

ORACLE®



ORACLE®

MySQL High Availability Solutions

Luca Olivari

MySQL Sales Consulting Manager, EMEA

Agenda

- Introduction to High Availability
- MySQL High Availability Solutions
- MySQL Cluster CGE
- Operational Best Practices
- Conclusion and Questions



Why do I need High Availability?



- In the old days
 - Local time
 - Monday thru Friday
 - 8 am to 6 pm



- Current environment
 - Global
 - 365 x 24 x 7
 - No Downtime

Eyjafjallajökull

http://en.wikipedia.org/wiki/2010_eruptions_of_Eyjafjallaj%C3%B6kull

Why do we need to provide High Availability?


Remembering the ash cloud...

- Planned Downtime
 - Database backup/upgrade/patching
 - Operating system upgrade/patching
 - Hardware and Network maintenance
- Unplanned Downtime
 - Corruptions
 - Logical corruptions
 - Physical corruption
 - Human Errors
 - Accidentally drops a table ...
 - Disasters
 - Volcanoes, terrorism, earthquakes ...



High Availability from Customers' Point of View

Hourly Costs of Downtime per Industry



Sources: IT Performance Engineering and Measurement Strategies: Quantifying Performance and Loss, Meta Group, Oct. 2000;

ORACLE

Reputation Matters

September 23, 2010 – Facebook is Down!



@alqaeda
Al-Qaeda

#facebook is down. Not sure if we did that, but we should claim credit anyway. Hitting the infidels where it hurts, etc.

19 hours ago via Tweetie for Mac ☆ Favorite ↺ Retweet ↻ Reply



@VladimirRS
Vladimir

BREAKING NEWS: Facebook down due to explosion in Farmville's manure composter. Five million people forced to work.



@RWW
ReadWriteWeb


Even Like buttons are down on Facebook. We can't live in a world where we can't arbitrarily like things. The end is nigh.
#FacebookDown

7 hours ago via TweetDeck ☆ Favorite ↺ Retweet ↻ Reply

Not Everything Needs 99.999% Uptime

- As increased availability needs more resources, be sure to adopt a tiered approach:
 - Tier 1, Mission-critical services
 - Tier 2, Business-critical services
 - Tier 3, Task-critical services
 - Tier 4, Non-critical services

Mapping Architectures to Availability



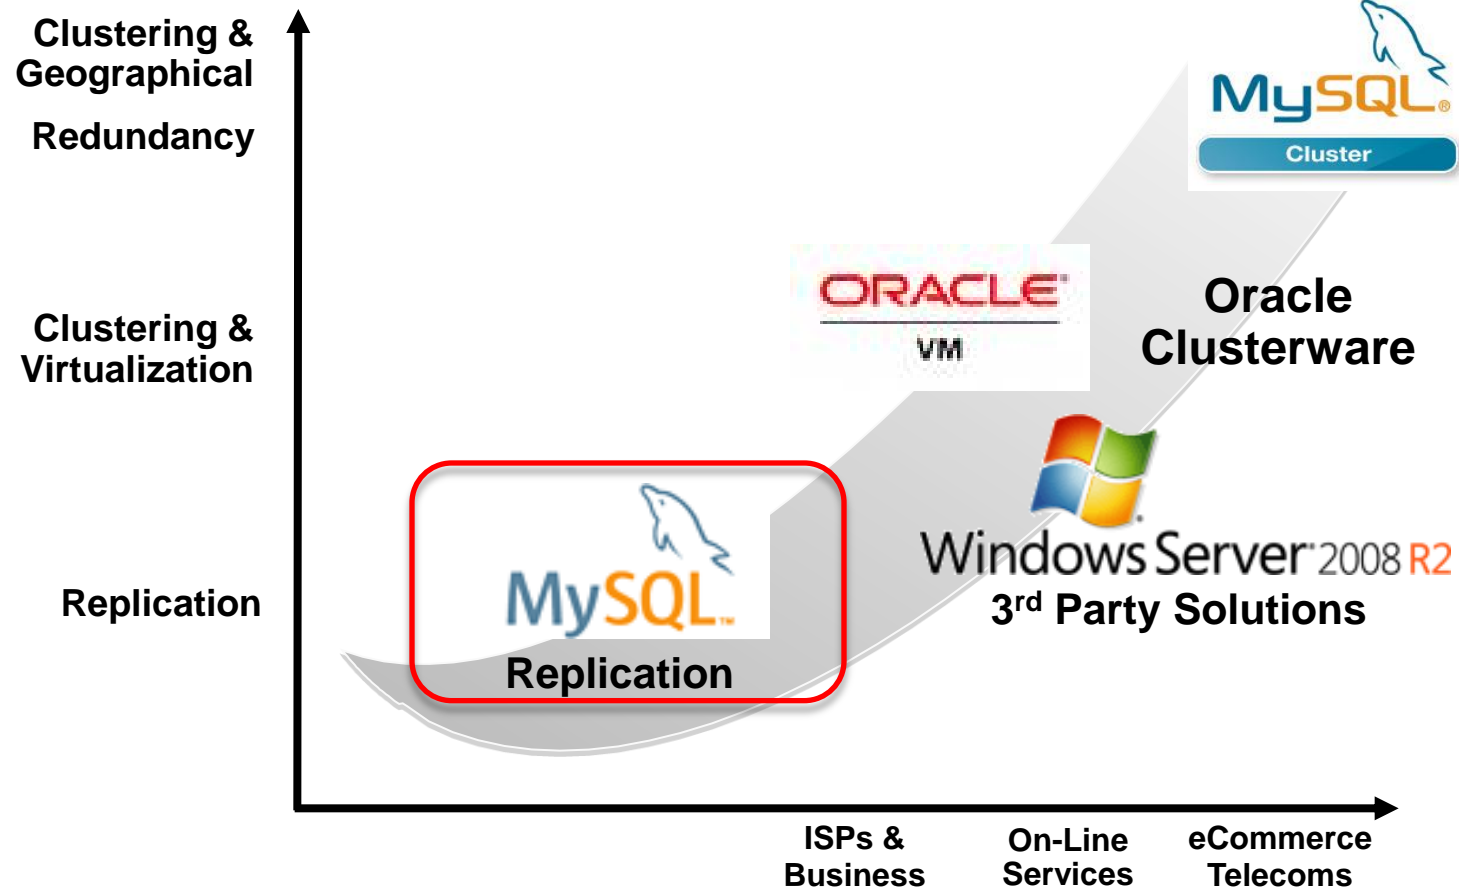
| | | | | | |
|--------------|-------------|--------------|--------------|-------------|---|
| 9 35 days | 9 4 days | 9 8 hours | 9 50 mins | 9 5 mins | % |
|--------------|-------------|--------------|--------------|-------------|---|

Agenda

- Introduction to High Availability
- MySQL High Availability Solutions
- MySQL Cluster CGE
- Operational Best Practices
- Conclusion and Questions




