Scaling Zabbix with MySQL InnoDB Cluster

Vittorio Cioe
MySQL Sr. Sales Consultant
vittorio.cioe@oracle.com

Introduction

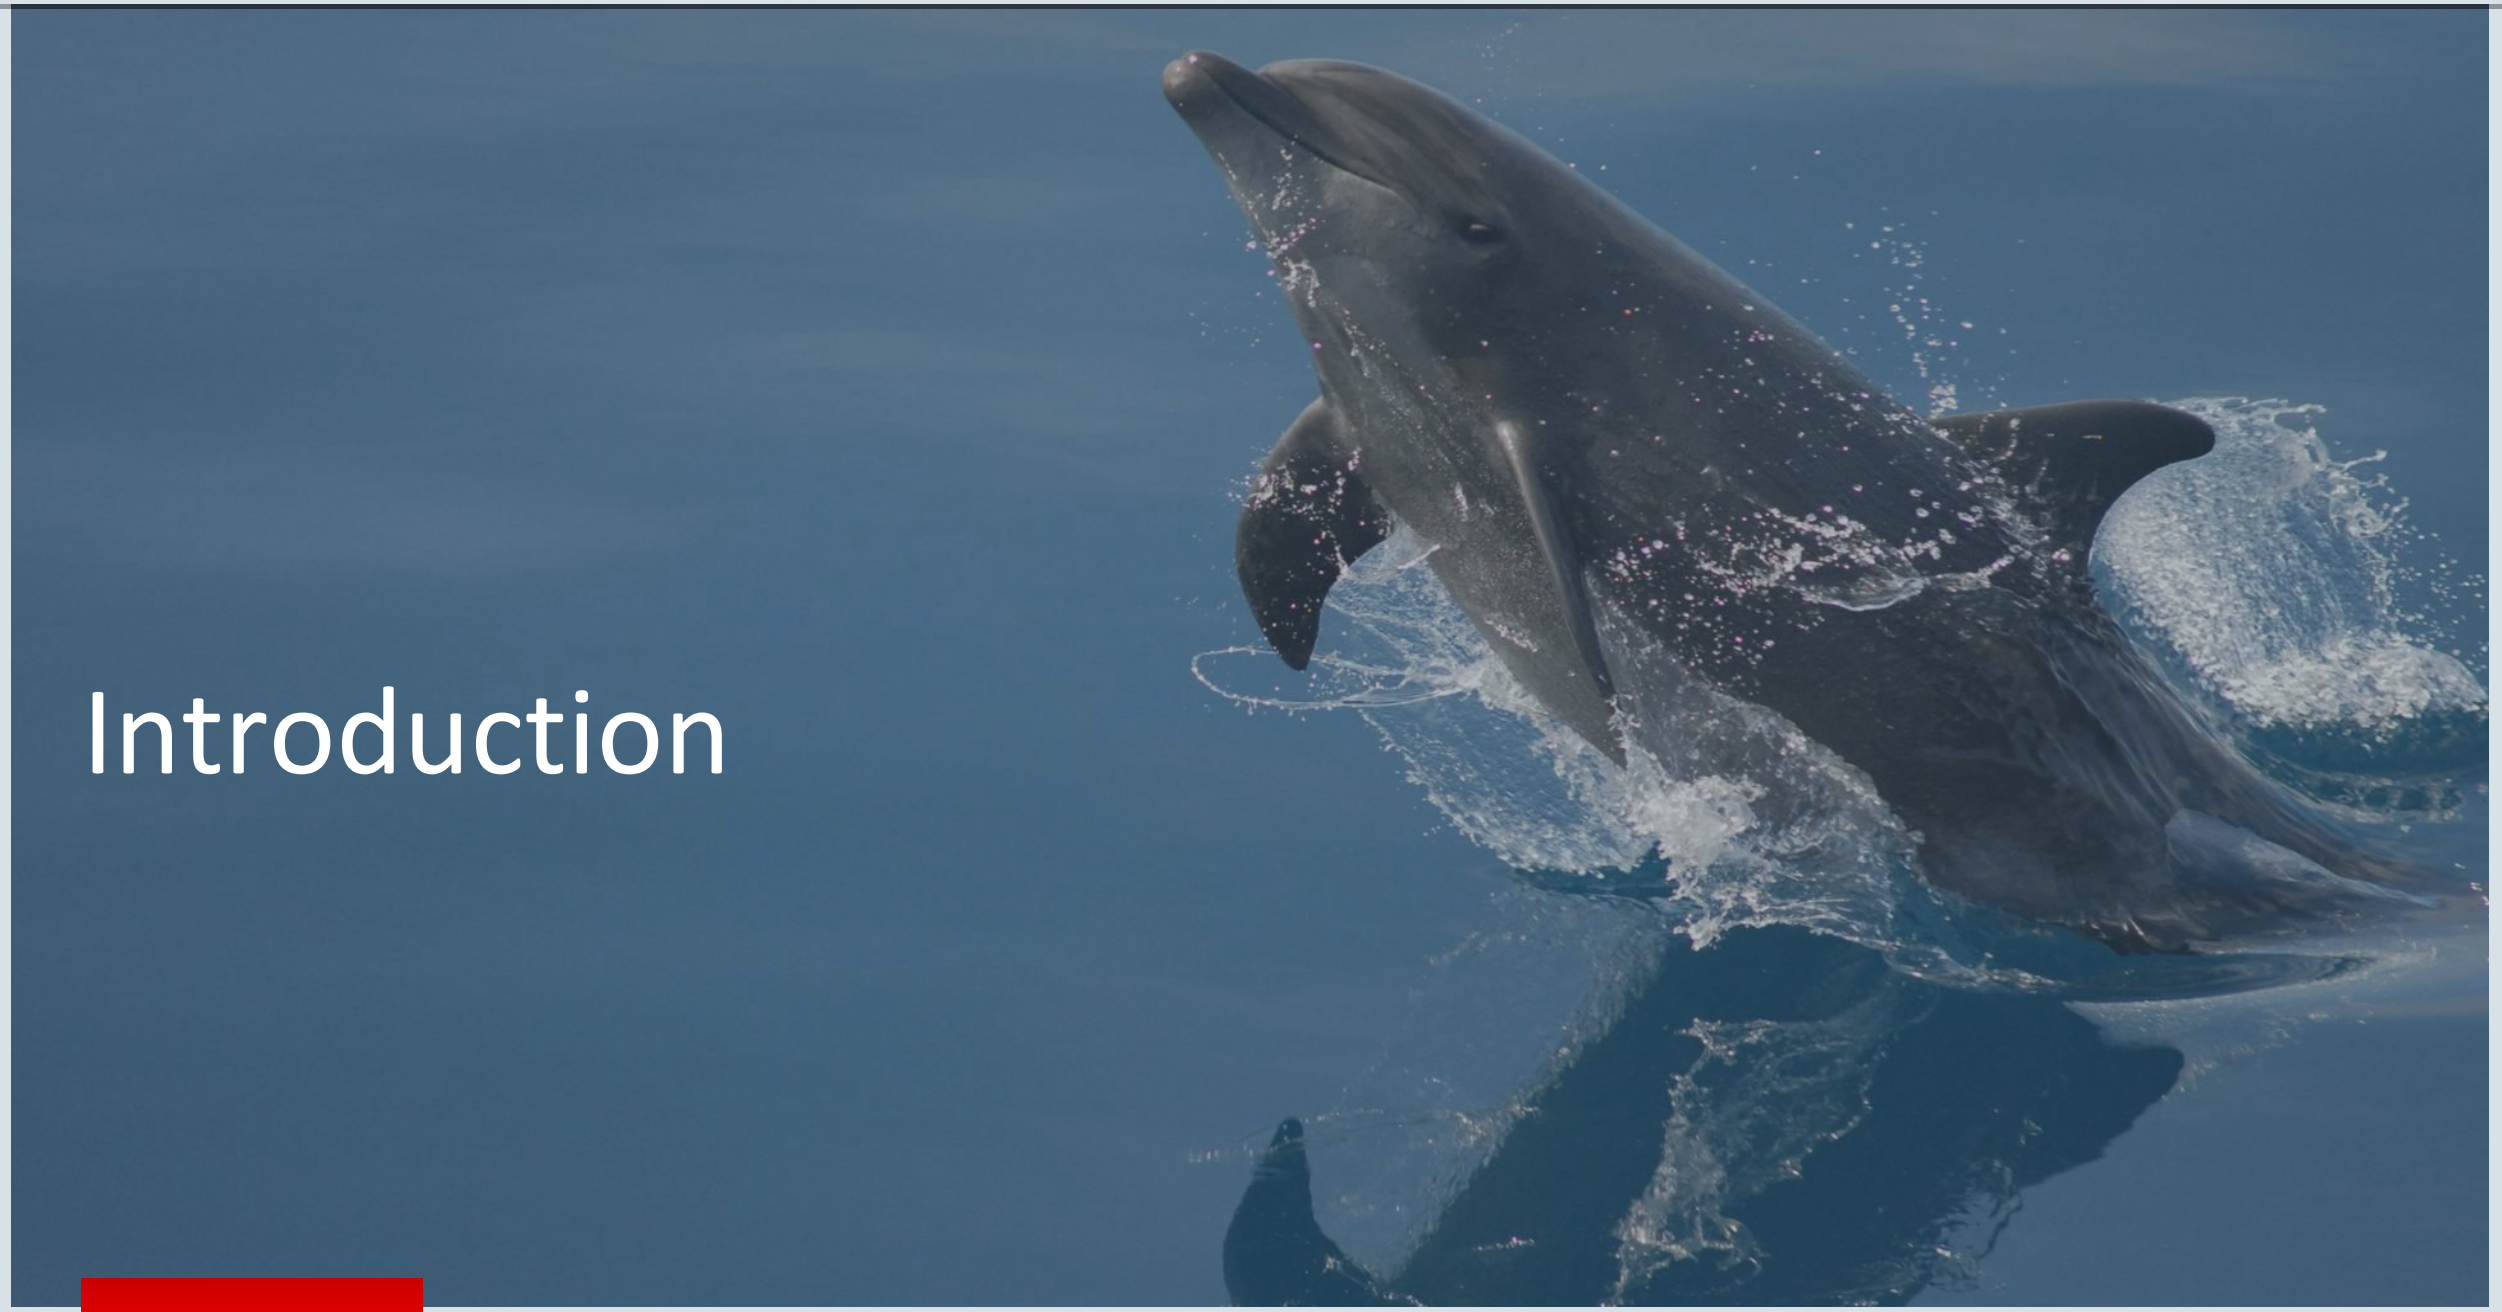

100%

All organizations require their most critical systems to be **highly available**

99.9999...% Availability: Monitoring is the key

Does Zabbix needs to be highly available?

