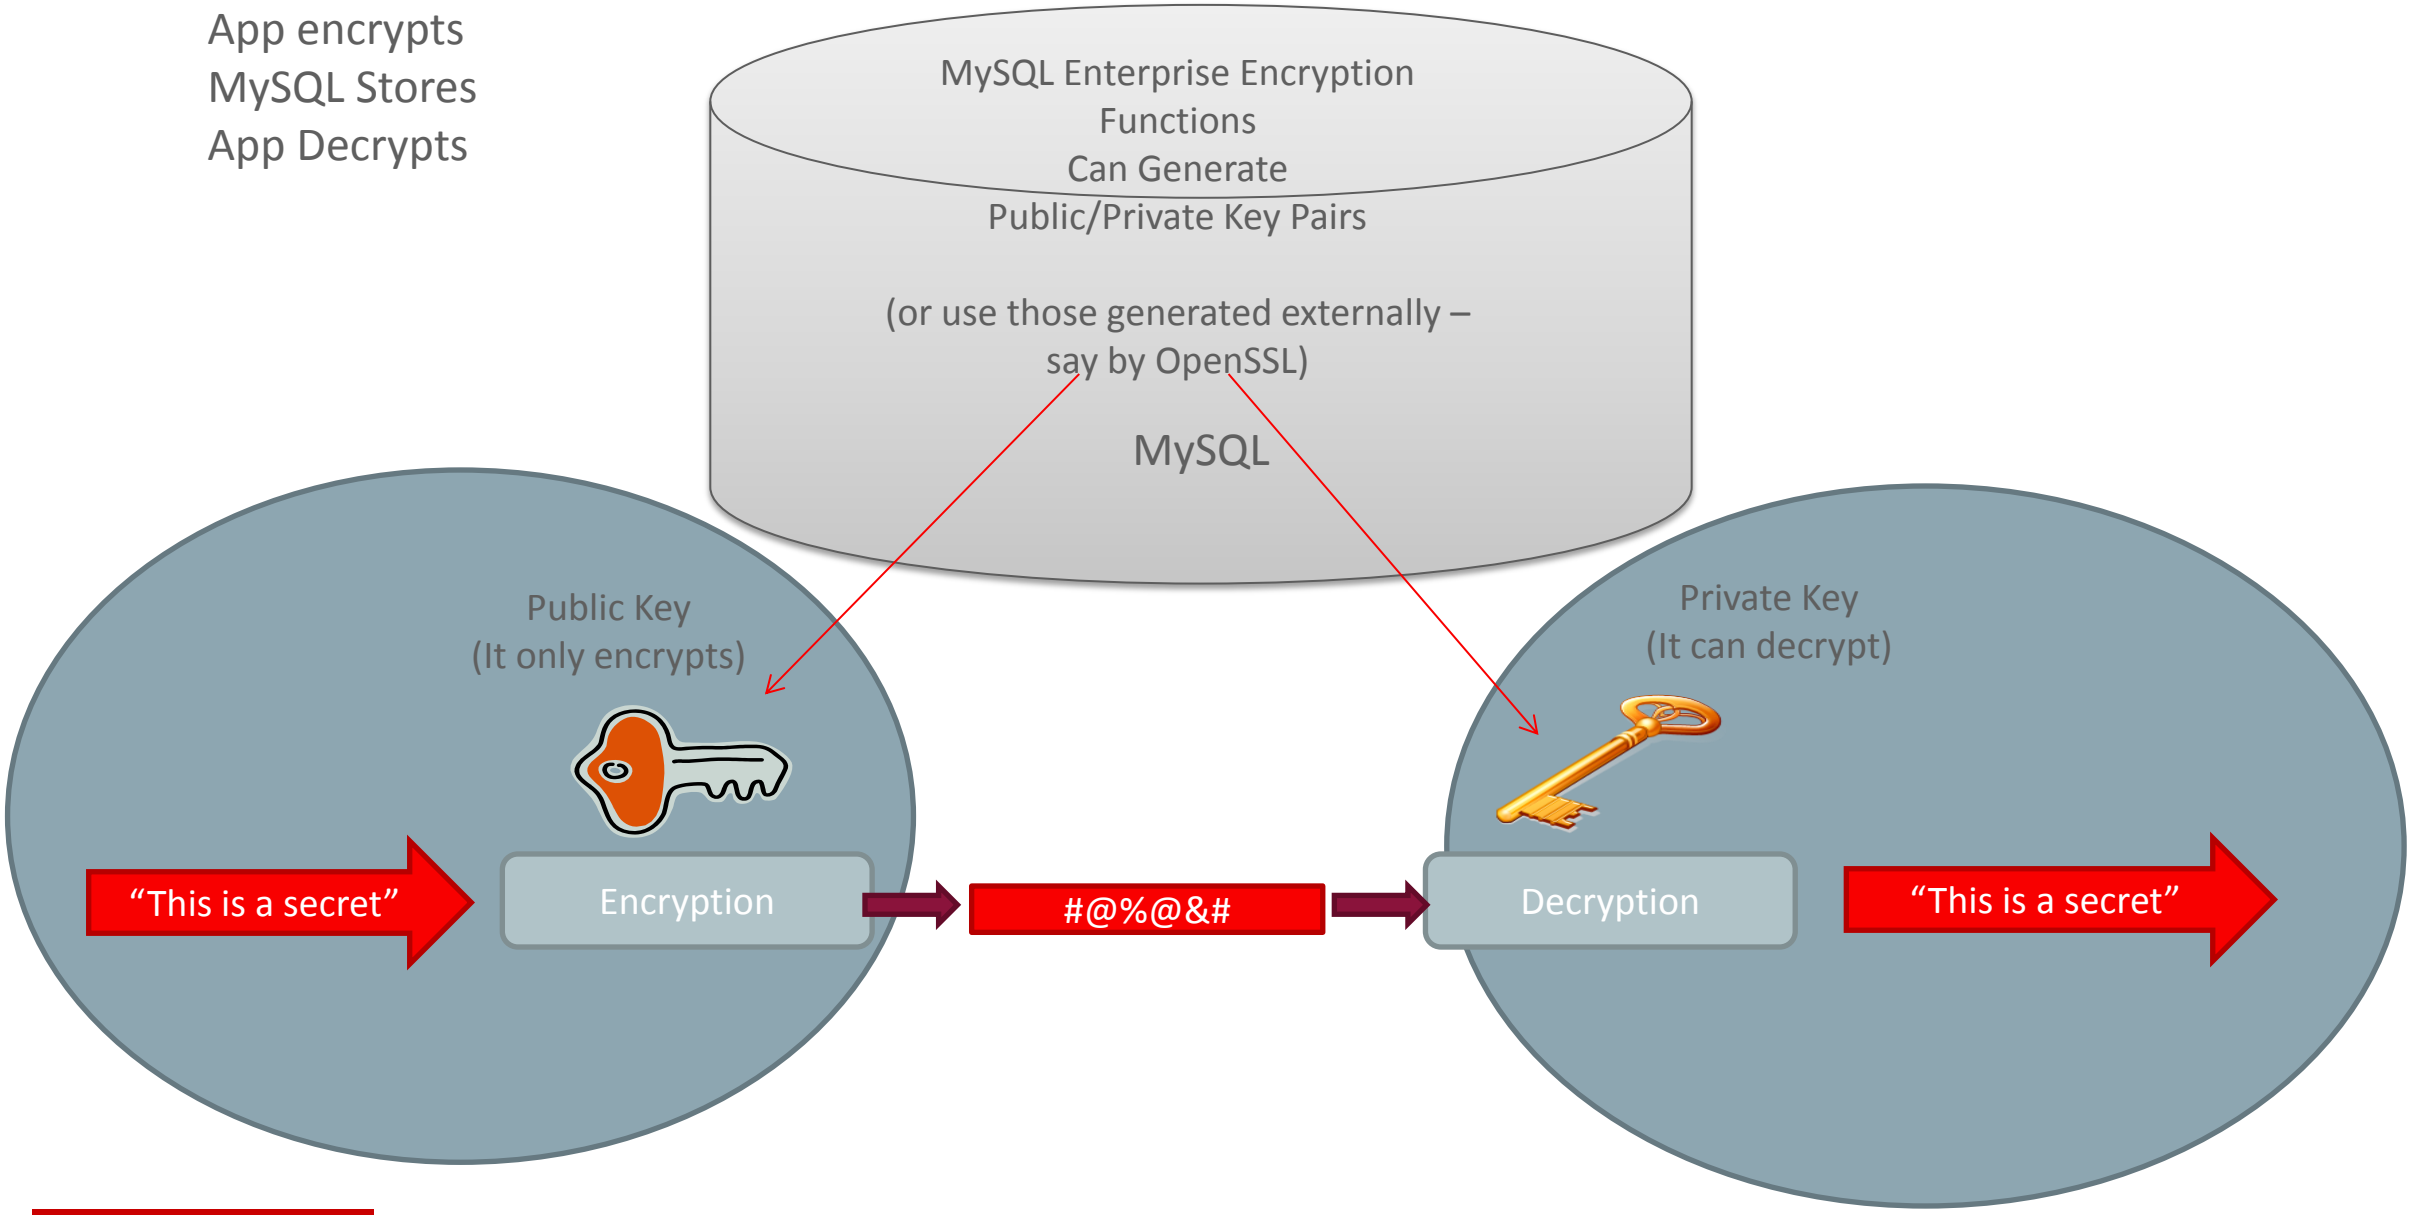

Oracle (or other) Key Vault Generates Keys
App Encrypts (only has public Key)
MySQL Stores Decrypts