Mapping Architectures to Availability



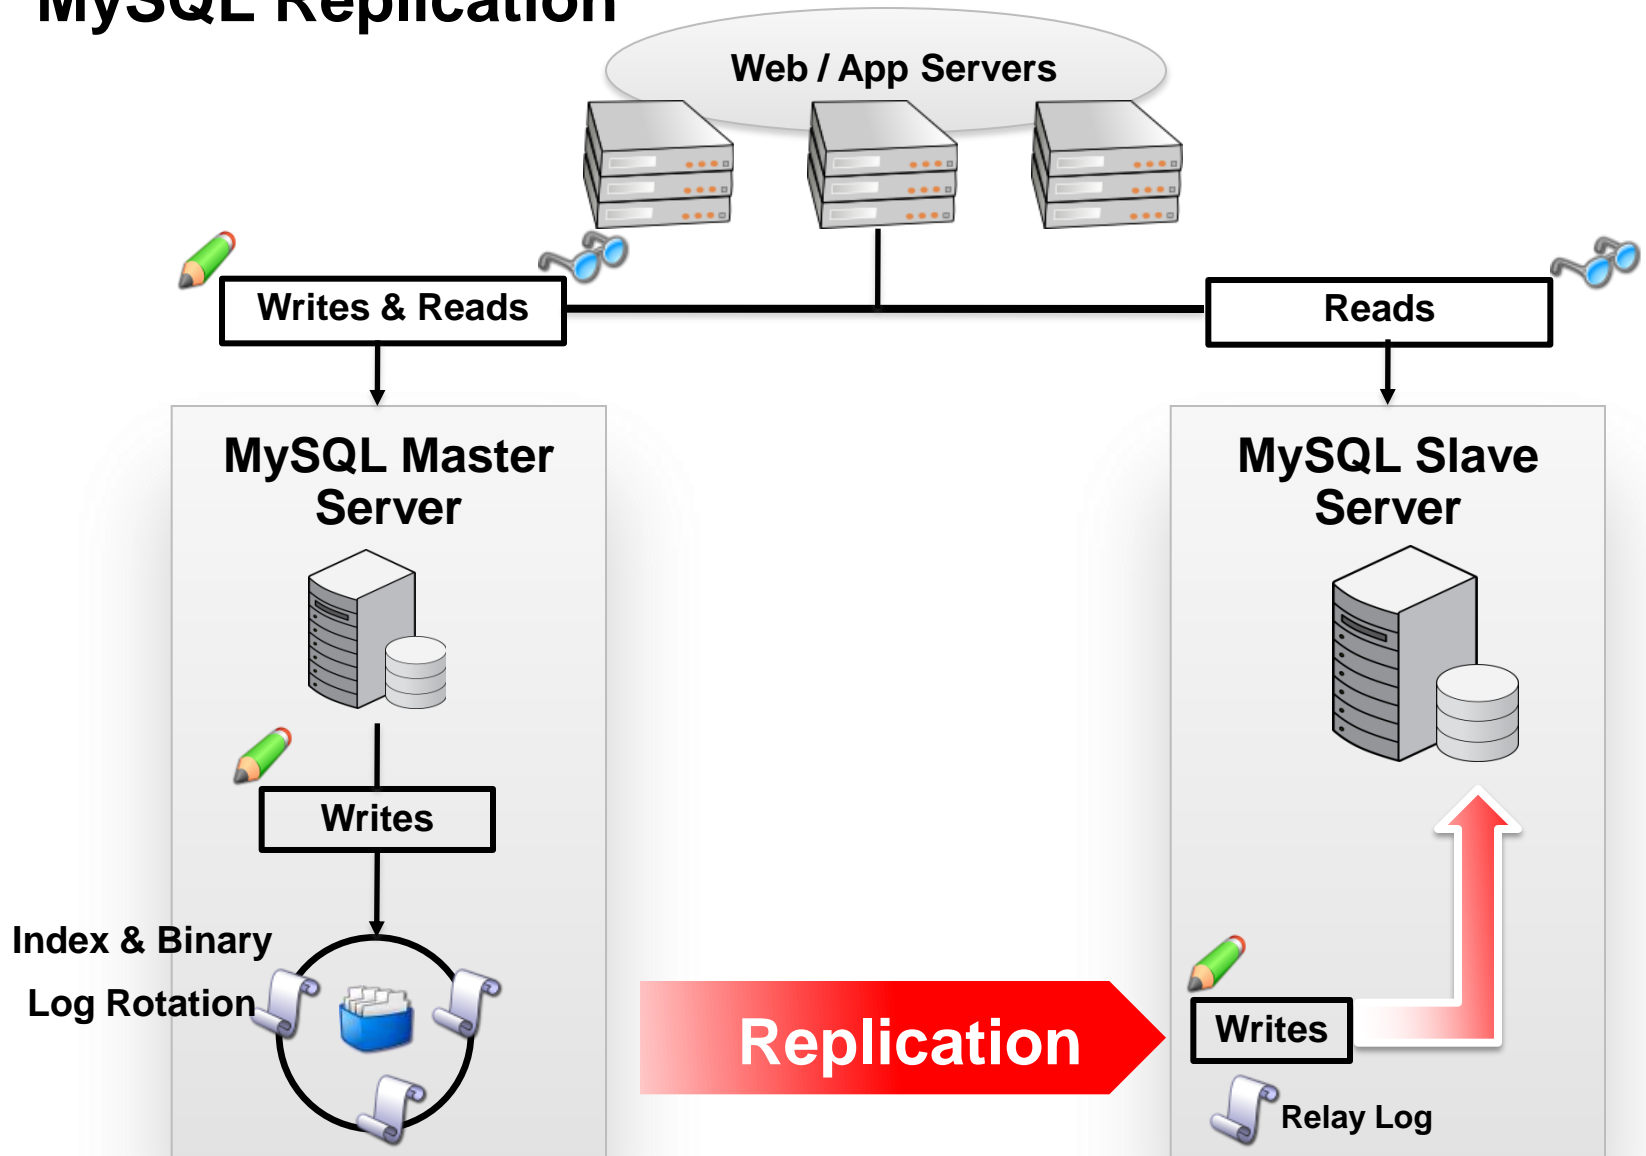
| | | | | | |
|--------------|-------------|--------------|--------------|-------------|---|
| 9 35 days | 9 4 days | 9 8 hours | 9 50 mins | 9 5 mins | % |
|--------------|-------------|--------------|--------------|-------------|---|

MySQL Replication

- Native to MySQL
- Duplicates updates from a “master” to a “slave”
- Most common solution for HA
 - Used to scale out as well
- Failover is either scripted or provided by additional middleware
- Supports Asynchronous and Semi-Synchronous replication




