



MySQL Cluster

Introduction and latest news from OOW

Ted Wennmark
ted.wennmark@oracle.com

Osaka (2nd) and Tokyo (4th) of December 2014

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.

Today's Database Requirements



Extreme Write Scalability



Rock Solid Reliability



Real-Time User Experience



Rapid Service Innovation

Today's Database Requirements

Transactional Integrity
OLTP + Real-Time Analytics
Standards & Skillsets

Extreme

Performance

Rock Solid Reliability

Rapid Service Innovation

Who's Using MySQL Cluster



MySQL Cluster Overview

HIGH SCALE, READS + WRITES

- Auto-Sharding, Multi-Master
- ACID Compliant, OLTP + Real-Time Analytics

99.999% AVAILABILITY

- Shared nothing, no Single Point of Failure
- Self Healing + On-Line Operations

REAL-TIME

- In-Memory Optimization + Disk-Data
- Predictable Low-Latency, Bounded Access Time

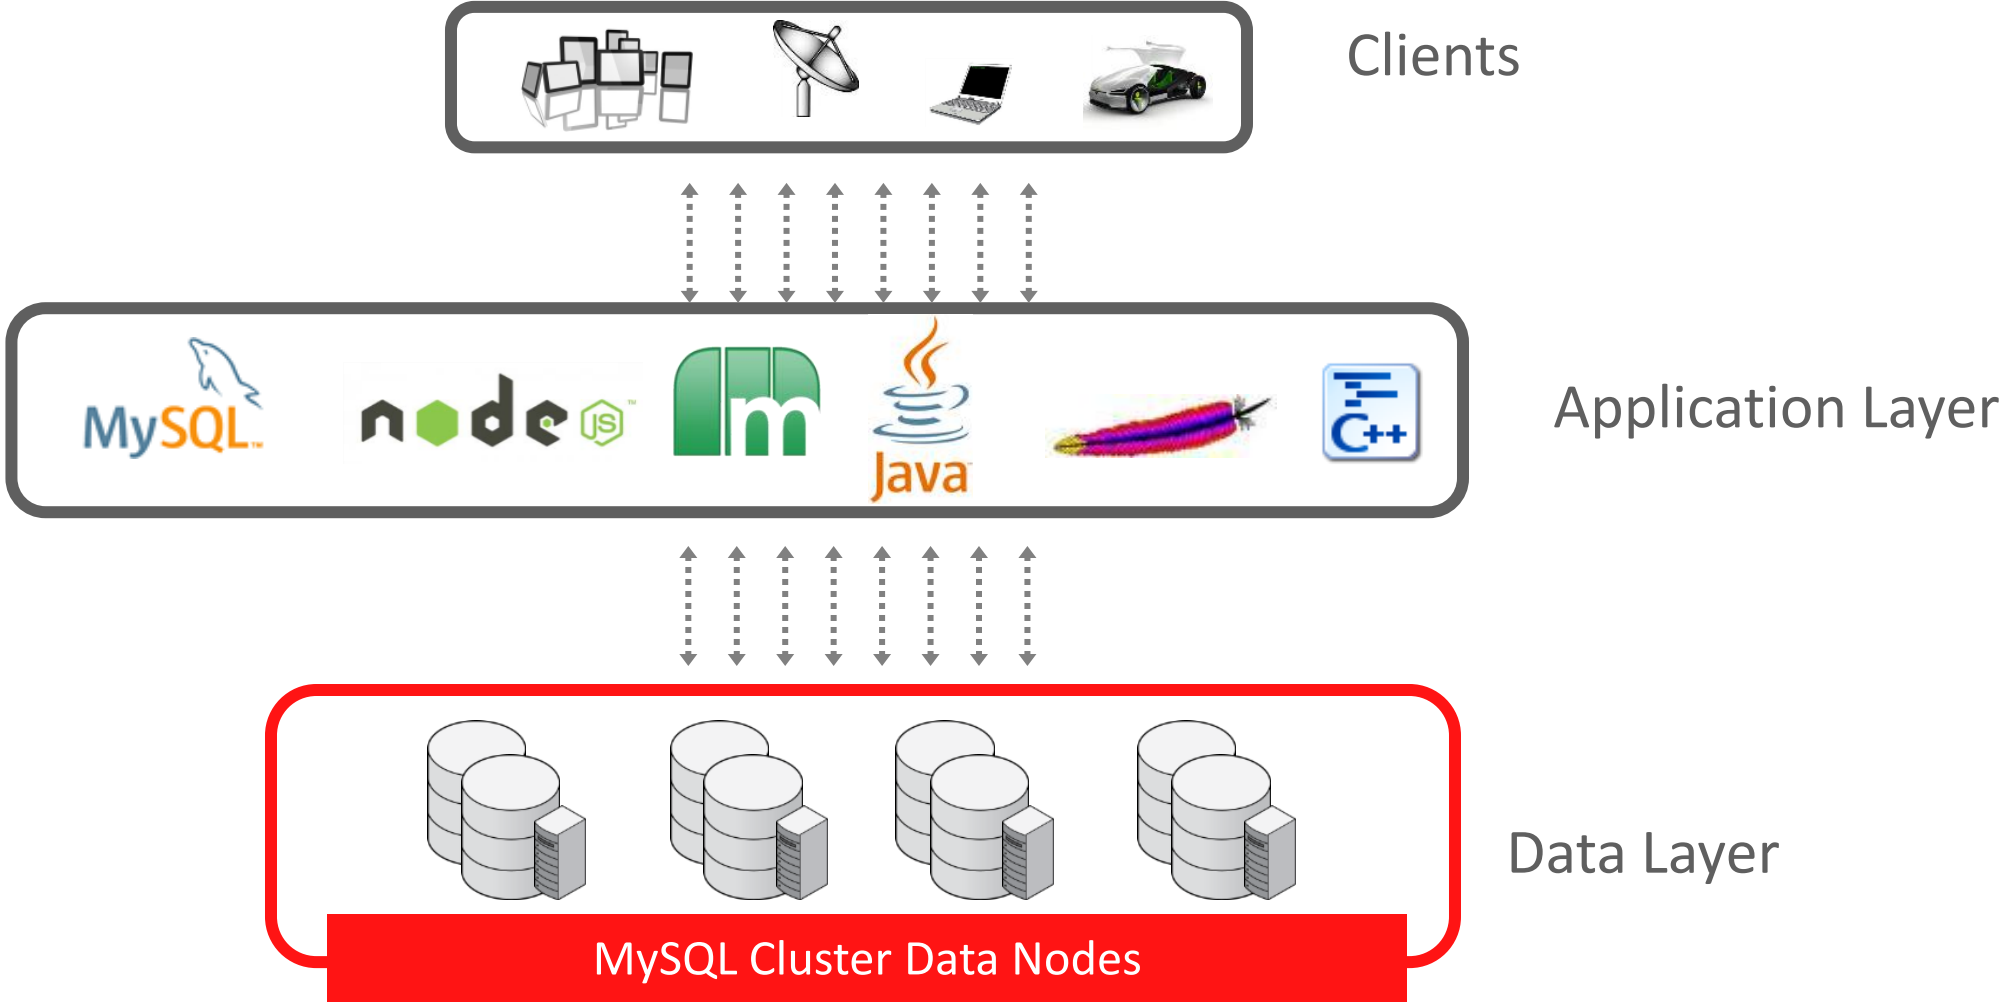
SQL + NoSQL

- Key/Value + Complex, Relational Queries
- SQL + Memcached + JavaScript + Java + HTTP/REST & C++