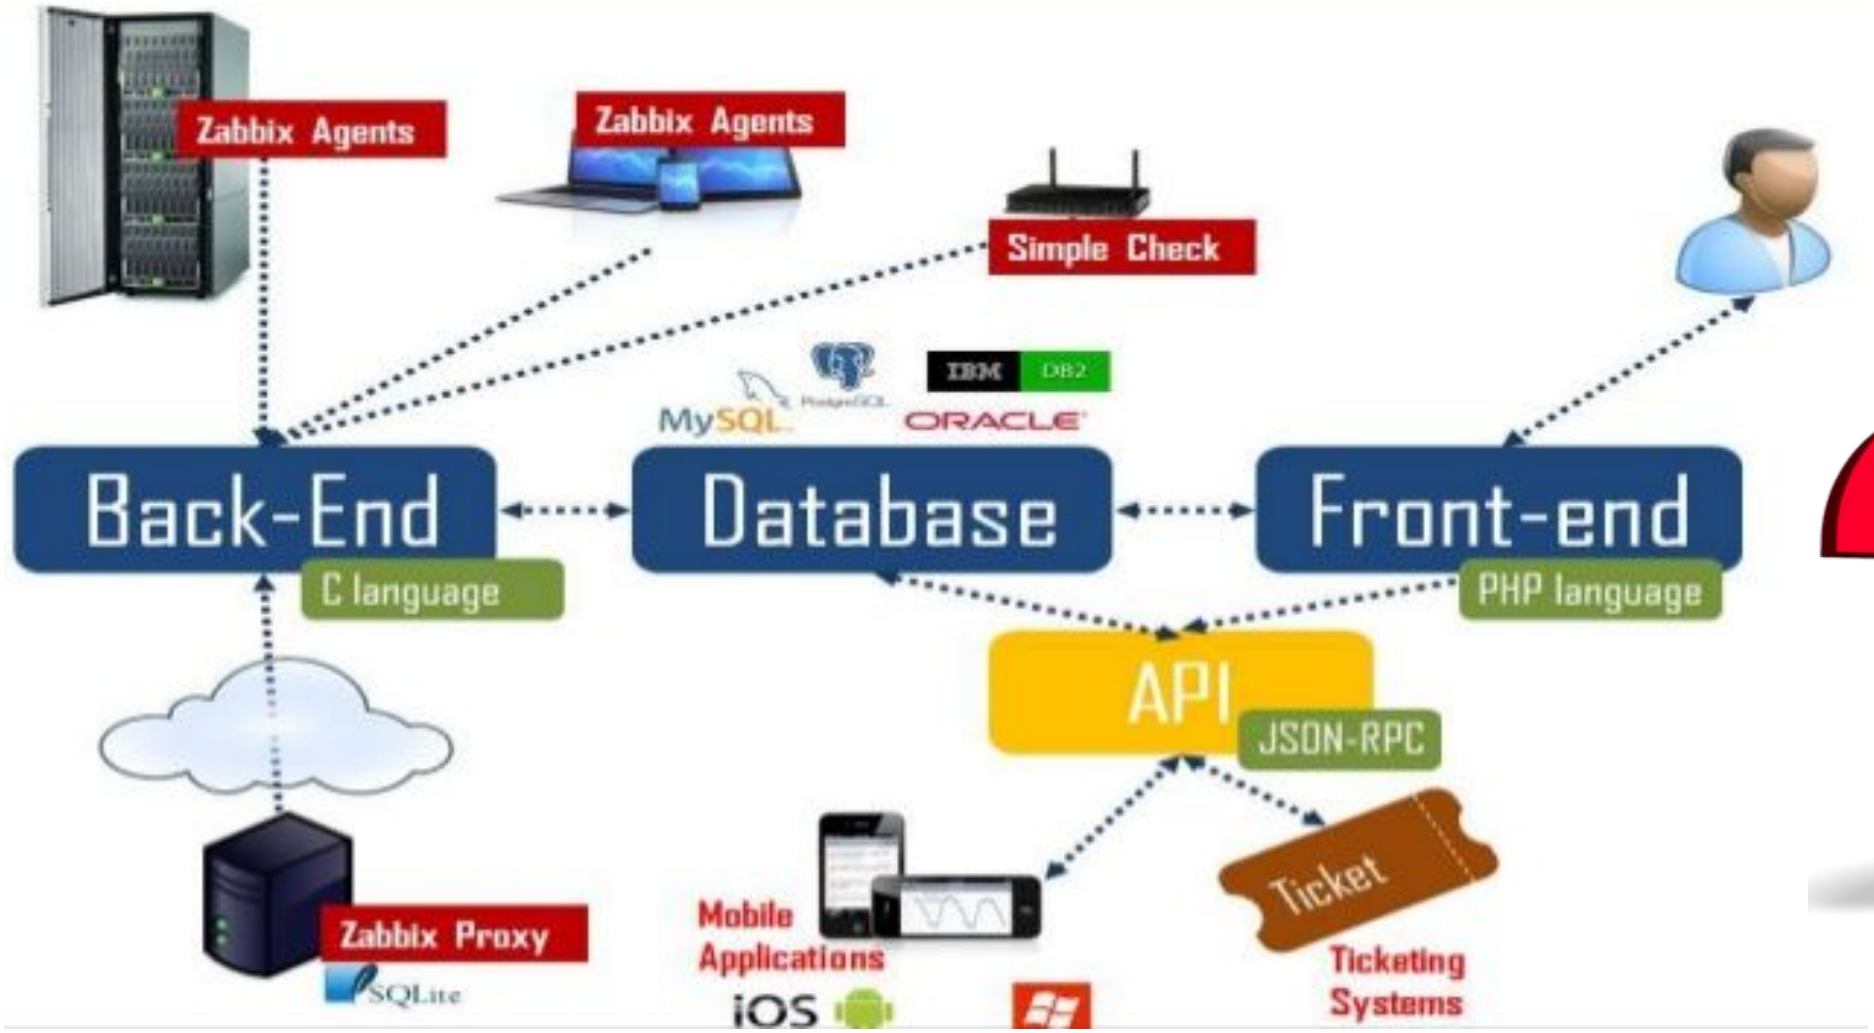

Application HA vs Database HA