MySQL Replication




Replication Topologies

A Flexible Way to Scale


Single




Chain



Multiple

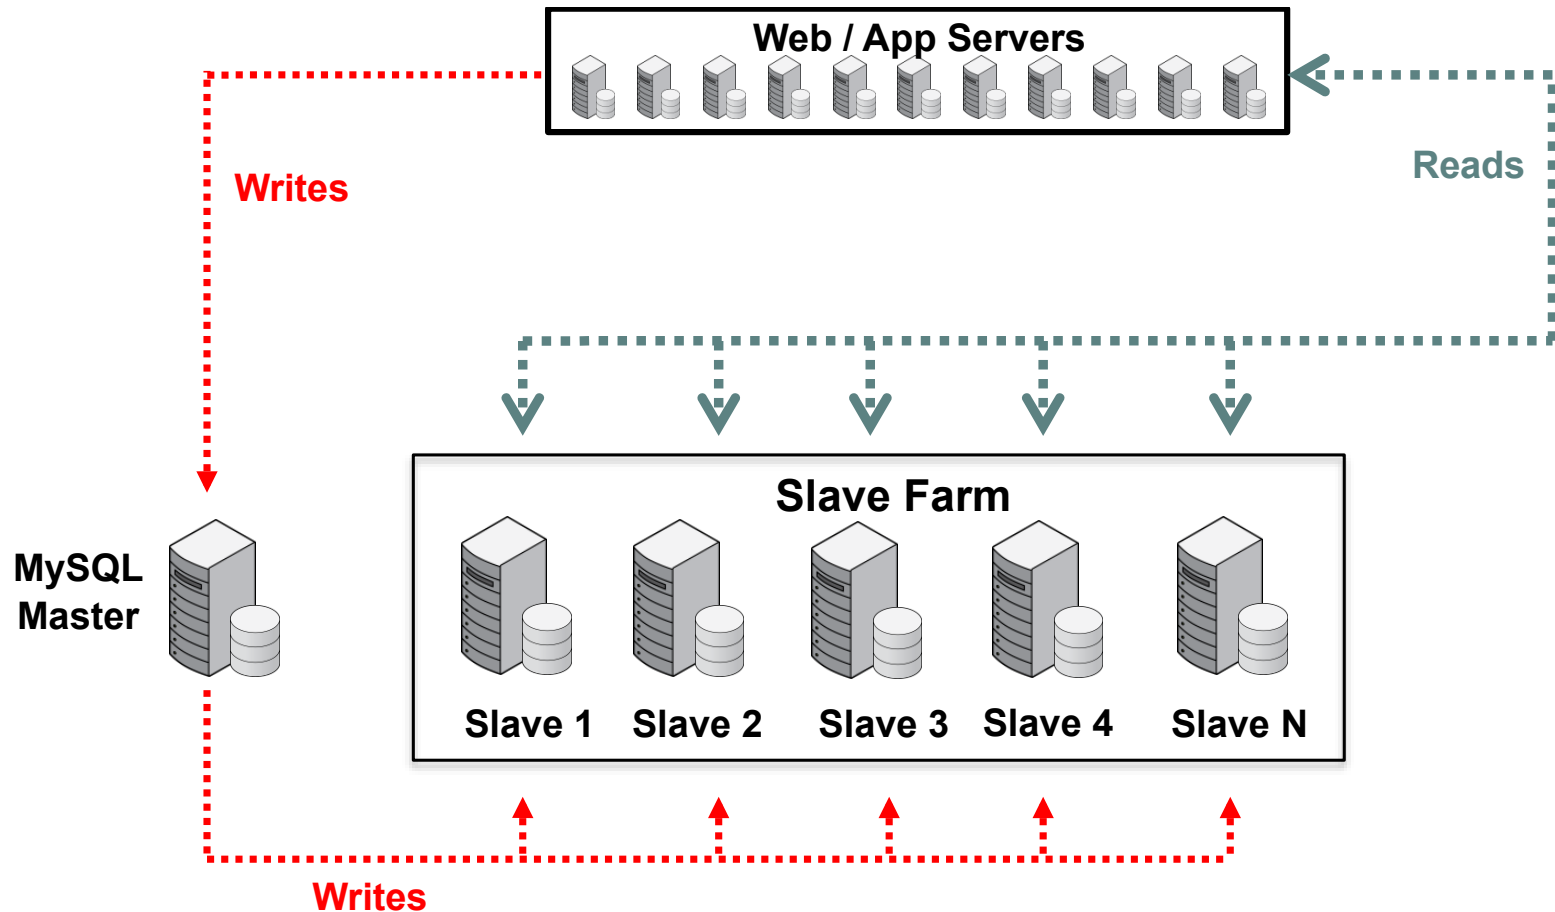


Circular



Master ●
Slave ●

Scaling out with MySQL Replication



Read more: <http://mysql.com/why-mysql/white-papers/mysql-wp-replication.php>

MySQL 5.5

InnoDB becomes default

- ACID Transactions, FKs, Crash Recovery

Improved Availability

- Semi-synchronous Replication
- Replication Heartbeat

Improved Usability

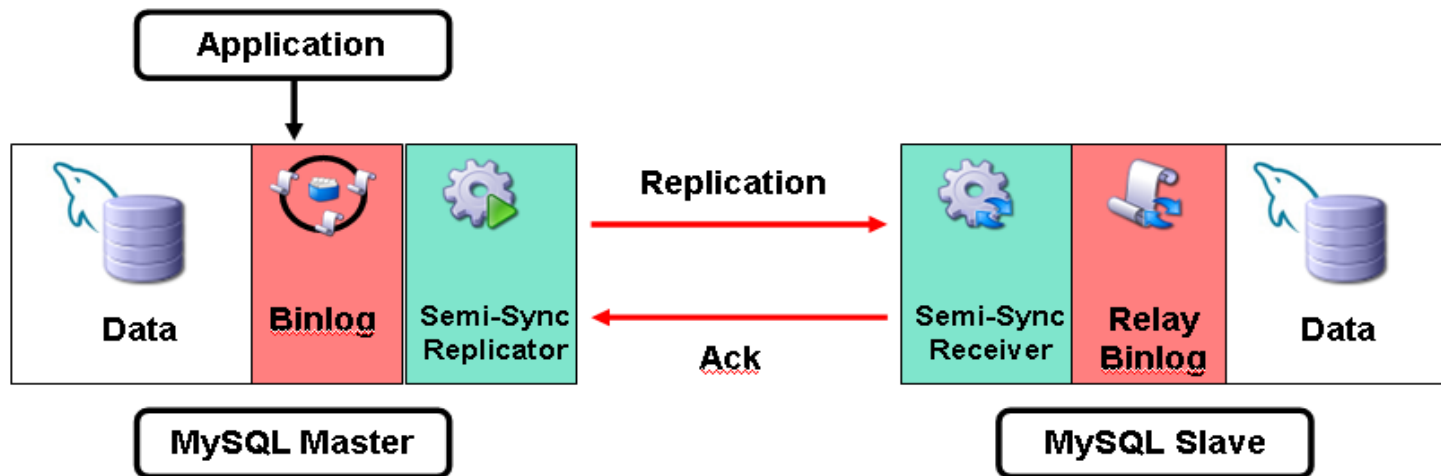
- SIGNAL/RESIGNAL
- More Partitioning Options
- PERFORMANCE_SCHEMA

Better Instrumentation/Diagnostics

- InnoDB stats in 5.5 PERFORMANCE_SCHEMA

MySQL 5.5 - Semisynchronous Replication


- Semisynchronous replication ensures redundancy
 - At least one slave acknowledges relay logging the transaction
 - Master waits for slave acknowledgement before commit returns
 - On time-out, master temporarily switches to async replication



Enhancing HA in MySQL Replication

MySQL 5.6 DMR & Early Access Labs

MySQL Masters



- Global Transaction IDs
- Crash-Safe Slaves
- Multi-Threaded Slaves
- Group Commit
- Replication Checksums
- Binlog API

Global Transaction IDs: <http://bit.ly/pqdkPZ>

General Replication Enhancements: <http://bit.ly/q3ZNga>

Automated Replication Monitor

MySQL Enterprise Monitor

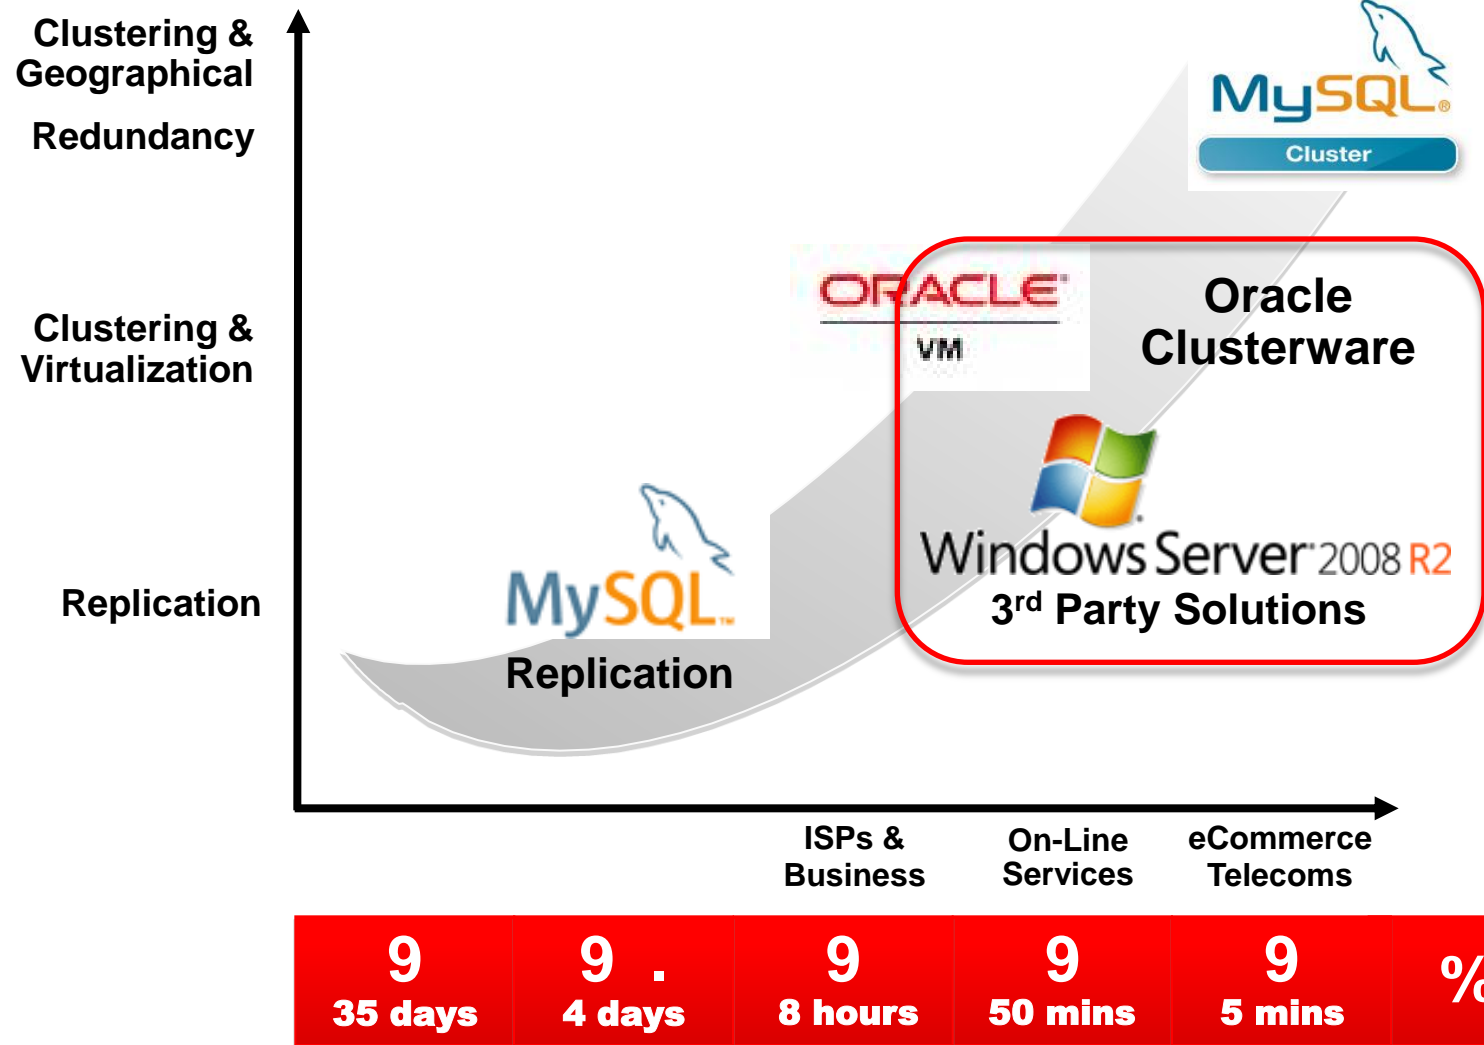
- Auto-detects, groups/maintains Master/Slave topologies
- Consolidated, real time status/synch check
- Notifications on Synch Issues
- Proactive vs reactive

Saves you time monitoring and collecting replication status/synch data from MySQL command line.



| Replication Monitoring | | | | |
|------------------------|--------------|----------|-----------|----------------|
| Servers | Type | Slave IO | Slave SQL | Seconds Behind |
| Basic (2) | TREE | Running | Running | |
| master:10101 | master | | | |
| slave:10100 | slave | Running | Running | 00:00:00 |
| Ringlet (2) | RING | Running | Running | |
| Yang:10120 | master/slave | Running | Running | 00:00:00 |
| Yin:10121 | master/slave | Running | Running | 00:00:00 |
| RingSpoke (4) | MIXED | Running | Running | |
| ring1:10183 | master/slave | Running | Running | 00:00:00 |
| ring2:10182 | master/slave | Running | Running | 00:00:00 |
| ring3:10181 | master/slave | Running | Running | 00:00:00 |
| ring3slave:10180 | slave | Running | Running | 00:00:00 |
| Tree 3 (5) | TREE | Running | Running | |
| master:10153 | master | | | |
| slave1:10150 | slave | Running | Running | 00:00:00 |
| slave2master:10152 | master/slave | Running | Running | 00:00:00 |
| slave2slave:10151 | slave | Running | Running | 00:00:00 |
| slave3:10154 | slave | Running | Running | 00:00:00 |

Mapping Architectures to Availability




Clustering Overview

- Tightly coupled clusters of servers/storage providing service to an application
 - Use heartbeating and management software to monitor hardware, OS, database and network
 - Heartbeat between nodes detects failures and automatically failover to redundant systems
 - Implements Virtual IP (VIP) so failover is transparent to the application
- Couple with Virtualization for higher agility
 - Migrate between servers with load balancing

Windows Server Failover Clustering


- Native Windows HA clustering with MySQL
- Failures & Planned Maintenance
- Quorum (3rd vote), data (InnoDB + schema) & binaries (optional) stored in shared storage (iSCSI & FCAL)
- Loss of service = couple of seconds + InnoDB recovery time
 - App sees temporary loss of connection and reconnects on same IP address
 - Replication slave recovers automatically
- Cluster managed through MS Failover Cluster Management snap-in GUI
- No new software/scripts required