LOW TCO

- Open Source + Commercial Editions
- Commodity hardware + Management, Monitoring Tools


MySQL Cluster Architecture





MySQL Cluster Scaling



Clients



Application Layer



Data Layer


MySQL Cluster Data Nodes





MySQL Cluster HA



Clients



Application Layer

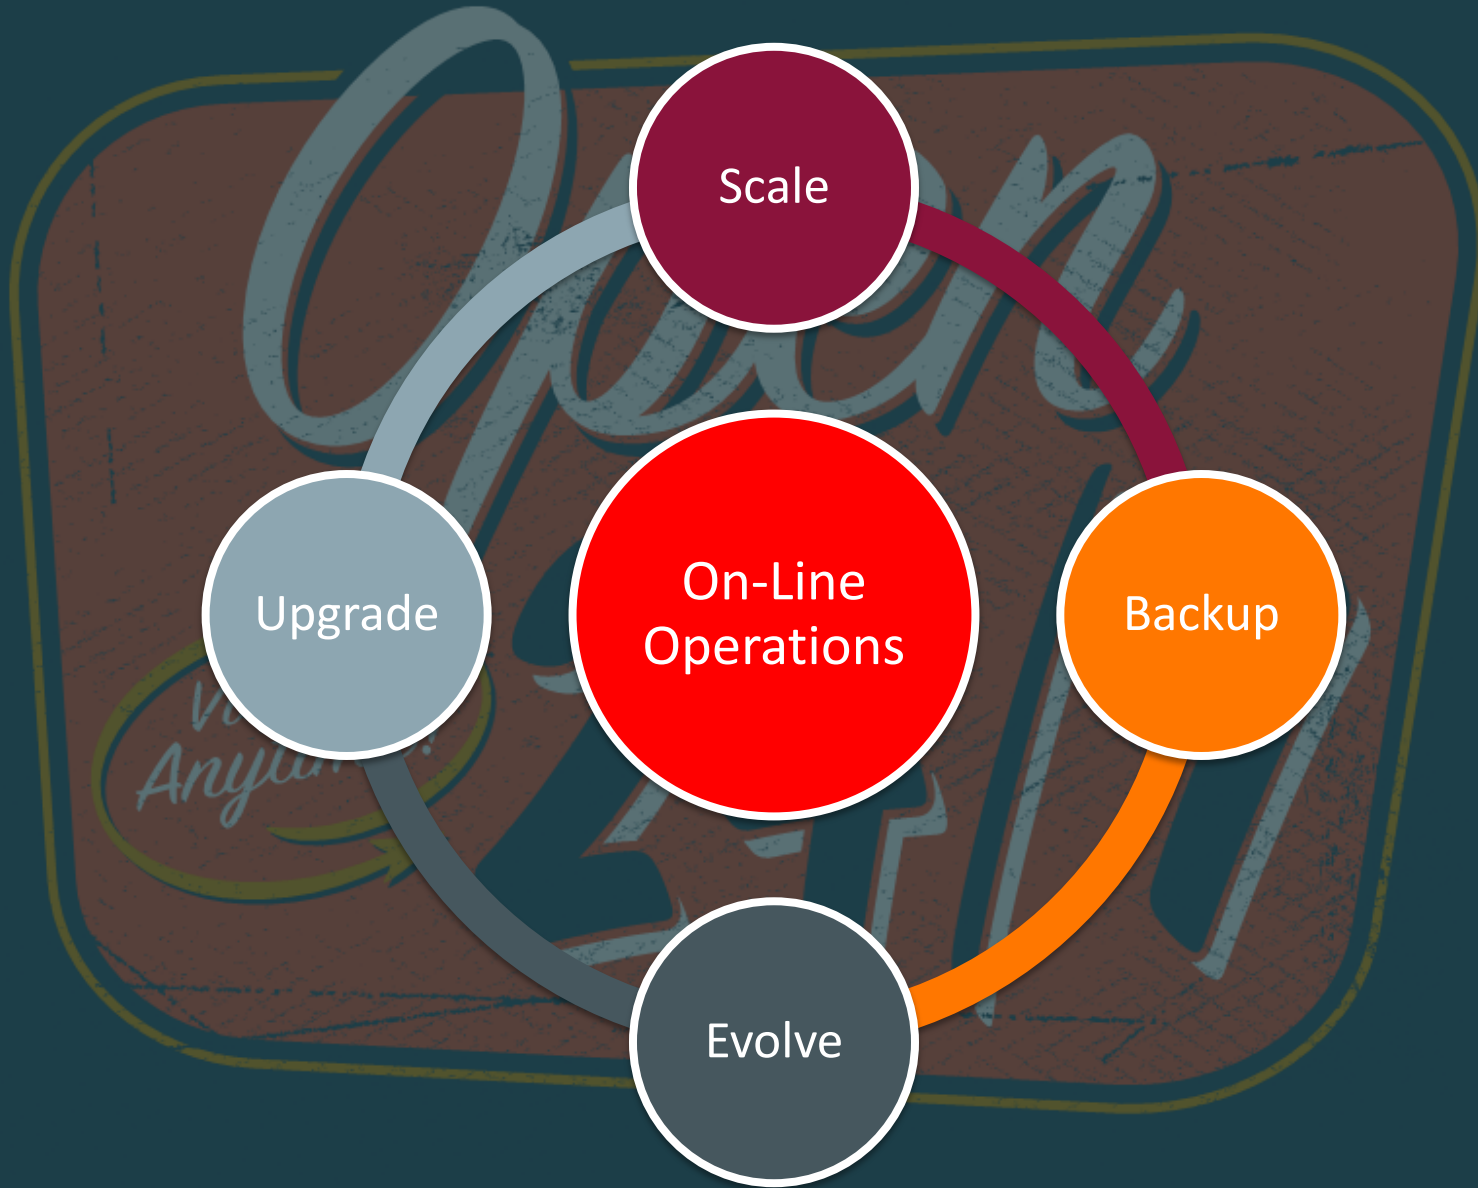


Data Layer

Open

*Visit Us
Anytime!*

24/7



Scale

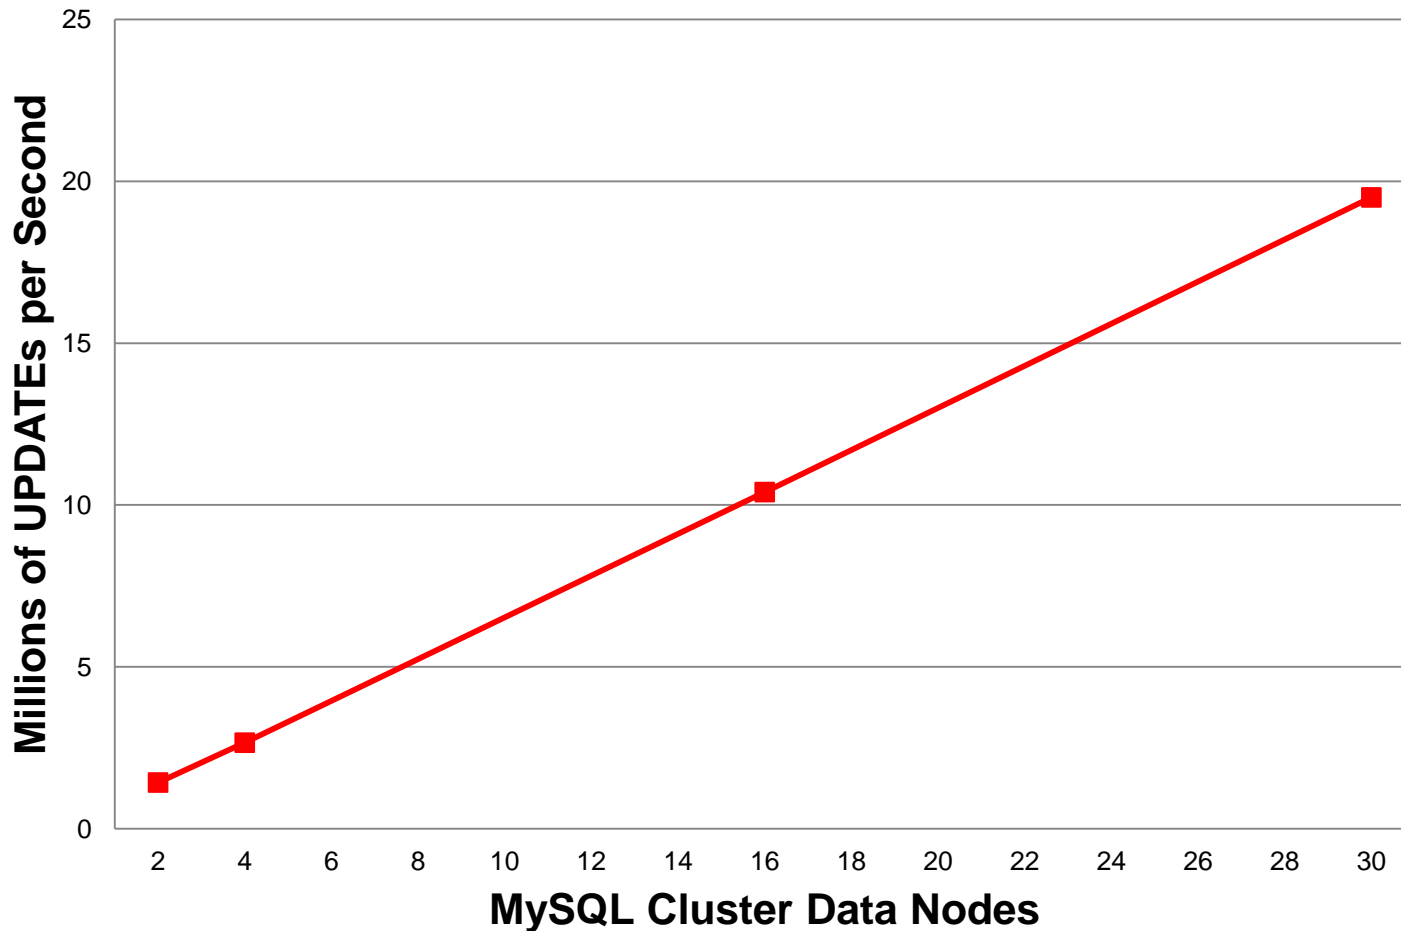
On-Line
Operations

Backup

Evolve

Upgrade

MySQL Cluster 7.3: 1.2 Billion UPDATES per Minute



- NoSQL C++ API, flexaSynch benchmark
- 30 x Intel E5-2600 Intel Servers, 2 socket, 64GB
- ACID Transactions, with Synchronous Replication