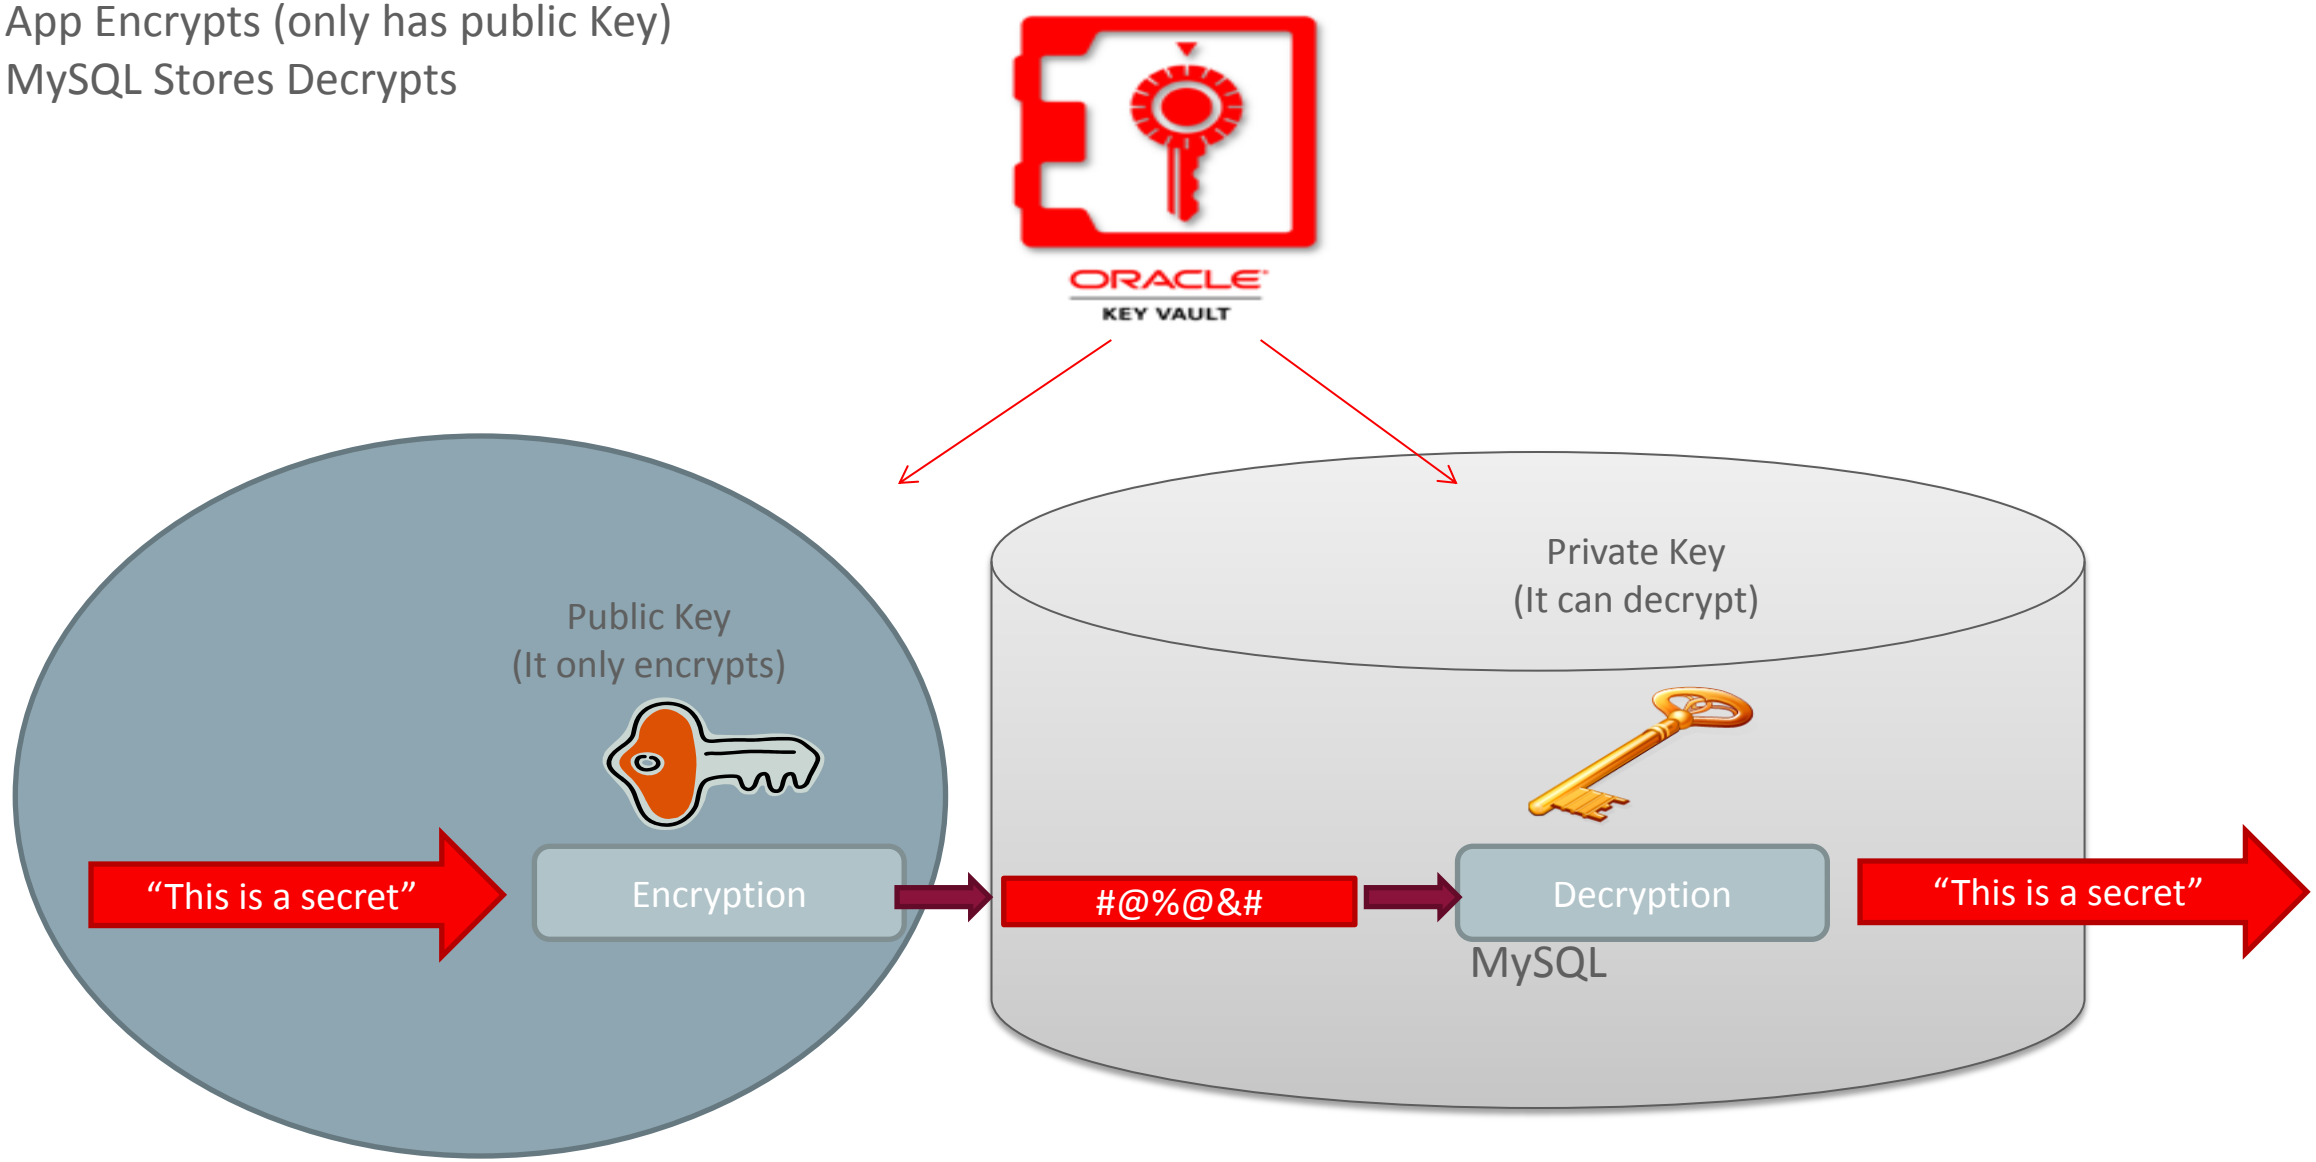

MySQL Enterprise Backup 4.0

New & Improved

- Online, non-locking backup and recovery
  - Complete MySQL instance backup (data and config)
  - Partial backup and restore
- Direct Cloud storage backups via Swift API
- Incremental backups & Point-in-time recovery
- Advanced compressed and encryption
- Optimistic backups
- Support for MySQL 5.7
  - General Tablespaces
- Improved SBT Backups to MMS Systems