Download the Guide: <http://bit.ly/w3rSoP>



**** Technical support for Windows Server Failover Clustering must be sourced from Microsoft.**

Solaris Clustering



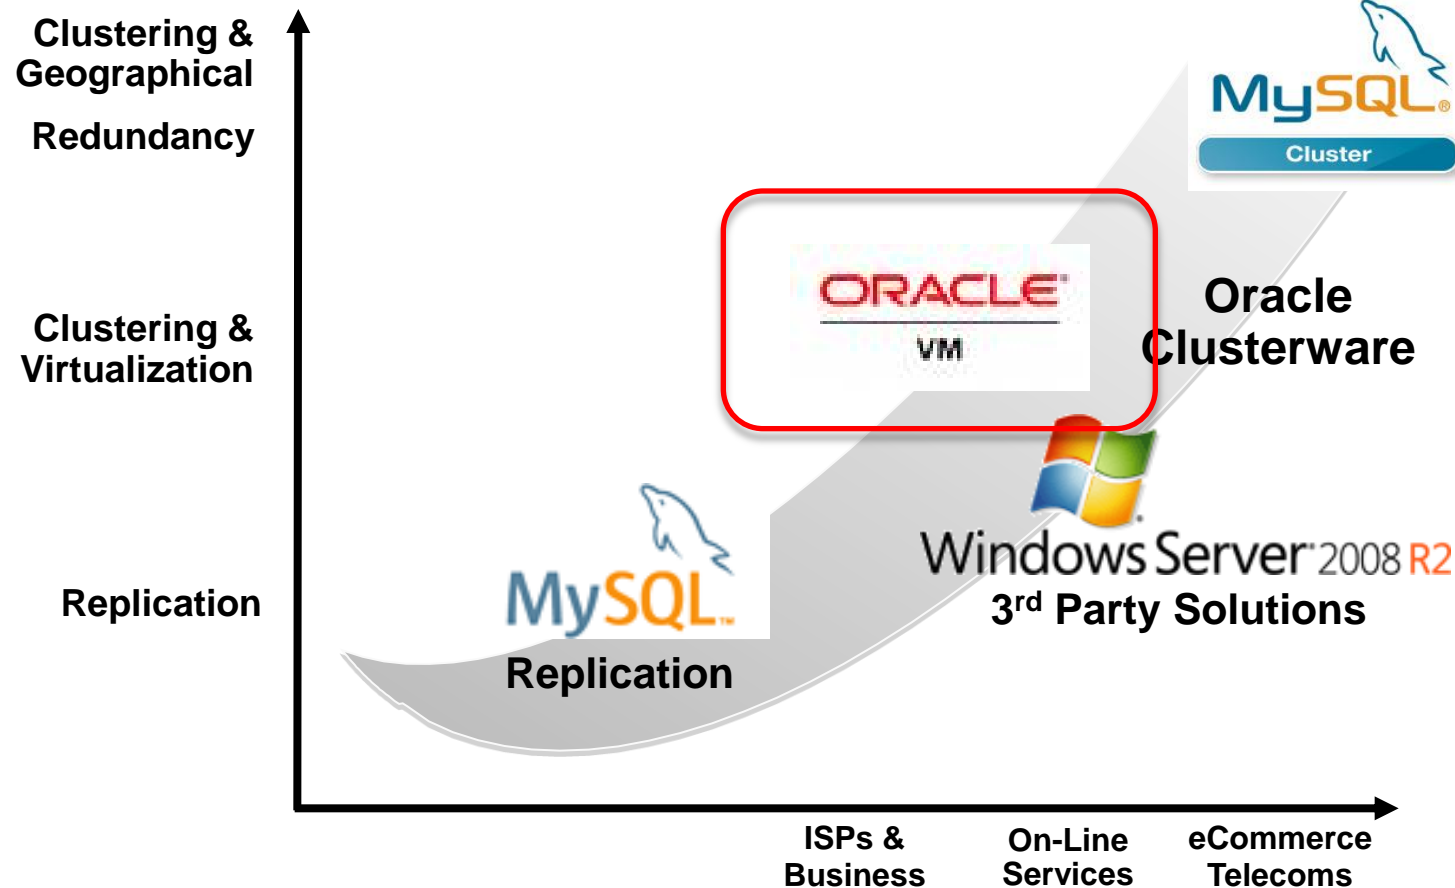
- Kernel based heartbeating and monitoring
- MySQL agent used to integrate with Solaris Clustering
 - Monitor MySQL Replication
- Learn more: <http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/index.html>

3rd Party HA Solutions

- Oracle supports MySQL on community / commercial HA technologies
 - Support for HA mechanisms from respective vendors
- Linux Heartbeat / Corosync with Pacemaker
 - Kernel level block replication using DRBD
- Symantec Veritas Cluster
- Red Hat Cluster Suite



Mapping Architectures to Availability

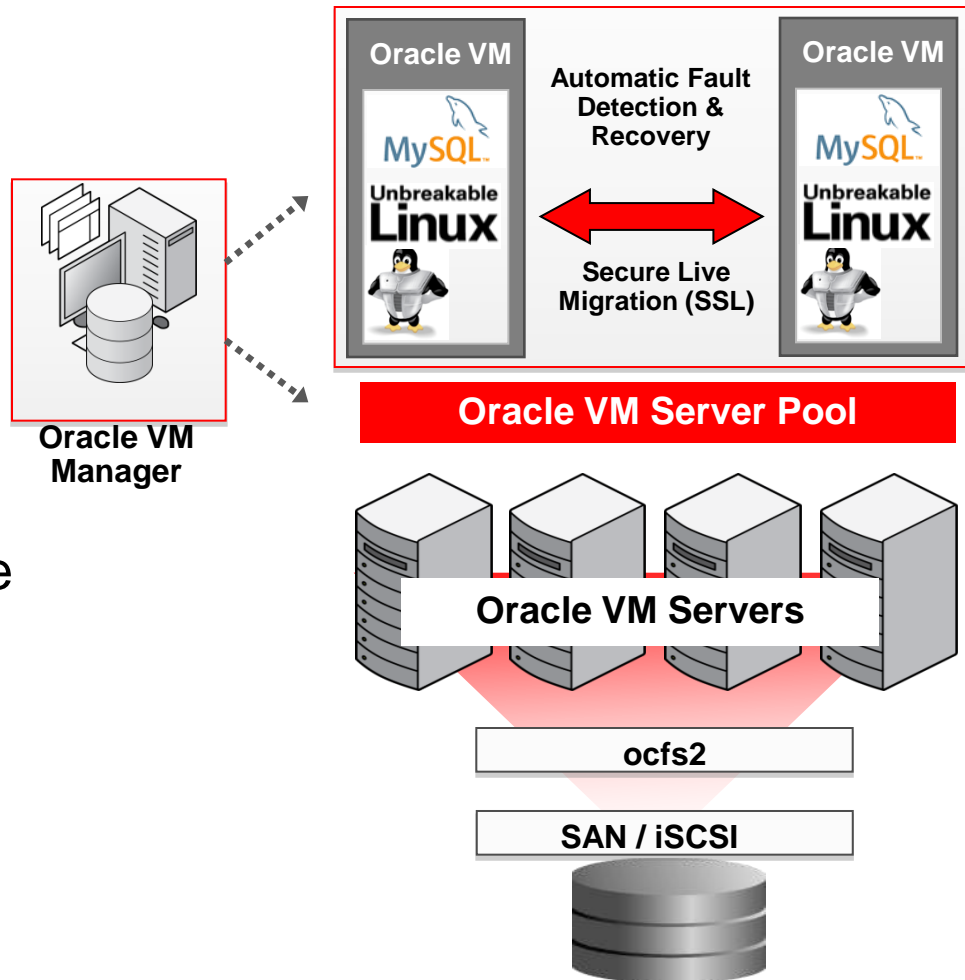


| | | | | | |
|--------------|-------------|--------------|--------------|-------------|---|
| 9 35 days | 9 4 days | 9 8 hours | 9 50 mins | 9 5 mins | % |
|--------------|-------------|--------------|--------------|-------------|---|

MySQL Enterprise High Availability

OVM Template for MySQL

- Oracle Linux with the Unbreakable Enterprise Kernel
- Oracle VM
- Oracle VM Manager
- Oracle Cluster File System 2 (OCFS2)
- MySQL Database (Enterprise Edition)
- Pre-Installed & Pre-Configured
- Full Integration & QA Testing
- Single Point of Support**




MySQL High Availability

Protecting Against Planned and Unplanned Downtime

Planned Events:

- Maintenance or upgrades
- Secure Live Migration
- Zero interruption



Failures

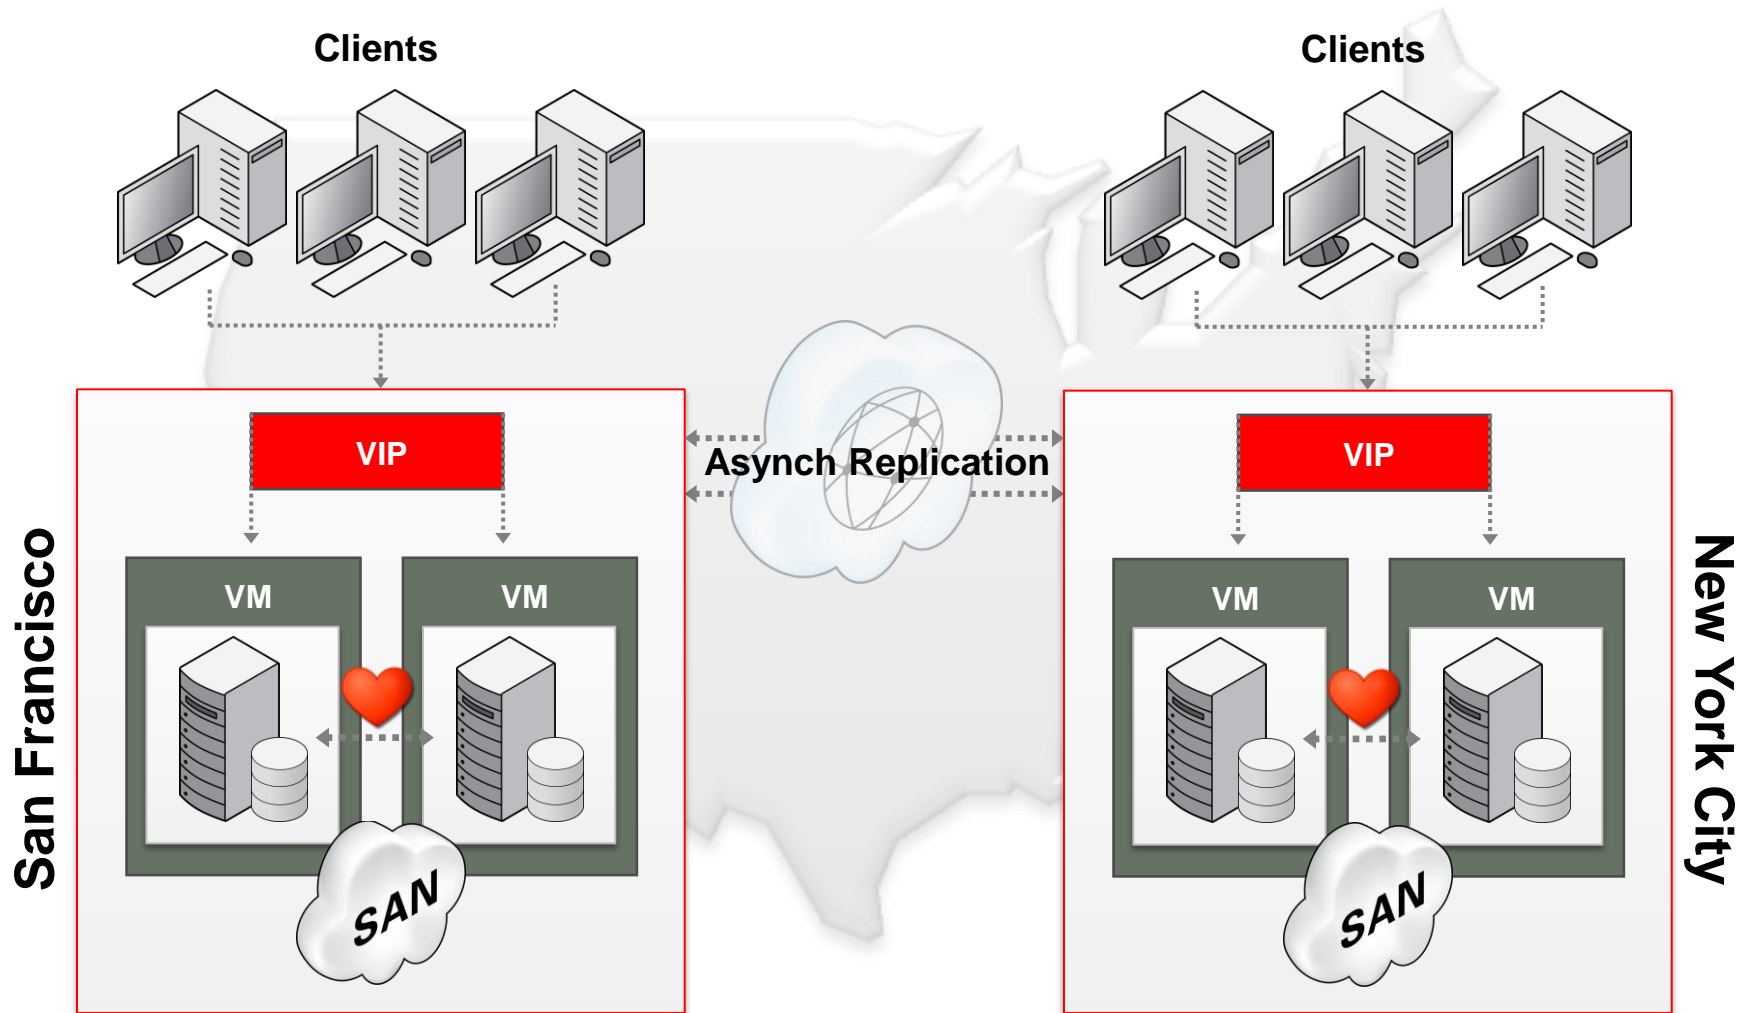
- Server, VM or database failure
- HA Auto-restart in pool
- Automatic failure detection & recovery

Disaster Recovery



- *“Disaster recovery is the process, policies and procedures related to preparing for recovery or continuation of technology infrastructure critical to an organization after a natural or human-induced disaster.”*