SQL

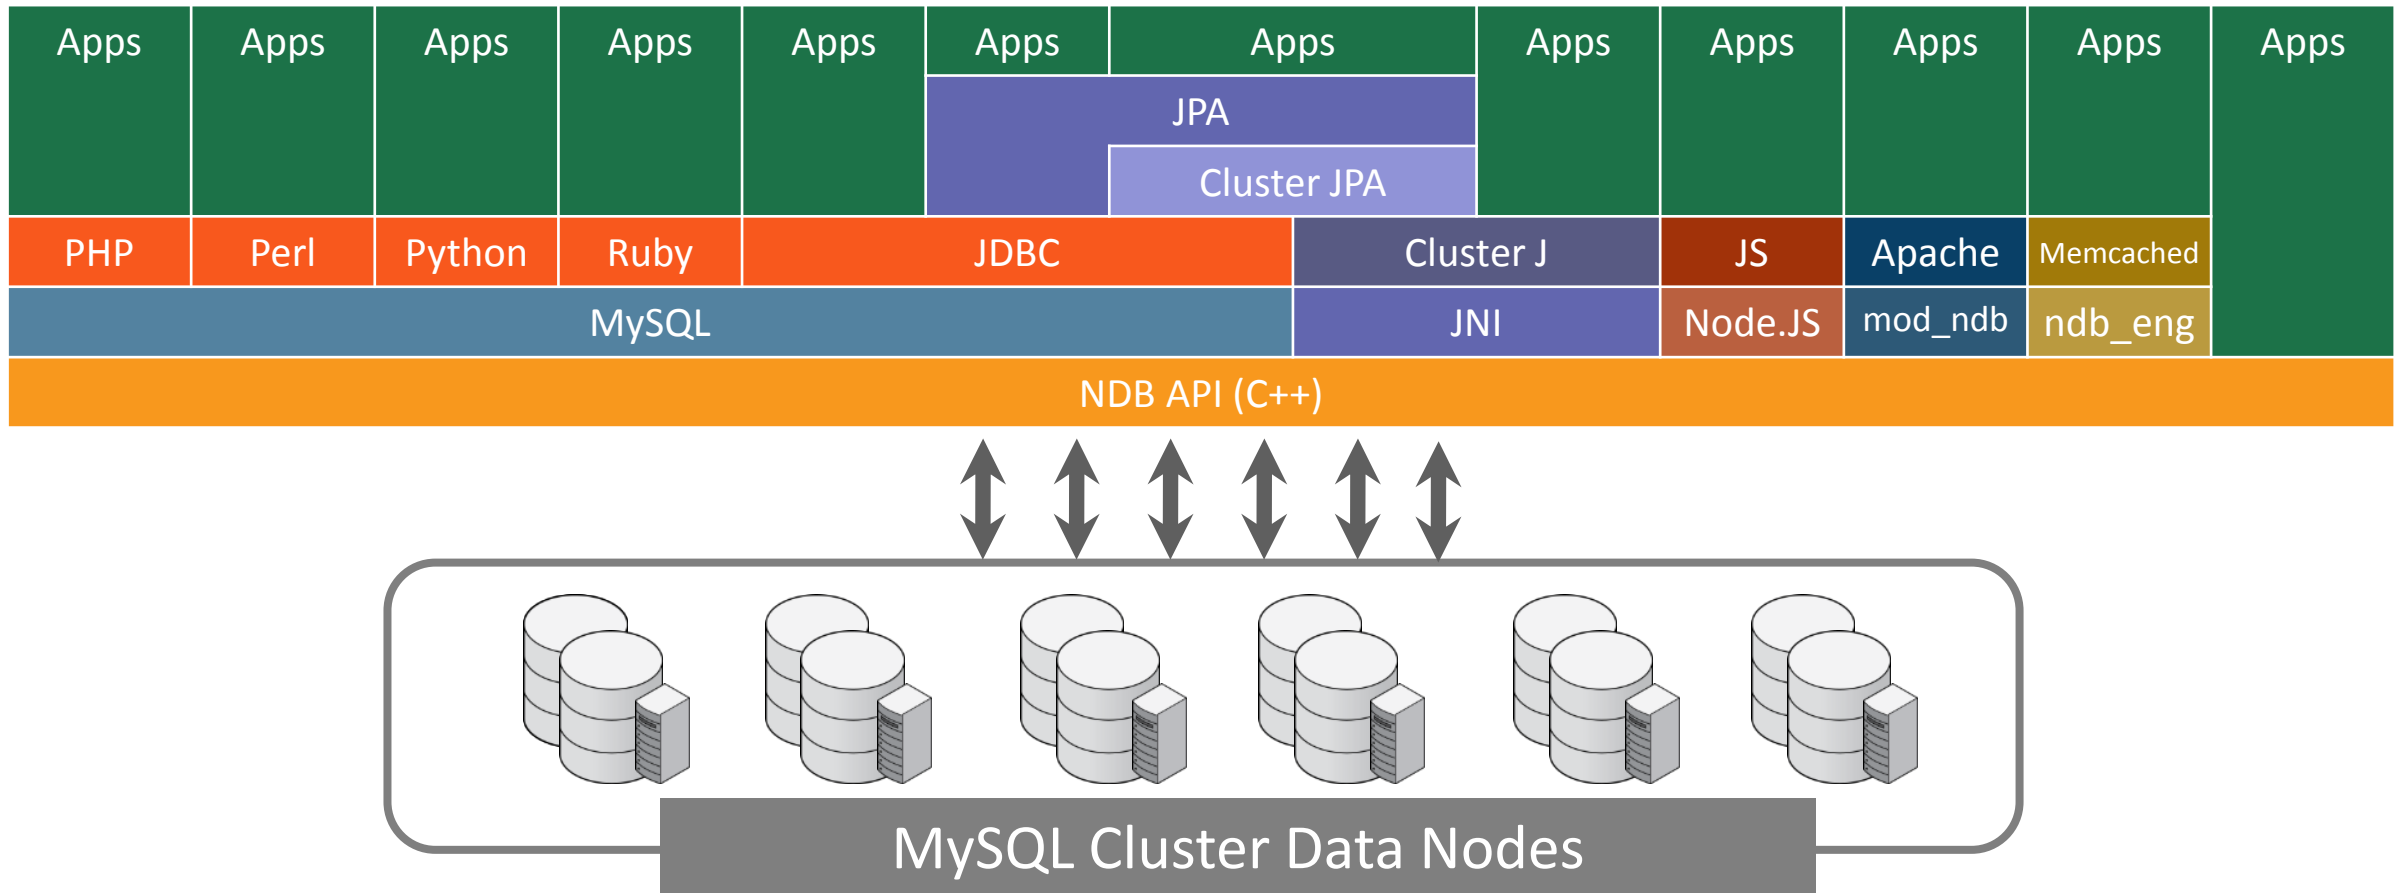


?

NosQL



NoSQL Access to MySQL Cluster data



Cluster & Memcached – Schema-Free

meal:lunch-cod random-96
age:fred-22 home:blog-clusterdb.com
nick:james-jimmy edges:triangle-3
town:maidenhead-SL6 town:reading-RG1
edges:square-4
hair:fred-mohawk

Application view

key value
<town:maidenhead, SL6>

SQL view

key value
<town:maidenhead, SL6>

Key	Value
town:maidenhead	SL6

generic table

Cluster & Memcached - Configured Schema

meal:lunch-cod random-96
 home:blog-clusterdb.com
 edges:triangle-3
 town:reading-RG1
 edges:square-4
 hair:Fred-mohawk
 age:fred-22
 nick:james-jimmy
 town:maidenhead-SL6

key value
<town:maidenhead, SL6>

Application view