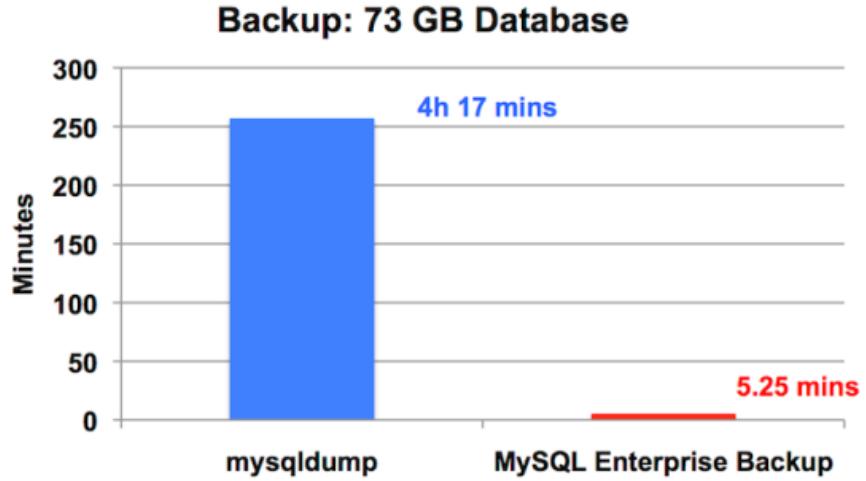

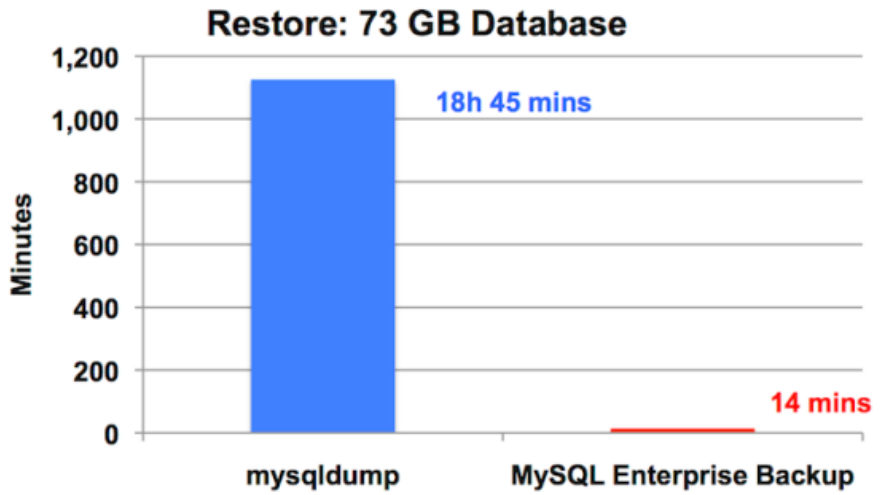

Logging For Audit

- Proper logging is always a requirement for security.
  - FIPS, HIPAA, PCI-DSS, SOX, DISA STIG, …
- MySQL built-in logging infrastructure:
  - general log, error log, NDB logs.
- MySQL Enterprise Audit plugin:
  - Granularity made for auditing.
  - Can be modified live.
  - Contains additional details not found in other logs.
  - Compatible with Oracle Audit Vault.

React:

KNOW IMMEDIATELY

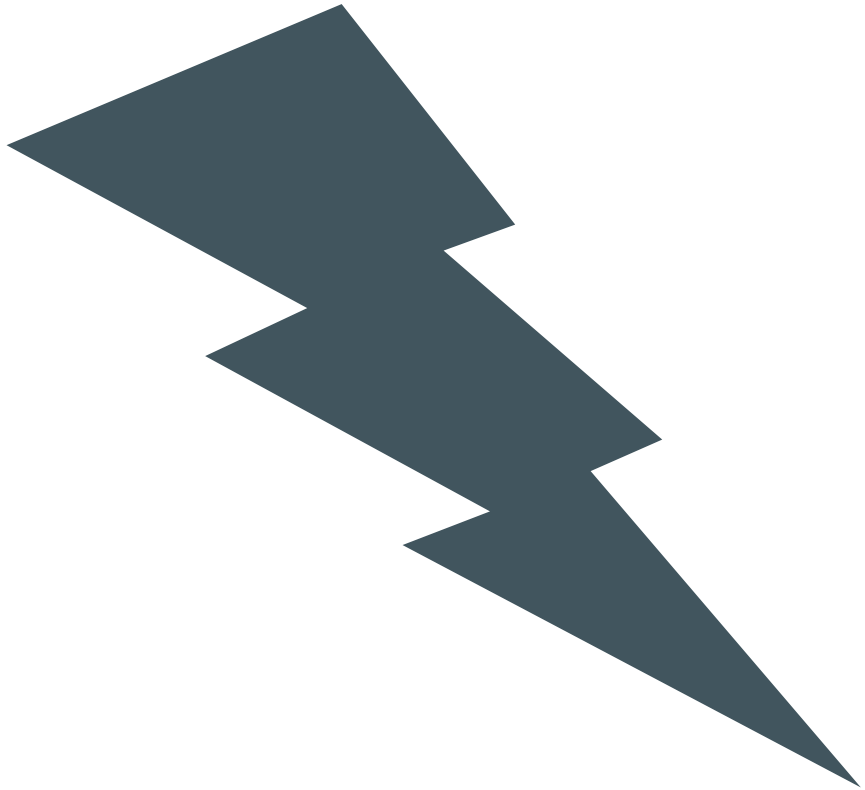

MySQL Enterprise Monitor

- Includes over 250 Best Practices Advisors.
- Many of those Identify 30+ Security Issues and Risks.
- Alerts DBA of security changes.
- Includes Query Analyzer.
- Available for Oracle EM.