Disaster Recovery with MySQL Replication



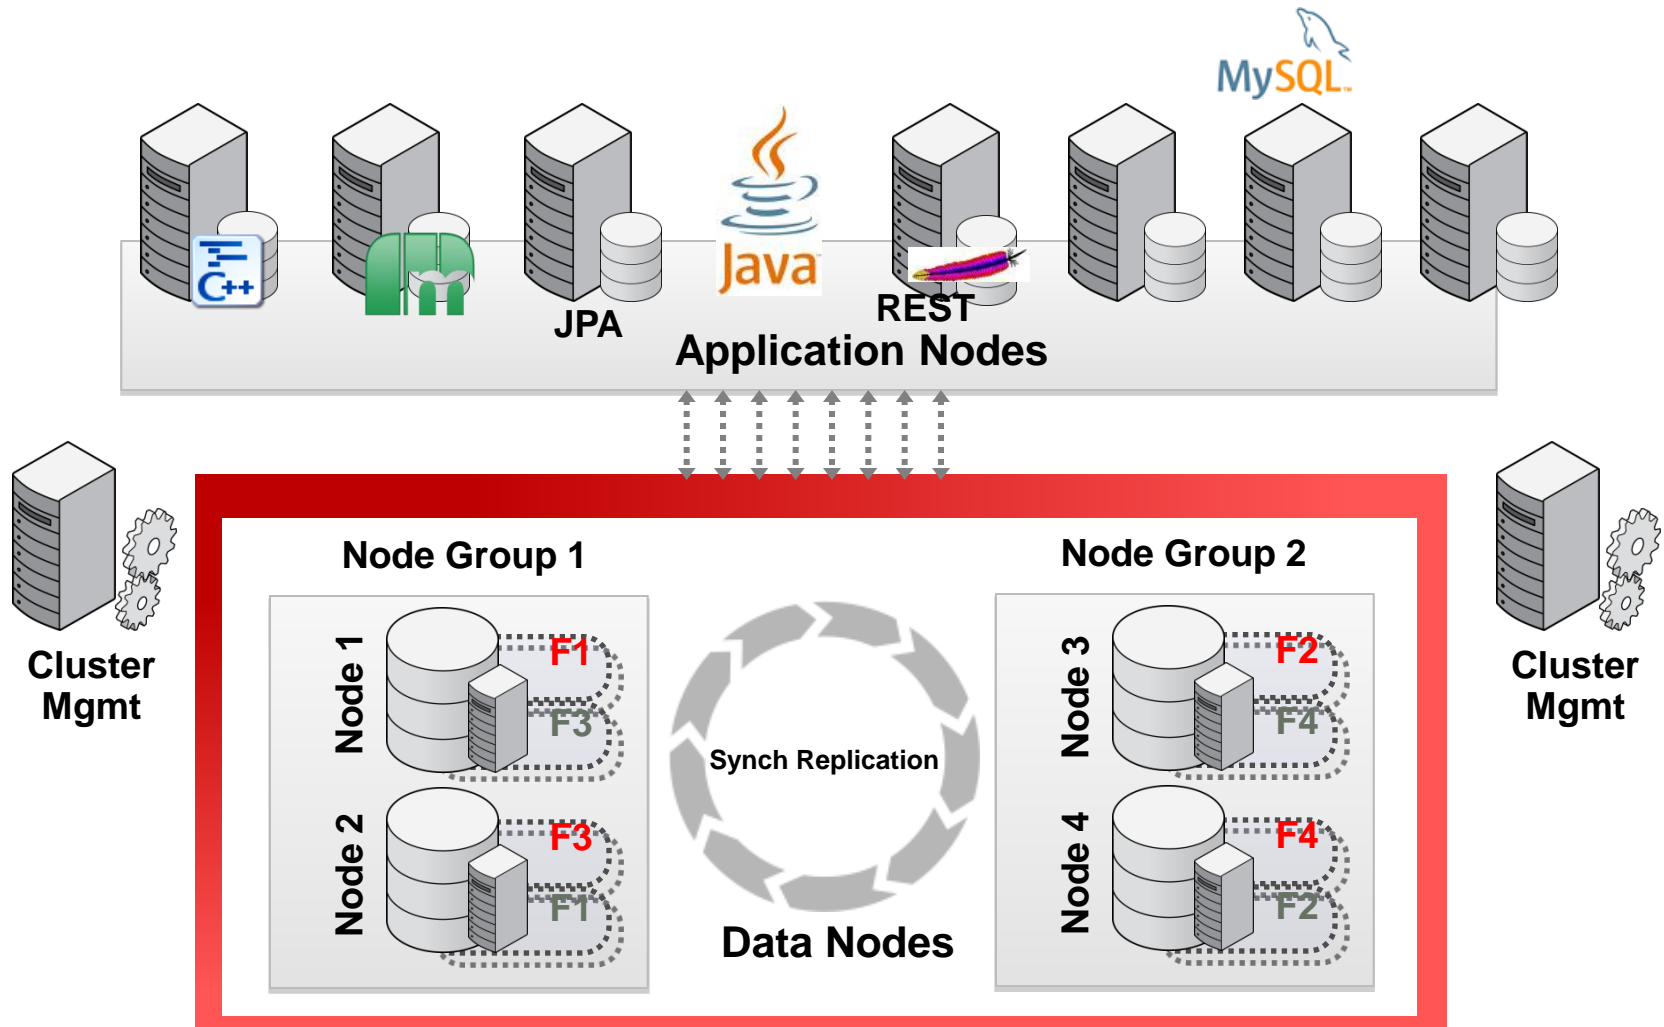
ORACLE®

Agenda

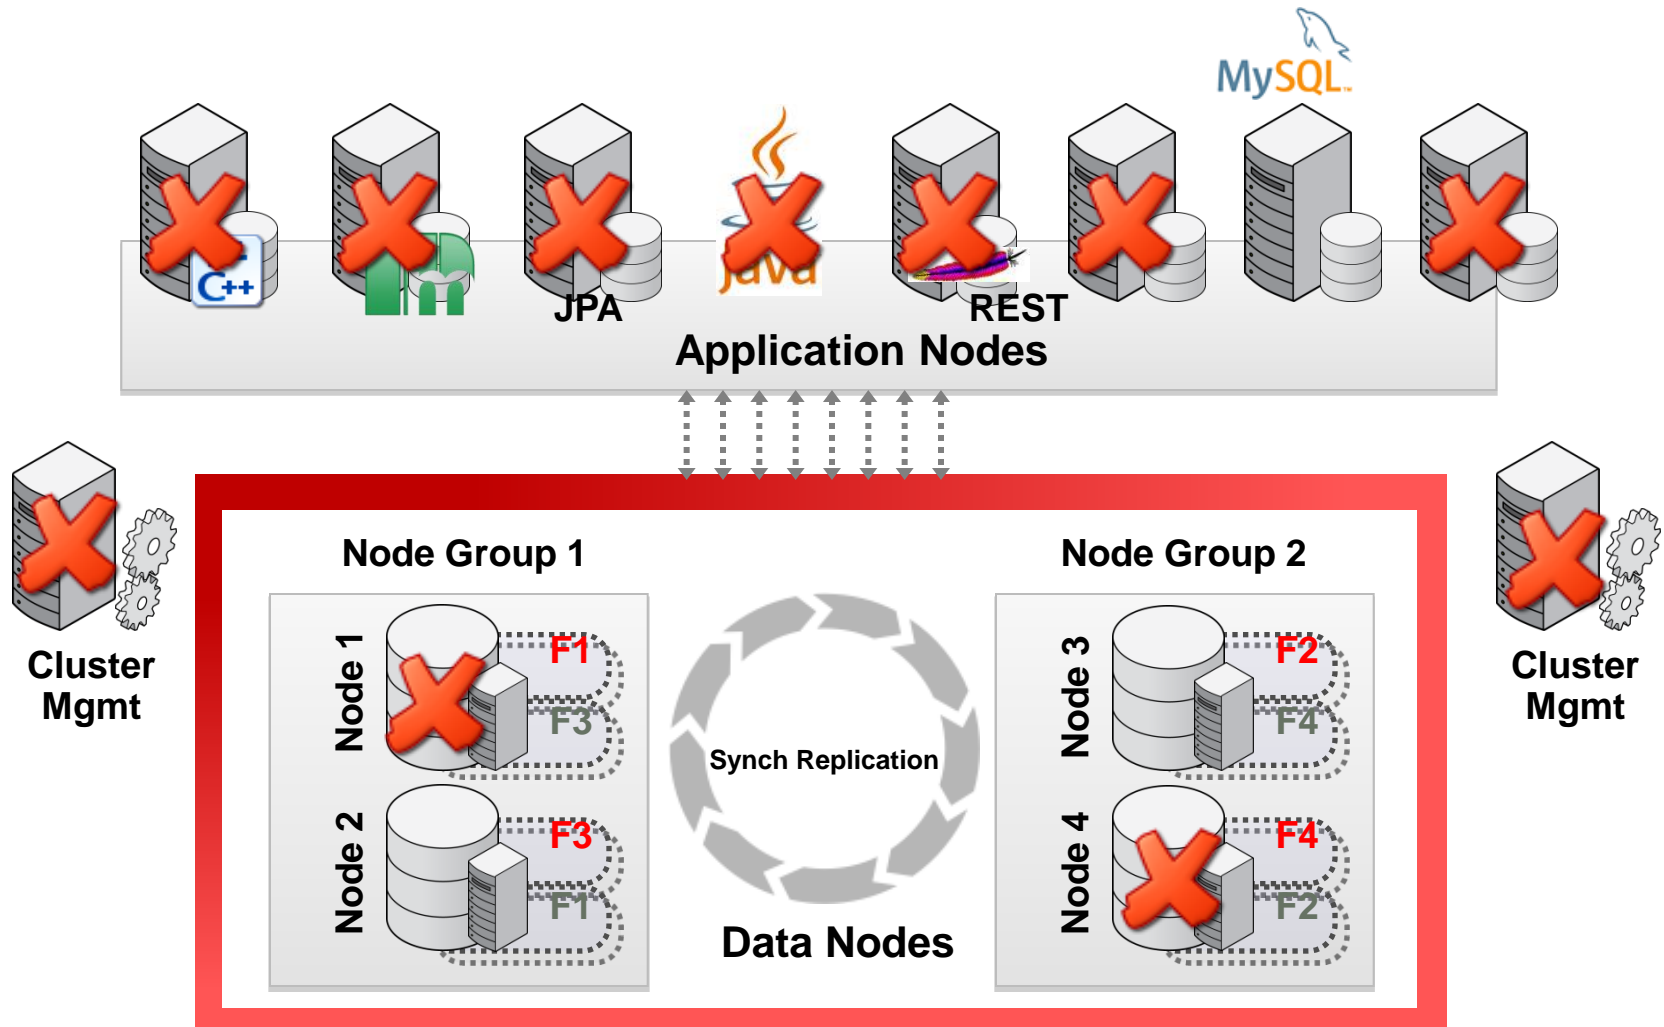
- Introduction to High Availability
- MySQL High Availability Solutions
- MySQL Cluster CGE
- Operational Best Practices
- Conclusion and Questions



MySQL Cluster Architecture



MySQL Cluster - Extreme Resilience



Handling Scheduled Maintenance

On-Line Operations

- Scale the cluster (add & remove nodes on-line)
- Repartition tables
- Upgrade / patch servers & OS
- Upgrade / patch MySQL Cluster
- Back-Up
- Evolve the schema on-line, in real-time

Automatic Data Sharding


4 Data Nodes and 2 Replicas

| IDCol | Col1 | Col2 | Col3 | Col4 | ... | |
|-------|------|------|------|------|-----|----|
| 1 | | | | | | F1 |
| 2 | | | | | | |
| 3 | | | | | | F2 |
| 4 | | | | | | |
| 5 | | | | | | F3 |
| 6 | | | | | | |
| 7 | | | | | | F4 |
| 8 | | | | | | |


F_n
F_n

Primary Replica
Secondary Replica

Node Group 1



Node Group 2





MySQL Cluster 7.2

1 Billion Queries per Minute

GA Now!

[Learn More »](#)

Servicing the Most Performance-Intensive Workloads

MySQL Cluster 7.2 GA

Enabling Next Generation Web Services

- 70x Faster Complex Queries: Adaptive Query Localization
- Native Memcached API
- MySQL 5.5 Server Integration
- Multithreaded Data Node Extensions
- VM Support

Enhancing Cross Data Center Scalability

- Multi-Site Clustering
- Simplified Active / Active Replication

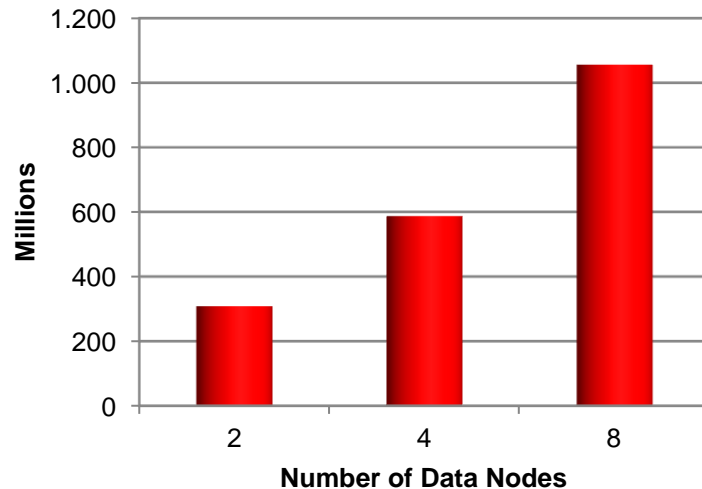
Ease of Use

- MySQL Cluster Manager 1.1.4
- Consolidated Privileges


<http://clusterdb.com/u/72wp>

MySQL Cluster 7.2 – 1 Billion Queries per Minute

SELECT Queries per Minute



UPDATE Queries per Minute



8 x Commodity Intel Servers
2 x 6-core processors
2.93GHz x5670 processors (24 threads)
48GB RAM
Linux OS

Infiniband networking
flexAsynch benchmark
C++ NoSQL API (NDB API)

Paggo

Customer Success Profile



COMPANY OVERVIEW

- Paggo has built an innovative payment solution making it easier for customers and merchants to buy and sell products and services using their mobile phones
- Enables buyers and sellers to conduct transactions using text messaging

CHALLENGES/OPPORTUNITIES

- Building a highly reliable system that delivers 24x7 availability
- Lost data results in lost business
- Build a system that could scale and perform as new merchants and buyers joined the network and transaction volumes increased

SOLUTIONS

- Oracle MySQL
- Java

KEY BUSINESS BENEFIT

- Since its launch with Oi, Brazil's largest GSM operator one year ago, MySQL Cluster Carrier Grade Edition has enabled Paggo to quickly add 40,000 merchants and 1.2 million subscribers to its network and handle over 750,000 financial transactions per month.