SQL view

prefix key value
<town:maidenhead, SL6>

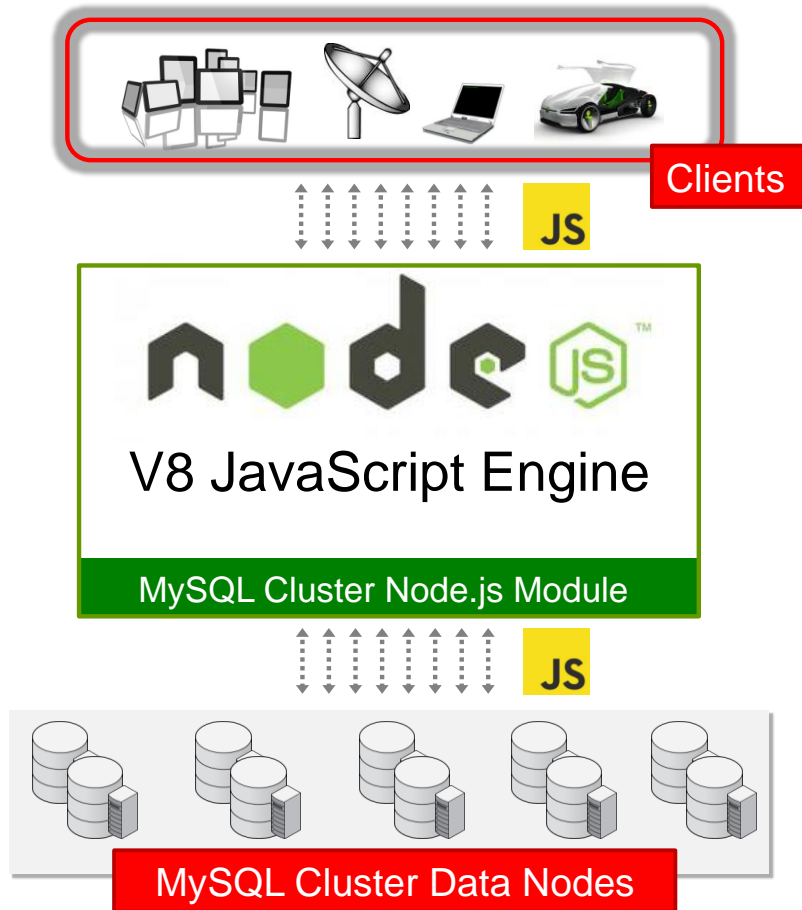
Prefix	Table	Key-col	Val-col	policy
town:	map.zip	town	code	cluster

Config tables

town	...	code	...
maidenhead	...	SL6	...

map.zip

Node.js NoSQL API





- Native JavaScript access to MySQL Cluster
 - End-to-End JavaScript: browser to the app & DB
 - Storing and retrieving JavaScript objects directly in MySQL Cluster
 - Eliminate SQL transformation
- Implemented as a module for node.js
 - Integrates Cluster API library within the web app
- Couple high performance, distributed apps, with high performance distributed database
- Optionally routes through MySQL Server

MySQL Cluster 7.4.1 DMR

Available Now!