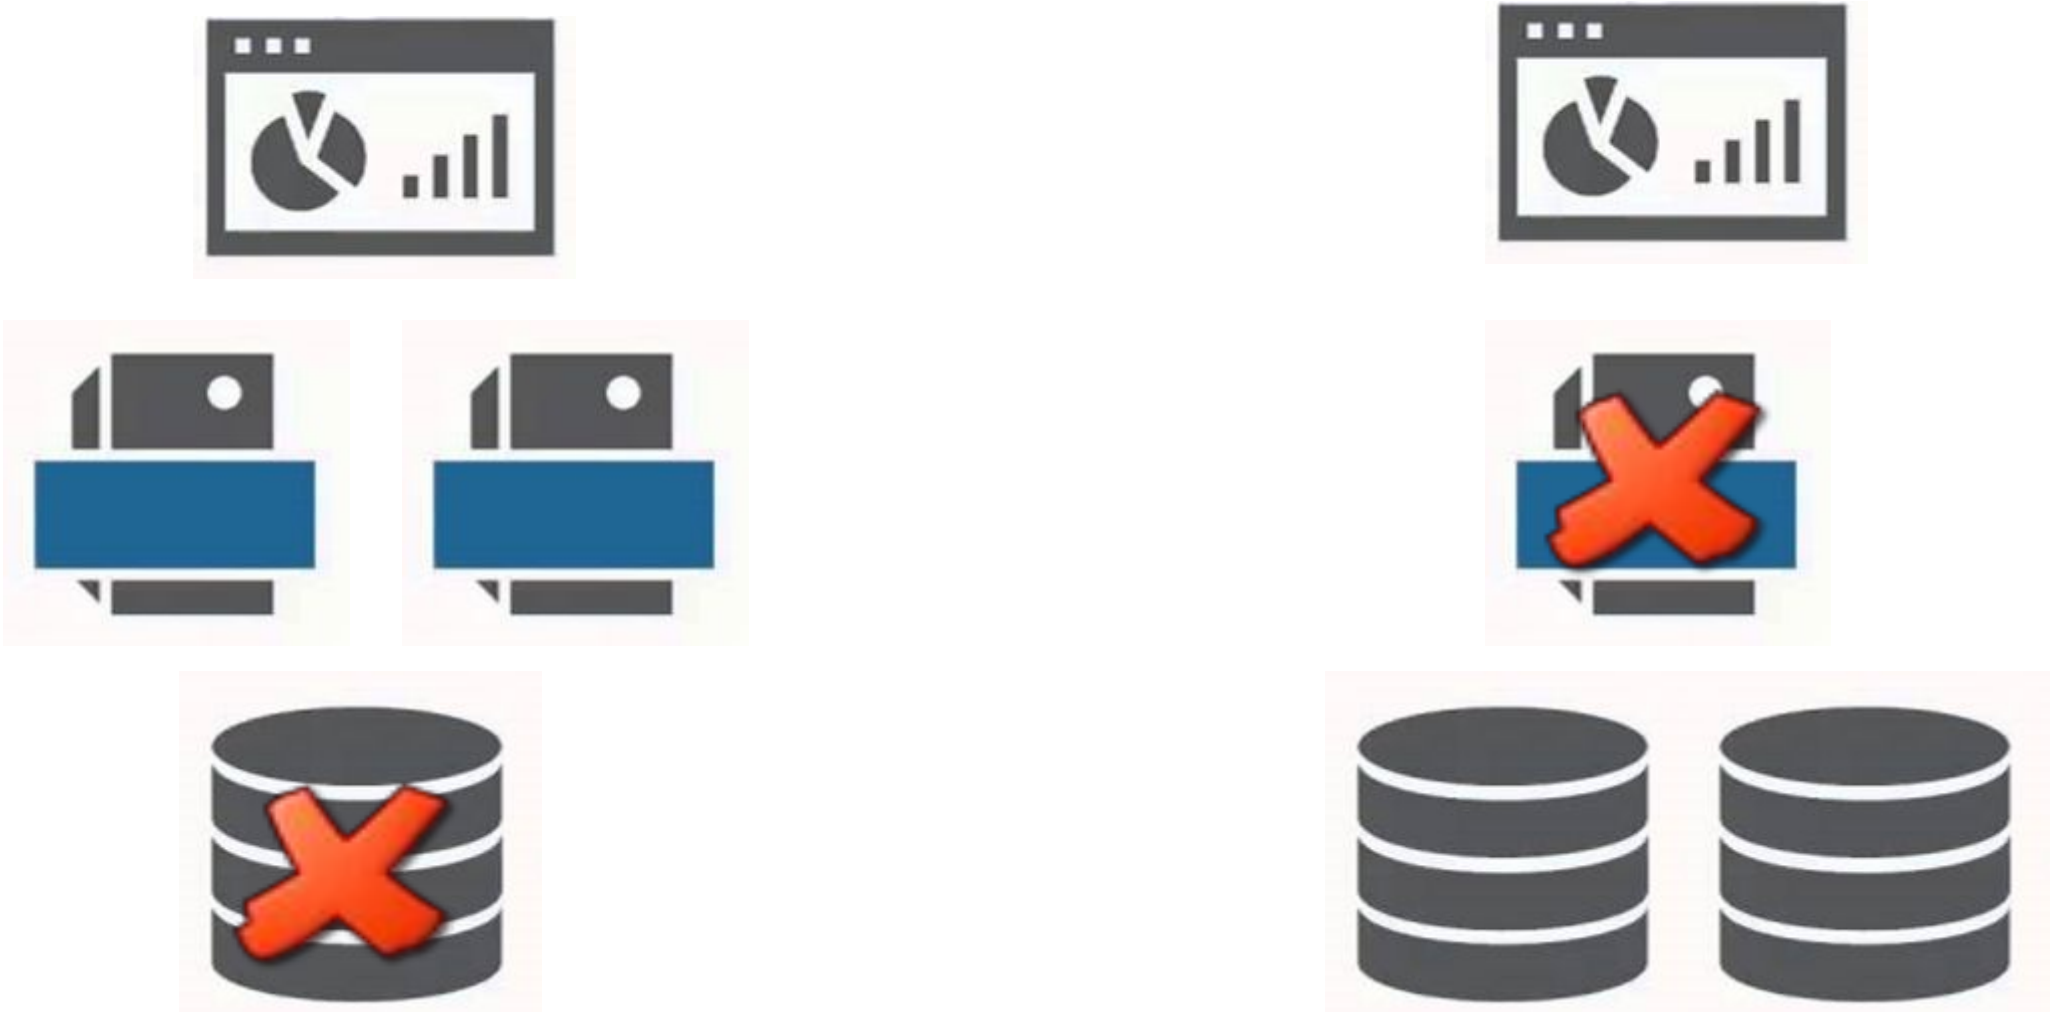


#1

Zabbix supports MySQL, the #1 database for the web, used by 10 of the top 10 websites