SUPPORTING QUOTE

"Paggo would not have been possible without MySQL Cluster. No other product would have give us the ability to start small, scale quickly and provide 24 x7 availability for our financial payment system."

Cicero Toreli
Founder - Paggo

ORACLE

Agenda

- Introduction to High Availability
- MySQL High Availability Solutions
- MySQL Cluster CGE
- Operational Best Practices
- Conclusion and Questions



Operational Best Practices

Training

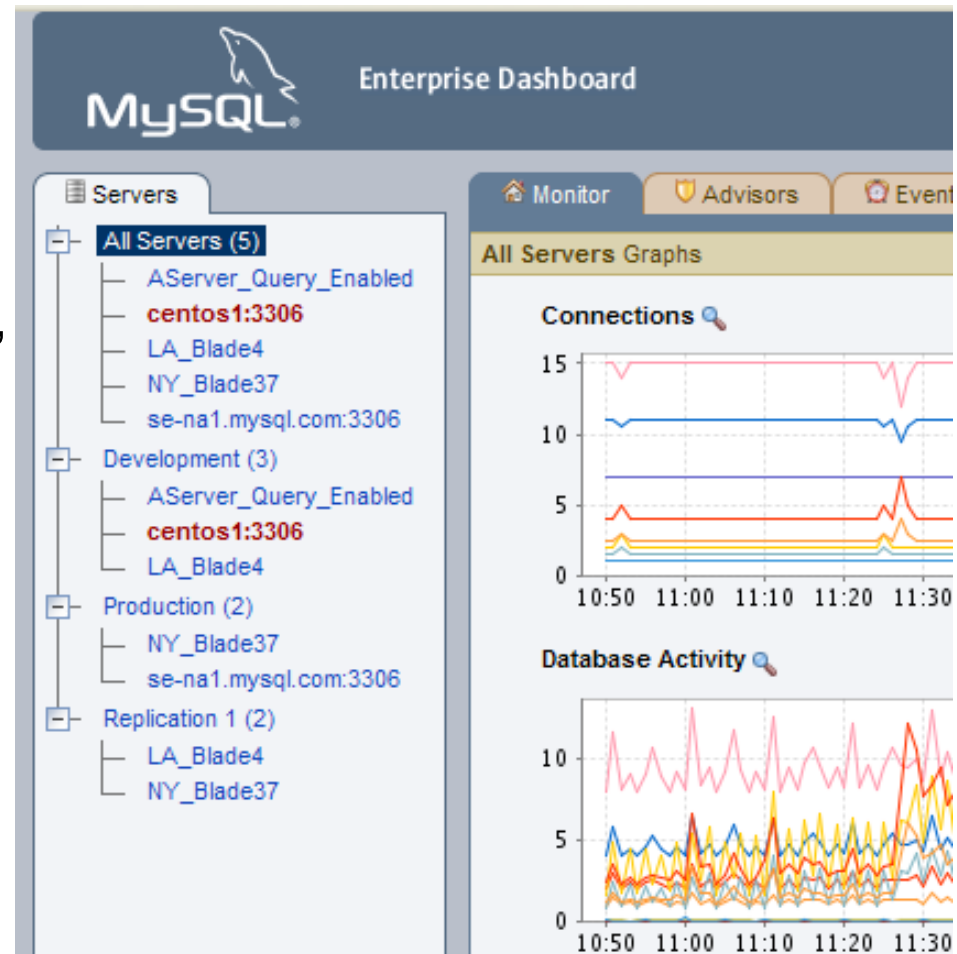
Consulting

24 x 7 x 365 Global Support

Database Management &
Monitoring

MySQL Enterprise Monitor

- Global view of MySQL environment
- Automated, rules-based monitoring and alerts (SMTP, SNMP enabled)
- Query capture, monitoring, analysis and tuning, correlated with Monitor graphs
- Visual monitoring of “hot” applications and servers
- Real-time Replication Monitor with auto-discovery of master-slave topologies
- Integrated with MySQL Support



A Virtual MySQL Tuning Assistant!

MySQL Expert Advisors



MySQL Cluster

- Monitors and Advises on status/ performance of MySQL Cluster Data Nodes.



Administration

- Monitors and Advises on Optimal Configuration



Performance

- Monitors and Advises on Optimal Performance Variable Settings



Custom

- Built by DBA to Enforce Organization specific best practices.



Upgrade

- Monitors and Advises on Bugs/Upgrades that affect current installation



Replication

- Monitors and Advises on Master/Slave Latency.



Security

- Monitors and Advises on Unplanned Security Changes/Loopholes



Schema

- Monitors and Advises on Unplanned Schema Change



Memory Usage

- Monitors and advises on optimal memory/cache settings

140+ Rules, 40+ MySQL, OS specific Graphs

Save you time writing, deploying, versioning, maintaining custom scripts.
Find problems and tuning opportunities you cannot find yourself.