MySQL Cluster 7.4.1 DMR

Better performance and operational simplicity



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

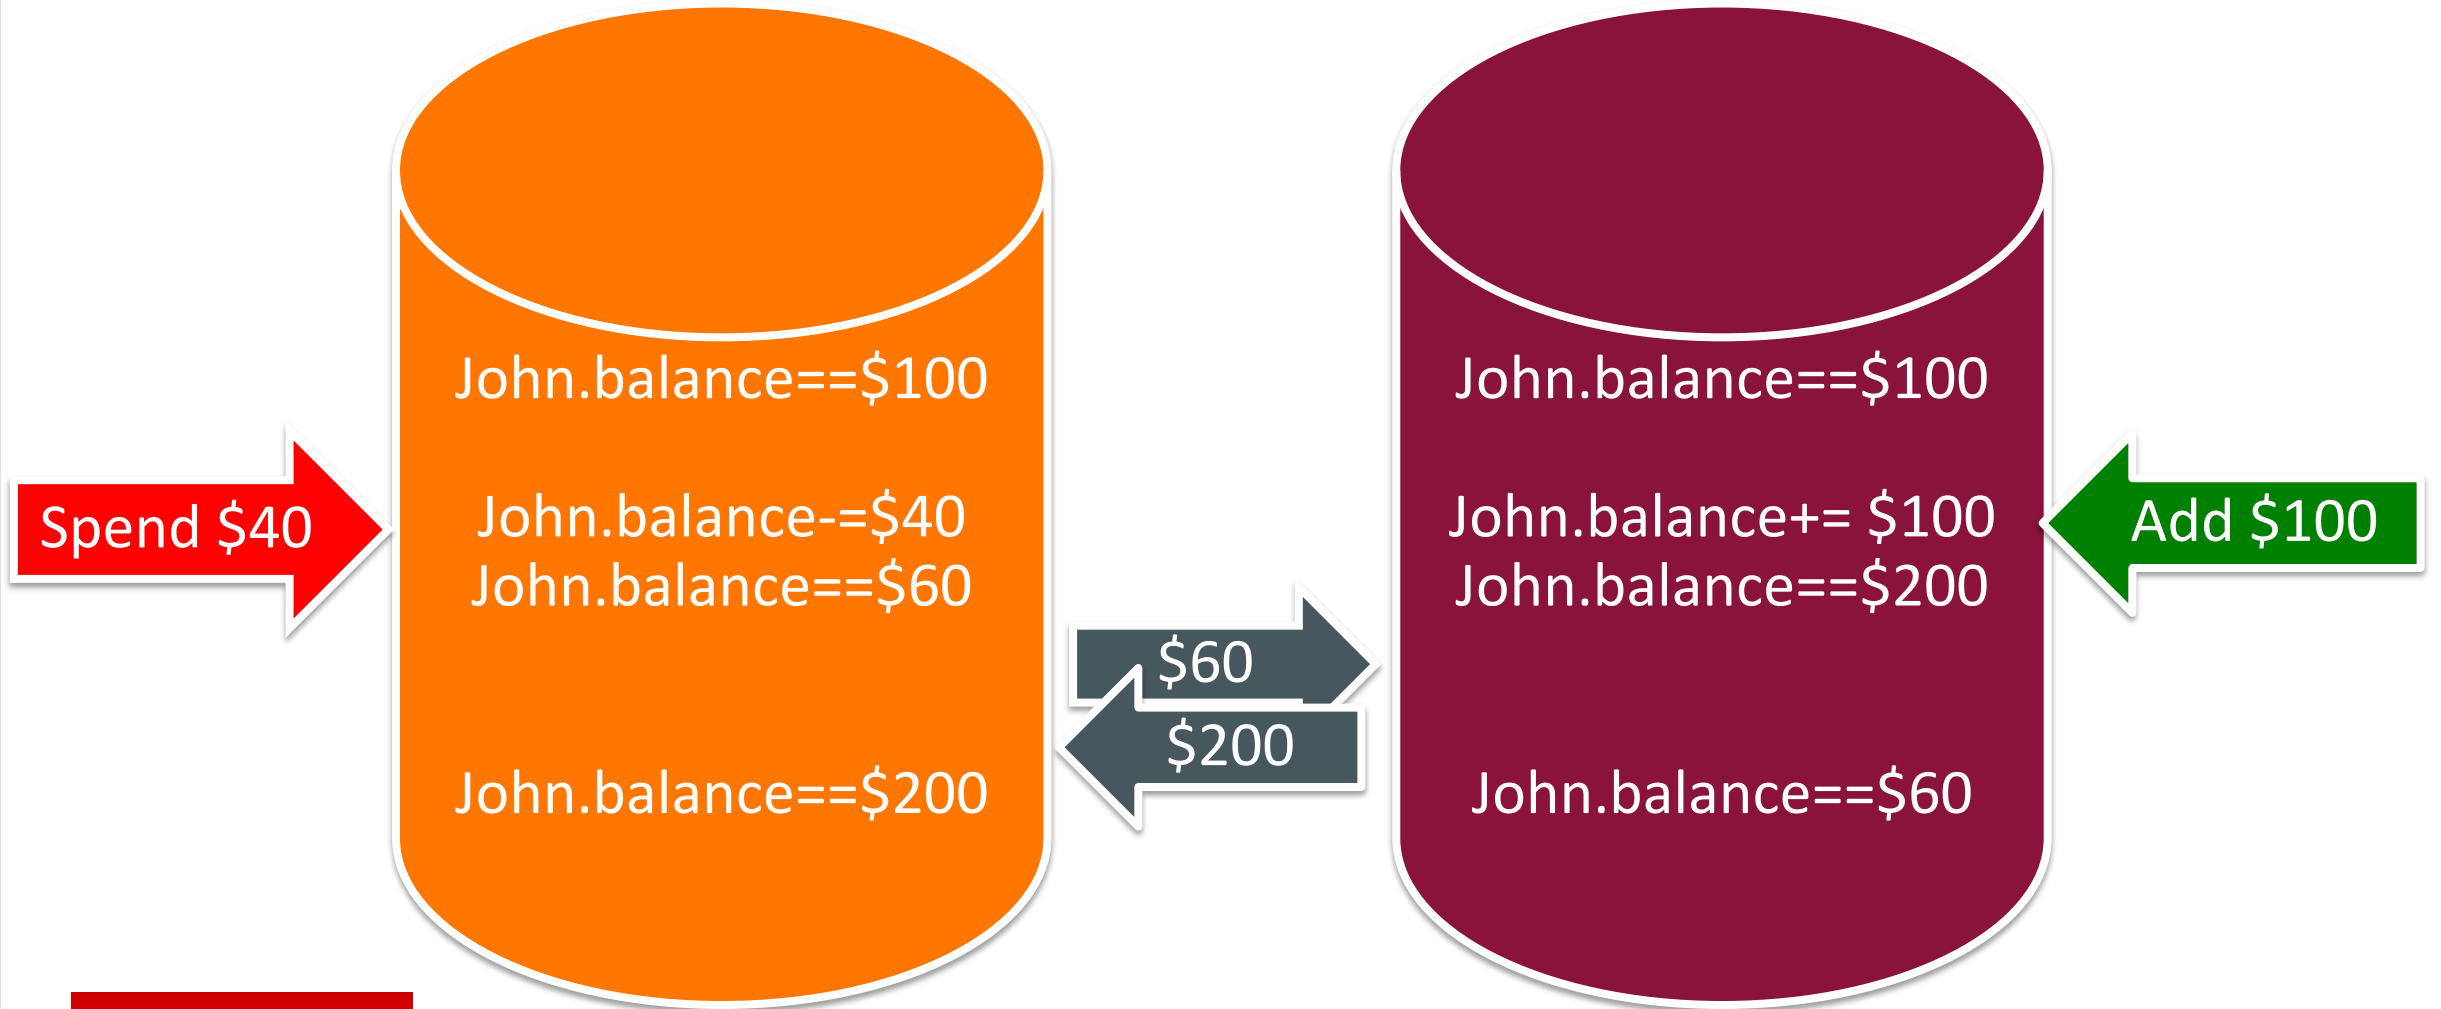
- Faster maintenance operations
 - Nodal & Rolling restarts
 - Upgrades