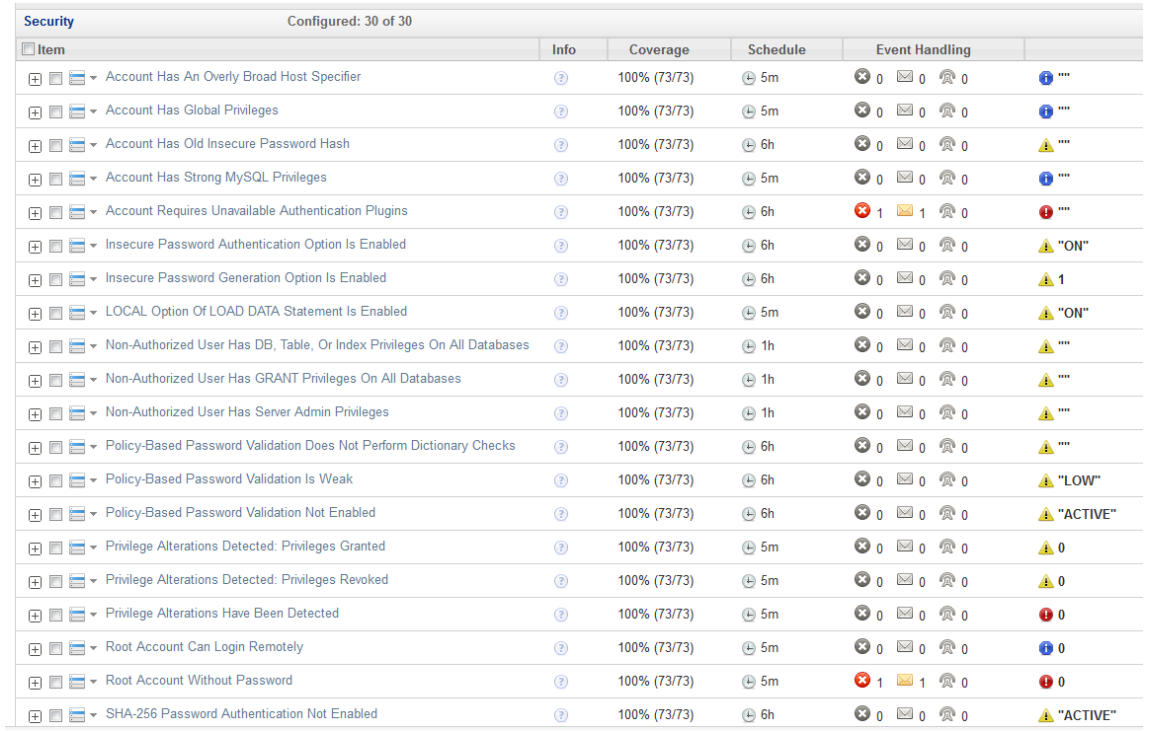

Security	Configured: 30 of 30				
Item	Info	Coverage	Schedule	Event Handling	
Account Has An Overly Broad Host Specifier		100% (73/73)	5m	0 0 0	""
Account Has Global Privileges		100% (73/73)	5m	0 0 0	""
Account Has Old Insecure Password Hash		100% (73/73)	6h	0 0 0	""
Account Has Strong MySQL Privileges		100% (73/73)	5m	0 0 0	""
Account Requires Unavailable Authentication Plugins		100% (73/73)	6h	1 1 0	""
Insecure Password Authentication Option Is Enabled		100% (73/73)	6h	0 0 0	"ON"
Insecure Password Generation Option Is Enabled		100% (73/73)	6h	0 0 0	1
LOCAL Option Of LOAD DATA Statement Is Enabled		100% (73/73)	5m	0 0 0	"ON"
Non-Authorized User Has DB, Table, Or Index Privileges On All Databases		100% (73/73)	1h	0 0 0	""
Non-Authorized User Has GRANT Privileges On All Databases		100% (73/73)	1h	0 0 0	""
Non-Authorized User Has Server Admin Privileges		100% (73/73)	1h	0 0 0	""
Policy-Based Password Validation Does Not Perform Dictionary Checks		100% (73/73)	6h	0 0 0	""
Policy-Based Password Validation Is Weak		100% (73/73)	6h	0 0 0	"LOW"
Policy-Based Password Validation Not Enabled		100% (73/73)	6h	0 0 0	"ACTIVE"
Privilege Alterations Detected: Privileges Granted		100% (73/73)	5m	0 0 0	0
Privilege Alterations Detected: Privileges Revoked		100% (73/73)	5m	0 0 0	0
Privilege Alterations Have Been Detected		100% (73/73)	5m	0 0 0	0
Root Account Can Login Remotely		100% (73/73)	5m	0 0 0	0
Root Account Without Password		100% (73/73)	5m	1 1 0	0
SHA-256 Password Authentication Not Enabled		100% (73/73)	6h	0 0 0	"ACTIVE"

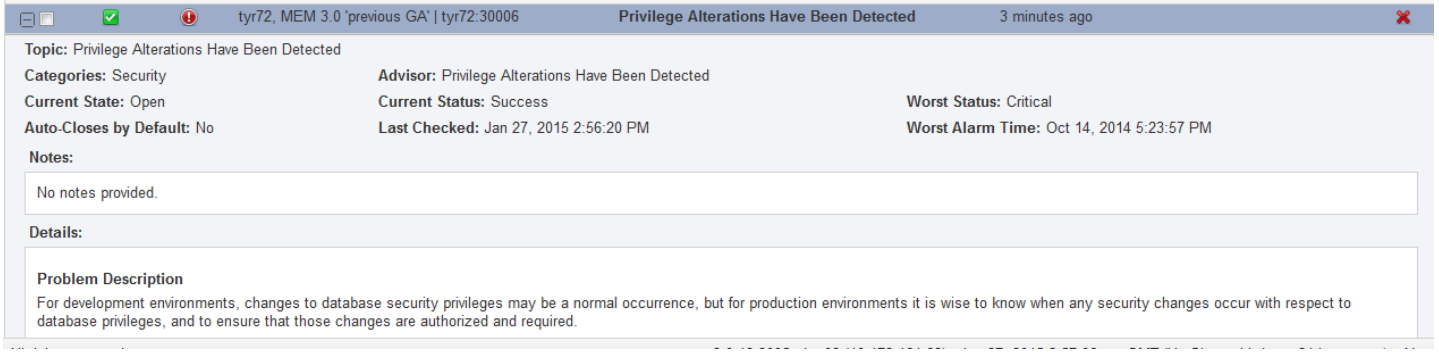

tyr72, MEM 3.0 'previous GA' | tyr72:30006 **Privilege Alterations Have Been Detected** 3 minutes ago

Topic: Privilege Alterations Have Been Detected

Categories: Security **Advisor:** Privilege Alterations Have Been Detected

Current State: Open **Current Status:** Success **Worst Status:** Critical

Auto-Closes by Default: No **Last Checked:** Jan 27, 2015 2:56:20 PM **Worst Alarm Time:** Oct 14, 2014 5:23:57 PM

Notes:

No notes provided.