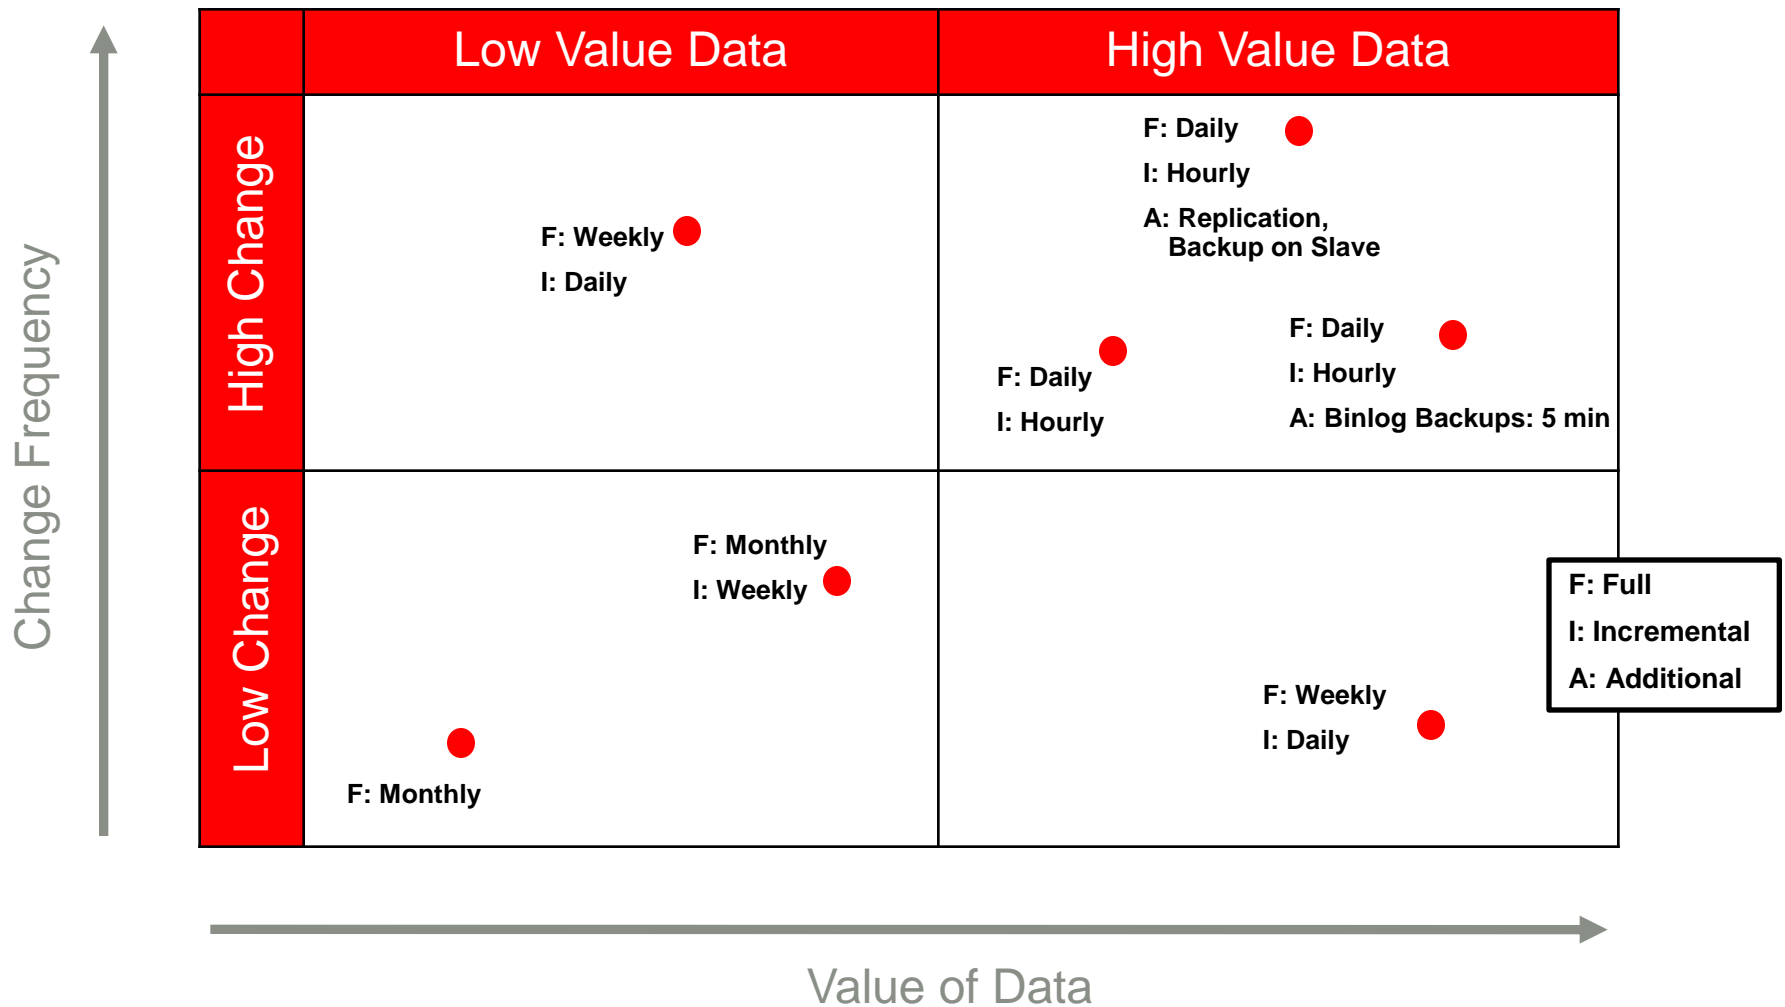
MySQL Enterprise Backup

- Online Backup for InnoDB
- Full, Incremental, Partial Backups (scriptable interface)
- Compression
- Point in Time, Full, Partial Recovery options
- Metadata on status, progress, history
- Unlimited Database Size
- Cross-Platform
 - Windows, Linux, Unix
- Certified with Oracle Secure Backup



Ensures quick, online backup and recovery of your MySQL apps.

Determining Backup Strategy



Backup Strategies Comparison

| Method | Backup Factors | Recovery Factors |
|---|--|---|
| Method 1: Full Backups | <ul style="list-style-type: none">• Longest Backup Times• Largest Storage Space• Save space with compression | <ul style="list-style-type: none">• Easy to Recover• Fastest Restore Times |
| Method 2: Full + Incremental Backup | <ul style="list-style-type: none">• Shortest Backup Time• Reduced Storage Requirements• Requires 1X production storage for copy | <ul style="list-style-type: none">• Finer-grained Recovery• Slower Restore Times• First Restore Full Backup• Then Restore Incrementals |
| Method 3: Full + Incremental + Log Backup | <ul style="list-style-type: none">• Added Storage Requirements• Requires more than 1X production storage for copy | <ul style="list-style-type: none">• Finest-grained Recovery• Slowest Restore Times• First Restore Full Backup• Then Restore Incrementals• Then Apply Logs |
| Method 4: Offload Backups Slave Replication | <ul style="list-style-type: none">• Used with 1 of the above• Frees Master for more workload• Requires 1X production hardware and storage for standby database | <ul style="list-style-type: none">• Fast failover to standby• Backups are last resort, in event of double site failure or need to perform PITR |

Oracle Premier Support for MySQL

Rely on The Experts - Get Unique Benefits

- Straight from the Source
- Largest Team of MySQL Experts
- Backed by MySQL Developers
- Forward Compatible Hot Fixes
- MySQL Maintenance Releases
- MySQL Support in 29 Languages
- 24/7/365
- Unlimited Incidents
- Knowledge Base
- MySQL Consultative Support

**Only From
Oracle**

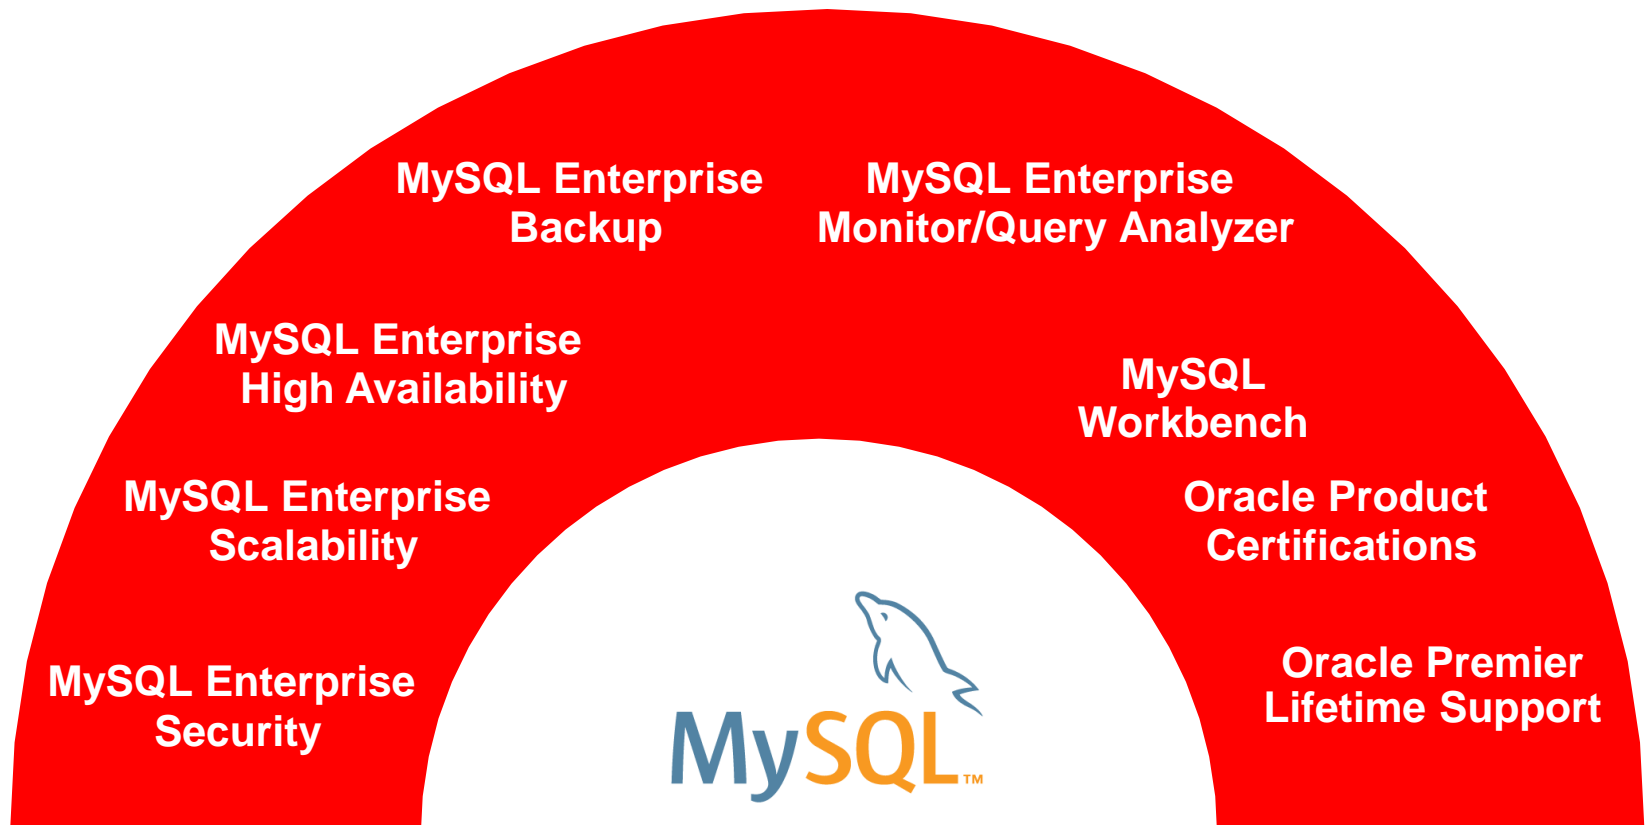
"The rep that assisted me was simply outstanding. He immediately recognized the cause of my problem and provided the resolution."

-- (July 27, 2011)

mysql.com/support/quotes

MySQL Enterprise Edition

Management & Monitoring Tools & Commercial Extensions, backed by Oracle Premier Lifetime Support



Agenda

- Introduction to High Availability
- MySQL High Availability Solutions
- MySQL Cluster CGE
- Operational Best Practices
- Conclusion and Questions

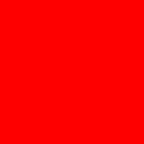


Mapping Applications to HA Technology

| Applications | Database [?] Replication | Clustered [?] / Virtualized | Shared-Nothing, [?] Geo-Replicated [?] Cluster |
|--|--------------------------------------|---|--|
| E-Commerce [?] /Trading | ???????1) [?] | ○ | ○ |
| Session [?] Management | ???????1) [?] | ○ | ○ |
| User [?] Authentication [?] /Accounting | ???????1) [?] | ○ | ○ |
| Feeds, [?] Blogs, [?] Wikis | ○ | | |
| OLTP | ???????1) [?] | ○ | ○ |
| Data [?] Warehouse/BI | ○ | ○ | ○ |
| Content [?] Management | ○ | ○ | |
| CRM | ○ | ○ | |
| Collaboration | ○ | ○ | |
| Packaged [?] Software | ○ | ○ | |
| Network [?] Infrastructure [?] | ○ | ○ | |
| Core [?] Telco [?] Apps [?] (HLR/HSS/SDP...) [?] | | | ○ |

1: Replication used in combination with cluster or virtualization – based HA





The preceding 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.

ORACLE®