Active-Active Geo-Replication

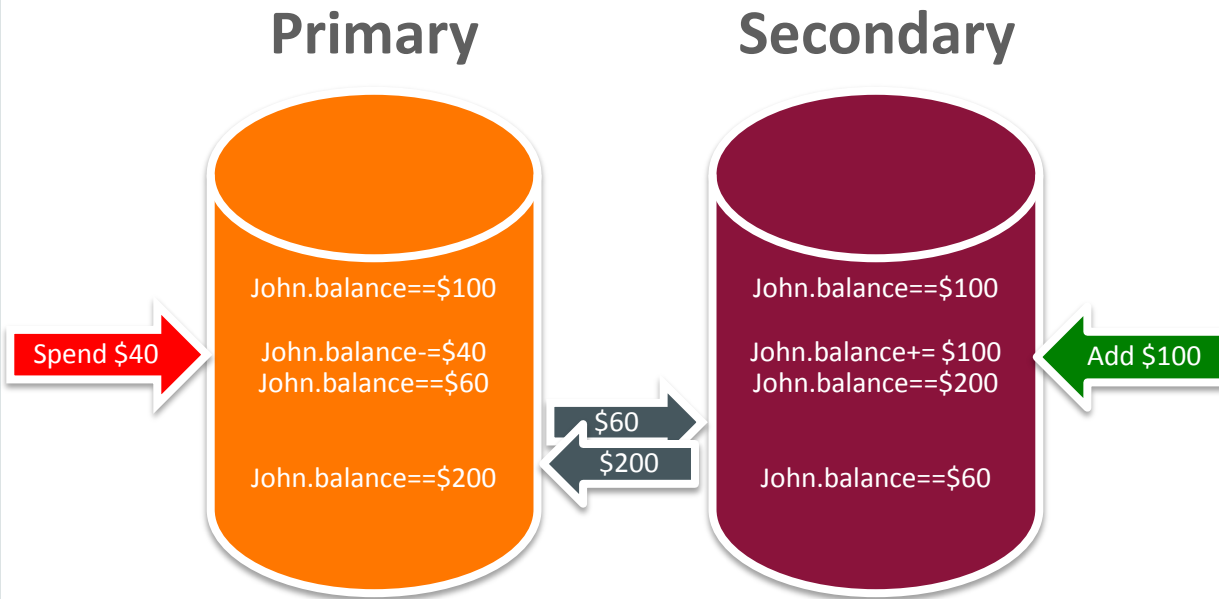


- Asynchronous replication between MySQL Clusters
- Active-Active
 - Update anywhere
 - Conflict detection
 - Application notified through exception tables
 - Can opt to have conflicts resolved automatically
 - Auto-conflict-resolution
 - NDB\$EPOCH, conflicting rows are rolled back
 - NDB\$EPOCH_TRANS, conflicting transactions are rolled back

What is a conflict?



Detecting Conflicts - Reflected GCI



- Primary store logical timestamp (GCI) against updated row
 - Window for conflict opens
- GCI replicated with updated row to Secondary
- The same row and GCI is replicated back (reflected) from Secondary to Primary after it has been applied
 - Closing window for conflict
- Primary checks every event originating from the Secondary to ensure it isn't for a 'conflictible' row