Details:

Problem Description

For development environments, changes to database security privileges may be a normal occurrence, but for production environments it is wise to know when any security changes occur with respect to database privileges, and to ensure that those changes are authorized and required.

MySQL Enterprise Monitor 3.1

Security

- Easily ensure all your MySQL assets are hardened and secure.
- Monitor MySQL Enterprise Firewall:
  - Detect SQL injection attacks and other common threats.
- Monitor MySQL Enterprise Audit:
  - Ensure regulatory compliance.
  - Know what happened when things go wrong.
- Change monitoring and tracking.
- Backup policy enforcement.

ORACLE

A Fully Integrated Secure System

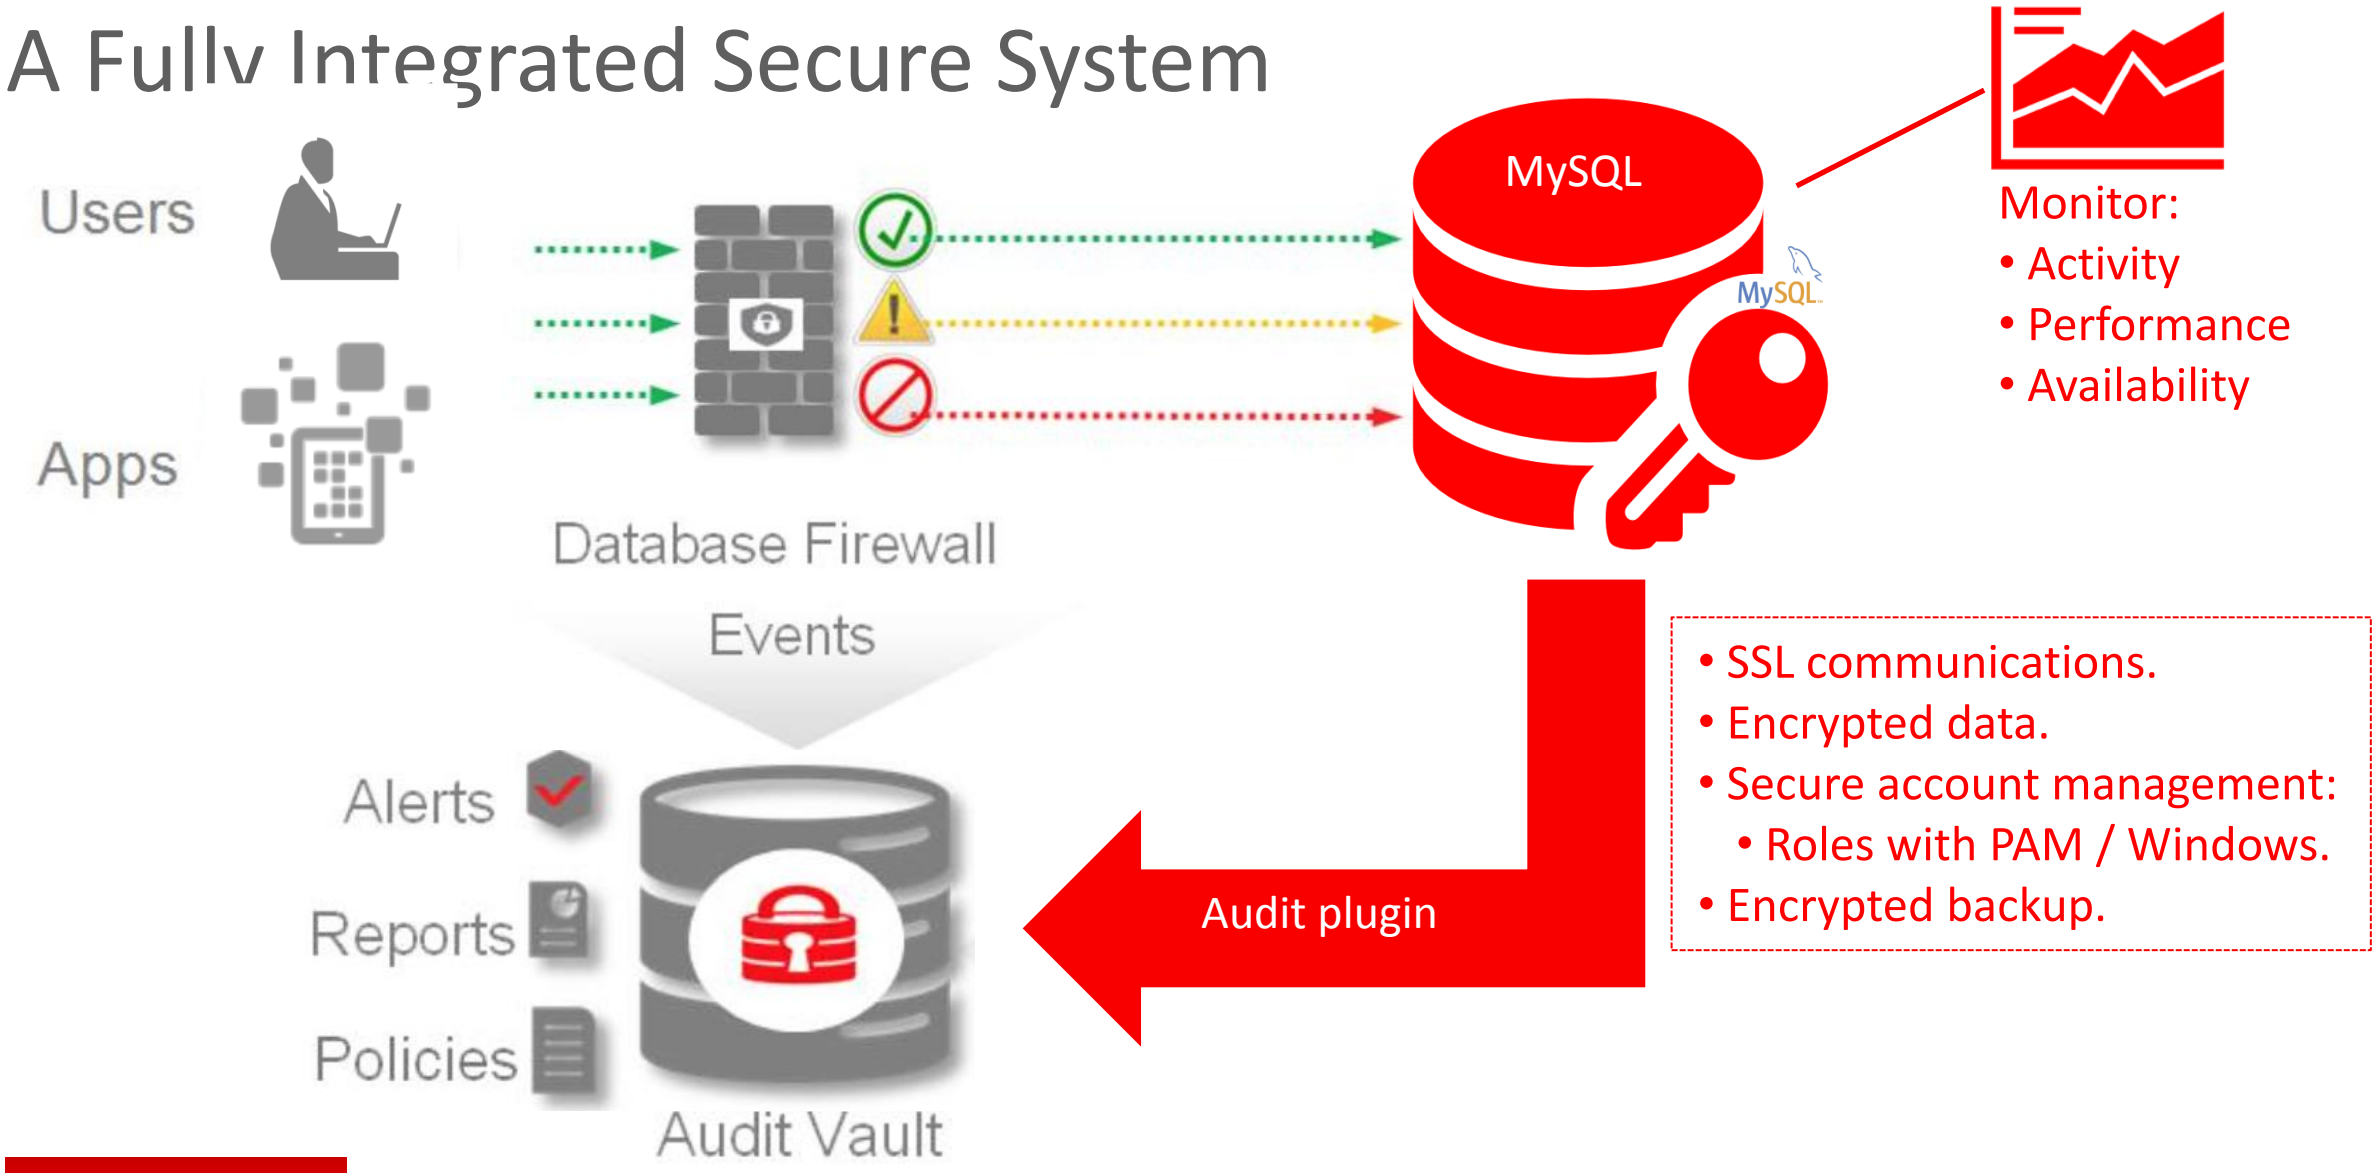


ORACLE®