Handling of Conflicts

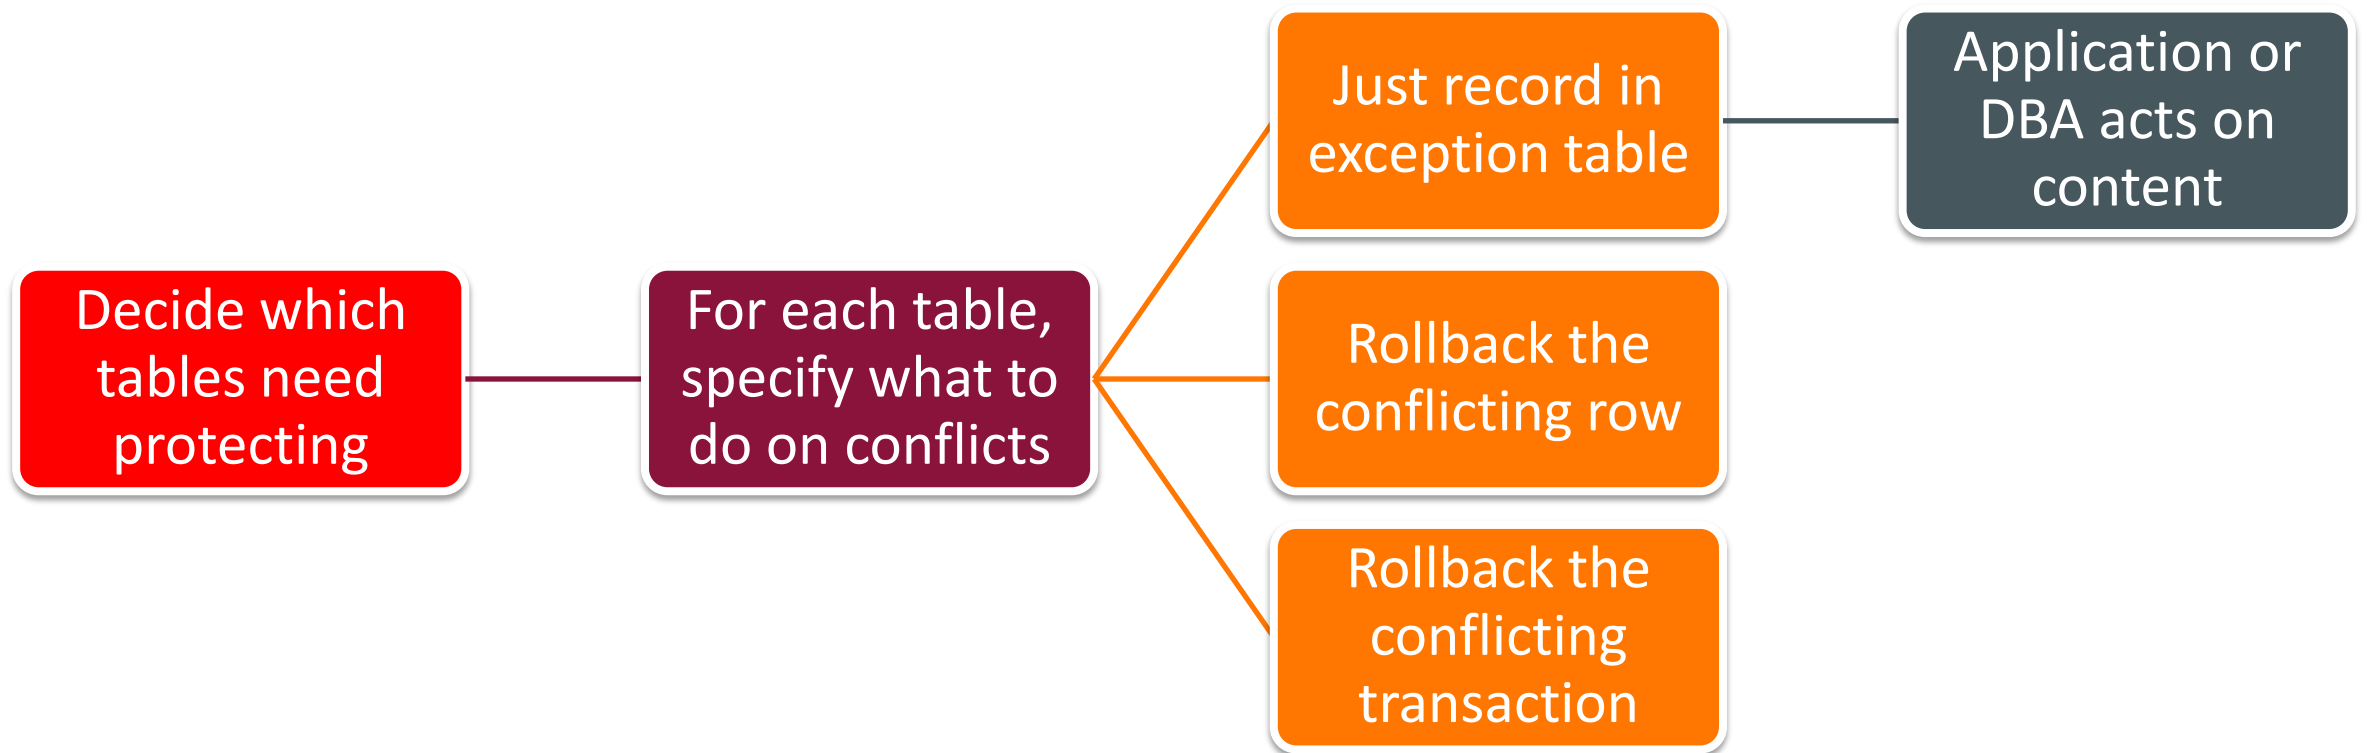
MySQL Cluster 7.4.1 DMR

- Rolling back of transactions that read conflicted data.
- Entire transactions (and dependent ones) rolled back.
- Information of conflict's type, cause, and originating transaction.

Later in MySQL Cluster 7.4

- Conflicting deletes

How to Use Conflict Detection/Resolution



Restart Times

What operations benefit?

- Restarting data node with locally checkpointed data
 - Major improvement
- Restarting data node which must recover data from peer
 - Major improvement
 - Further speedups to come in 7.4.X (greater parallelization)
- Upgrade/rolling restarts
 - Major improvement
- Cluster shutdown and restart
 - Minor improvement

Enhanced Memory Reporting

See how much memory a table is using

```
mysql> CREATE DATABASE clusterdb;USE clusterdb;
mysql> CREATE TABLE simples (id INT NOT NULL AUTO_INCREMENT PRIMARY KEY) ENGINE=NDB;
mysql> select node_id AS node, fragment_num AS frag, fixed_elem_alloc_bytes alloc_bytes,
fixed_elem_free_bytes AS free_bytes, fixed_elem_free_rows AS spare_rows from
memory_per_fragment where fq_name like '%simples%';
```

node	frag	alloc_bytes	free_bytes	spare_rows
1	0	131072	5504	172
1	2	131072	1280	40
2	0	131072	5504	172
2	2	131072	1280	40
3	1	131072	3104	97
3	3	131072	4256	133
4	1	131072	3104	97
4	3	131072	4256	133

Enhanced Memory Reporting

See how memory is made available after deleting rows

node	frag	alloc_bytes	free_bytes	spare_rows
1	0	131072	5504	172
1	2	131072	1280	40
2	0	131072	5504	172
2	2	131072	1280	40
3	1	131072	3104	97
3	3	131072	4256	133
4	1	131072	3104	97
4	3	131072	4256	133

```
mysql> DELETE FROM clusterdb.simples LIMIT 1;
```

node	frag	alloc_bytes	free_bytes	spare_rows
1	0	131072	5504	172
1	2	131072	1312	41
2	0	131072	5504	172
2	2	131072	1312	41
3	1	131072	3104	97
3	3	131072	4288	134
4	1	131072	3104	97
4	3	131072	4288	134

Enhanced Memory Reporting

Check how well partitioned/sharded a table is

```
mysql> CREATE TABLE simples (id INT NOT NULL AUTO_INCREMENT, species VARCHAR(20) DEFAULT "Human",  
PRIMARY KEY(id, species)) engine=ndb PARTITION BY KEY(species);
```

```
// Add some data
```

```
mysql> select node_id AS node, fragment_num AS frag, fixed_elem_alloc_bytes alloc_bytes,  
fixed_elem_free_bytes AS free_bytes, fixed_elem_free_rows AS spare_rows from ndbinfo.memory_per_fragment  
where fq_name like '%simples%';
```

node	frag	alloc_bytes	free_bytes	spare_rows
1	0	0	0	0
1	2	196608	11732	419
2	0	0	0	0
2	2	196608	11732	419
3	1	0	0	0
3	3	0	0	0
4	1	0	0	0
4	3	0	0	0

Oracle MySQL HA & Scaling Solutions

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


When to Consider MySQL Cluster

- Scalability demands
 - Sharding for write performance?
- Latency demands
 - Cost of each millisecond?
- Uptime requirements
 - Cost per minute of downtime?
 - Failure versus maintenance?
- Application agility
 - Developer languages and frameworks?
 - SQL or NoSQL?







MySQL Cluster Auto-Installer








- Fast configuration
- Auto-discovery
- Workload optimized
- Repeatable best practices



Deploy Configuration and start MySQL Cluster

Your MySQL Cluster configuration can be reviewed below. To the left are the processes you have defined, ordered by their startup sequence. Please select a process to view its startup command(s) and configuration file. Note that some processes do not have configuration files. At the bottom of the center panel, there are buttons to *Deploy*, *Start* and *Stop* your cluster. Please note that starting the cluster may take up to several minutes depending on the configuration you have defined. In the process tree, the icons reflect the status of the process as reported by the management daemon:  : *unknown* or if the management daemon does not reply,  : *connected* or *started*,  : *starting* or *shutting down*, and  : *not connected* or *stopped*

MyCluster processes




- Management layer
 -  Management node 1
 -  Management node 2
- Data layer
 -  Multi threaded data node 1
 -  Multi threaded data node 2
- SQL layer
 -  SQL node 1
 -  SQL node 2
 -  SQL node 3

Startup command

Host	blue
Path	/var/mysql/mysql-cluster-gpl-7.3.1-linux-x86_64/
Executable	mysql_install_db

Configuration file

No configuration file for this process

 Deploy cluster  Deploy and start cluster  Stop cluster

MySQL Cluster Manager

Enhancing DevOps Agility, Reducing Downtime



Automated Management

- Start / Stop node or whole cluster
- On-Line Scaling
- On-Line Reconfiguration
- On-Line Upgrades
- On-Line Backup & Restore
- Import Running Cluster

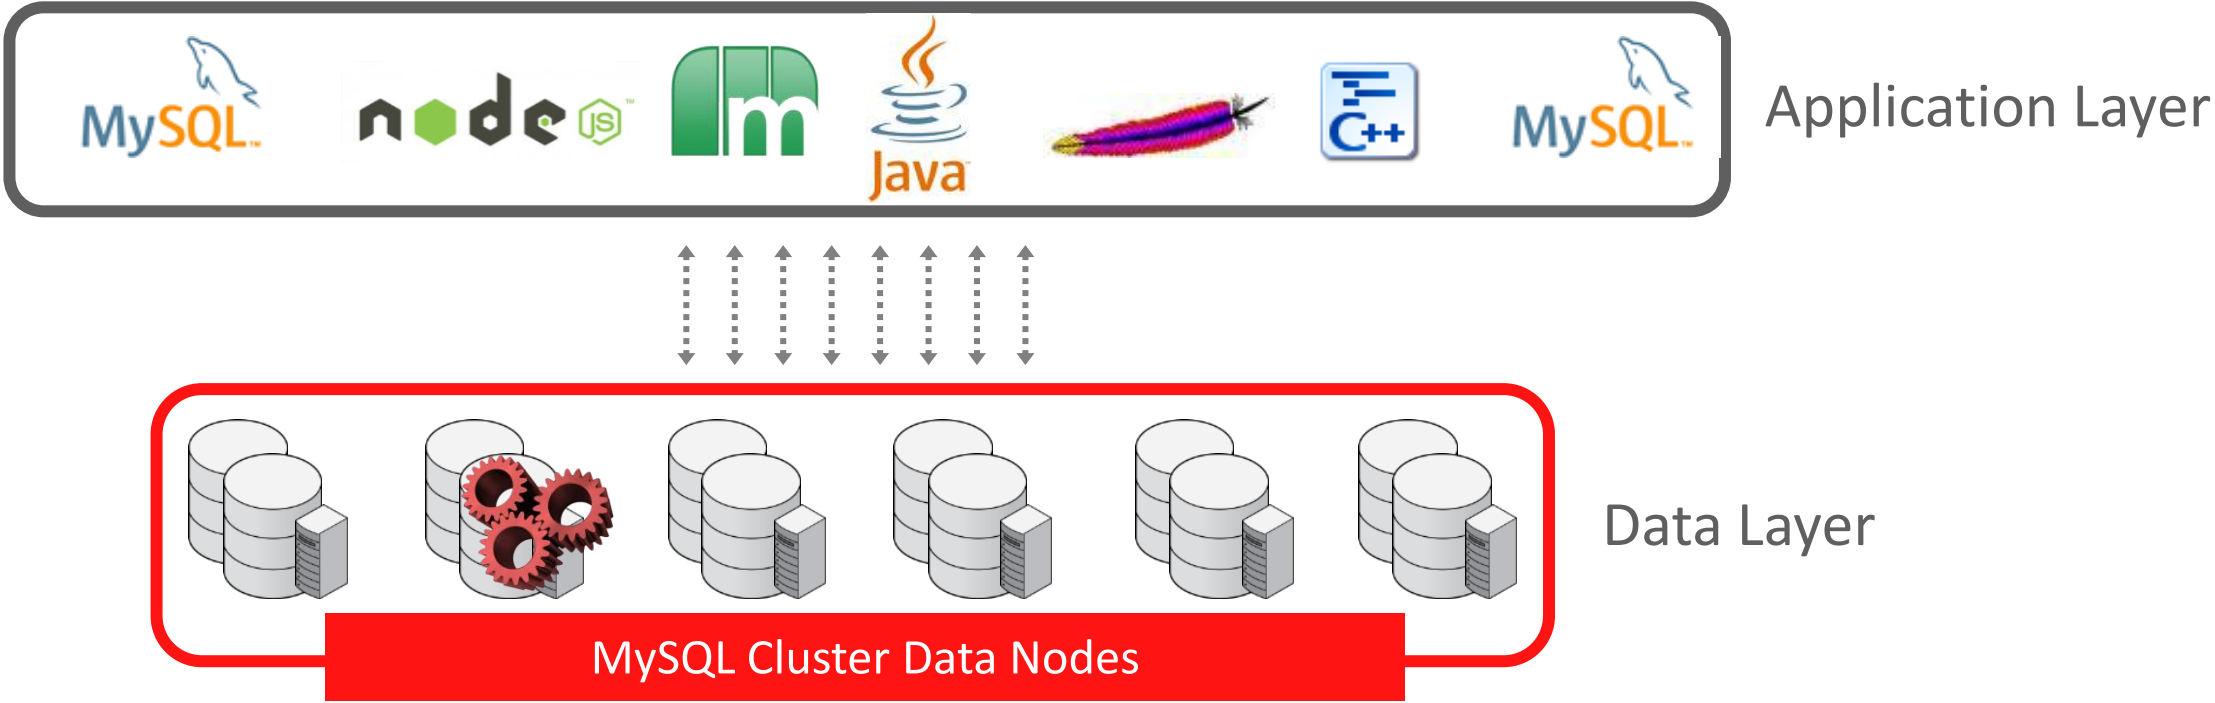
Self-Healing

- Node monitoring
- Auto-recovery extended to SQL + mgmt nodes

HA Operations

- Cluster-wide configuration consistency
- Persistent configurations
- HA Agents

Upgrade using MCM



```
upgrade cluster --package=7.3 mycluster;
```

MySQL Cluster Manager 1.3.2 GA

Available Now!

MySQL Cluster Manager 1.3.2 GA

Import a running Cluster into MCM

“Unmanaged” production Cluster

mcm> create cluster --import

mcm> import config [--dryrun]

mcm> import cluster [--dryrun]

Cluster now managed by MCM

Oracle University MySQL Training Services

Prepare Your Organization to Enable Reliable and High-Performance Web-Based Database Applications

RECENTLY RELEASED

!!ALL NEW!! MySQL Cluster Training

To Register your interest to influence the schedule on this newly released course – go to education.oracle.com/mysql and click on the **MySQL Cluster Course**

“Training and team skill

have the most significant impact on overall performance of technology and success of technology projects.” - IDC, 2013

Premier Support customers eligible to save 20% on learning credits.

Benefits

- Expert-led training to support your MySQL learning needs
- Flexibility to train in the classroom or online
- Hands-on experience to gain real world experience
- Key skills needed for database administrators and developers

Top Courses for Administrators and Developers

- MySQL for Beginners
- MySQL for Database Administrators
- MySQL Performance Tuning
- MySQL Cluster – **NEW - Register Your Interest!**
- MySQL and PHP - Developing Dynamic Web Applications
- MySQL for Developers
- MySQL Developer Techniques

Top Certifications

- MySQL 5.6 Database Administrator
- MySQL 5.6 Developer

To find out more about available MySQL Training & Certification offerings, go to: education.oracle.com/mysql

Next Steps



Learn More

- www.mysql.com/cluster
- Authentic MySQL Curriculum: <http://oracle.com/education/mysql>



Try it Out

- dev.mysql.com/downloads/cluster/



Let us know what you think

- bugs.mysql.com
- forums.mysql.com/list.php?25