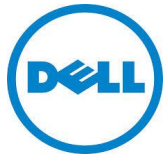

ORACLE

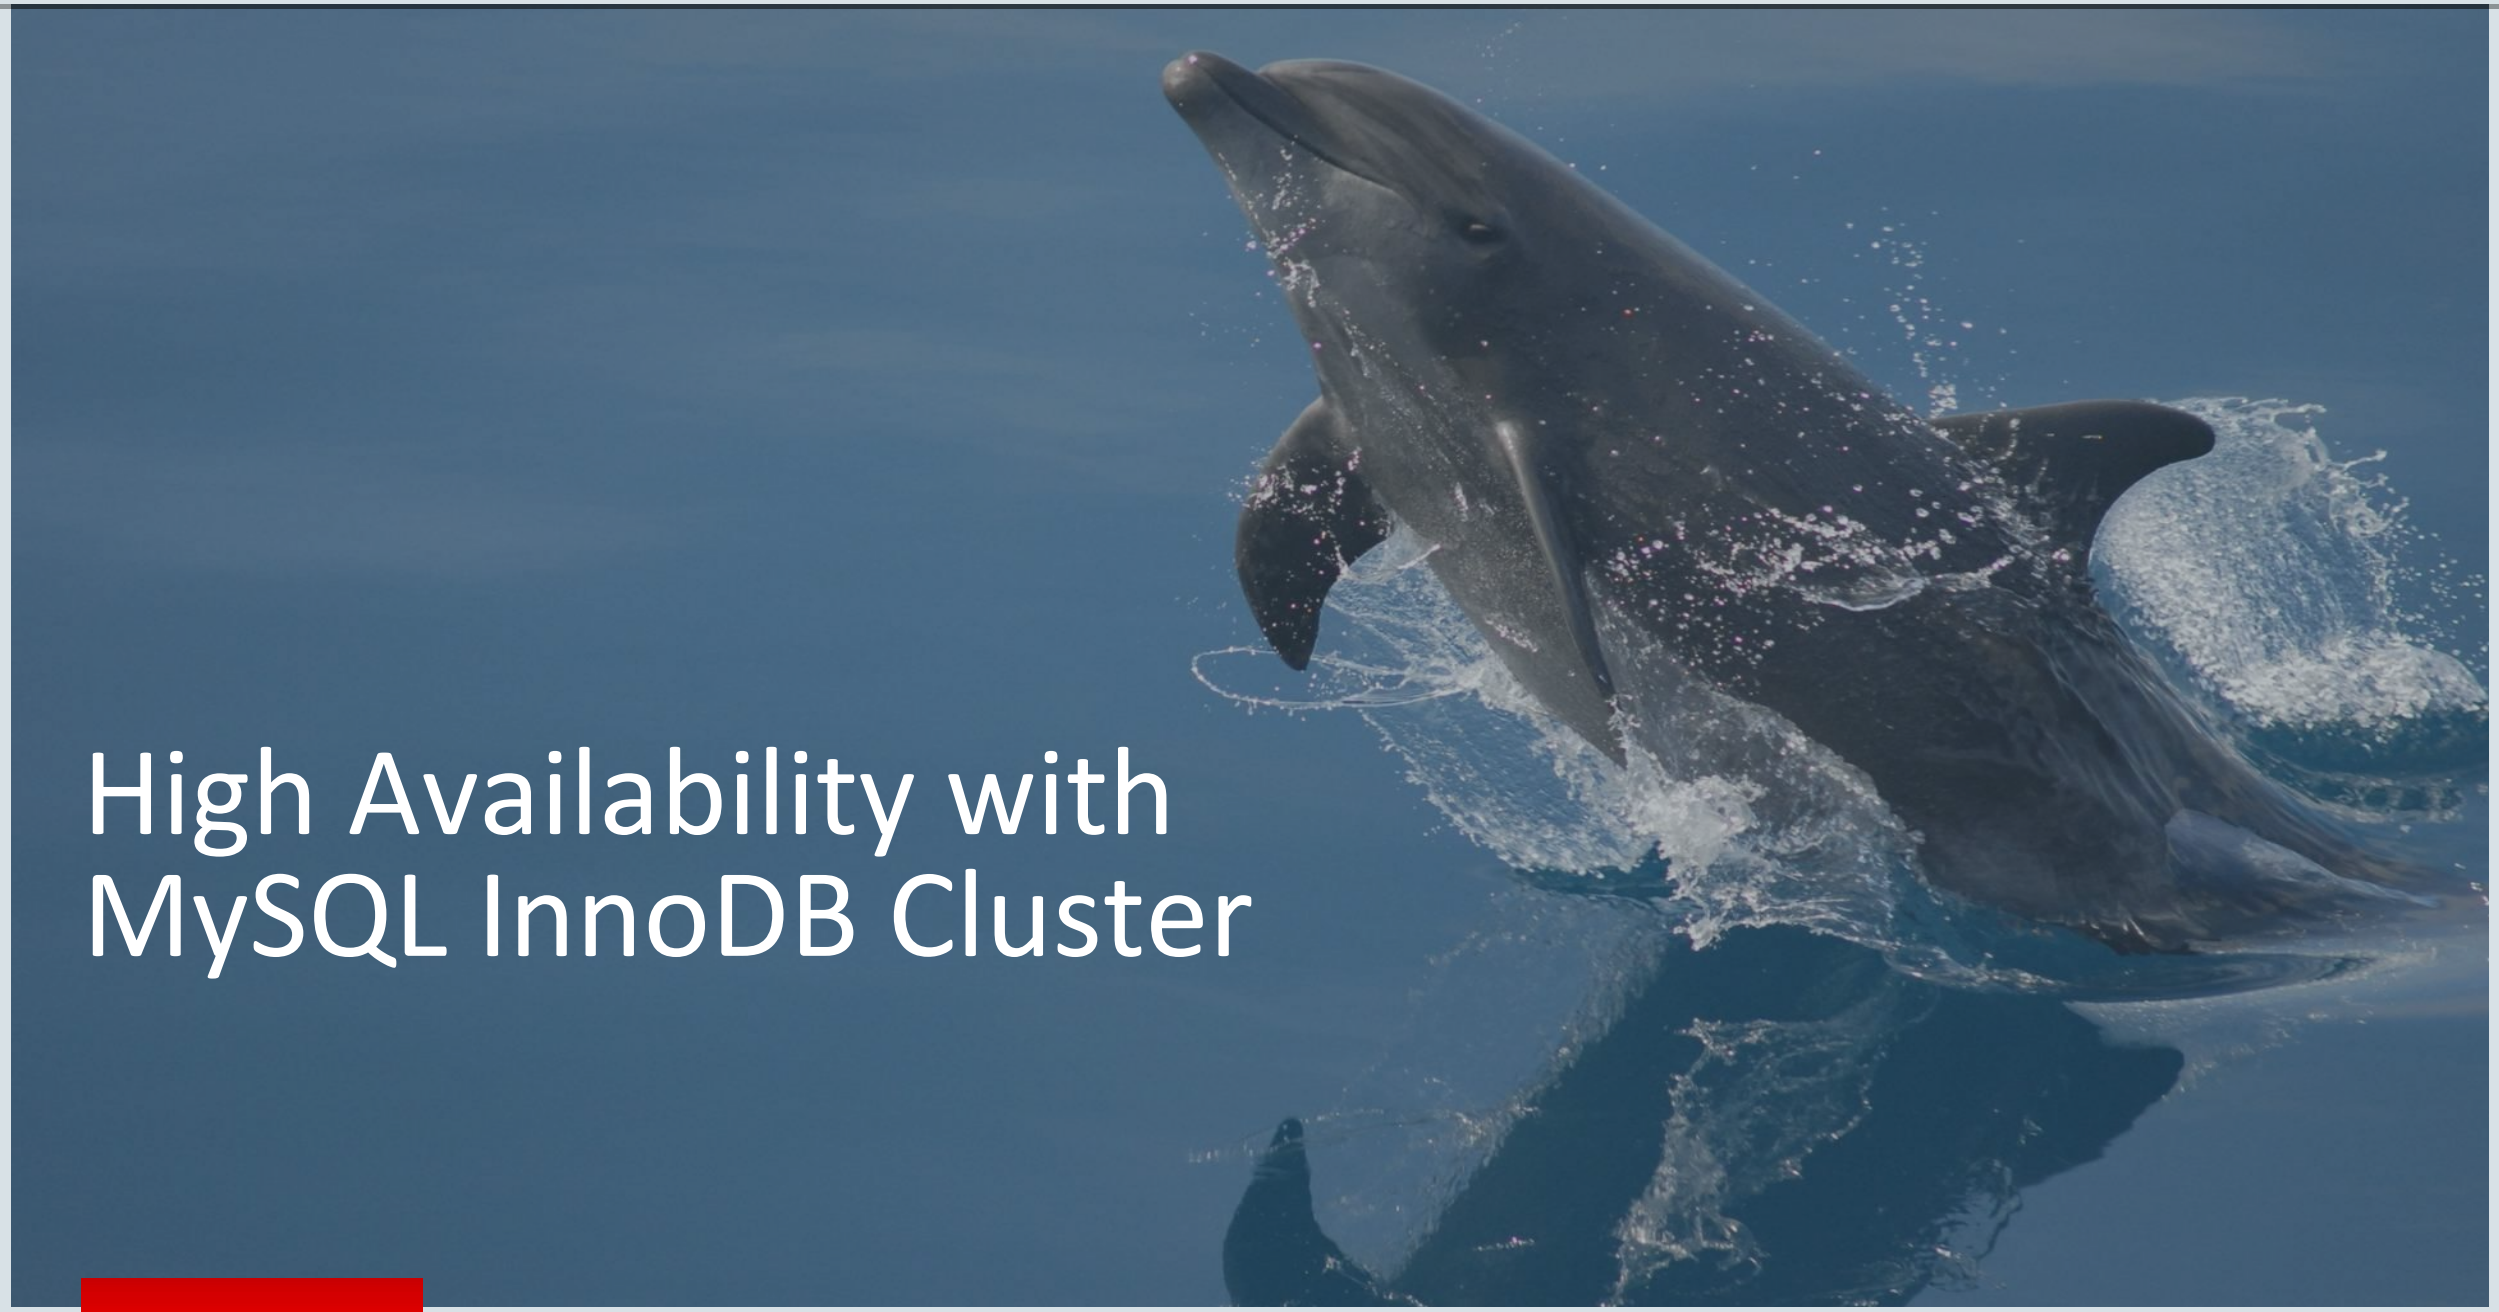

High Availability with MySQL InnoDB Cluster

ORACLE

Yesterday: Asyncronous replication

Read scale-out

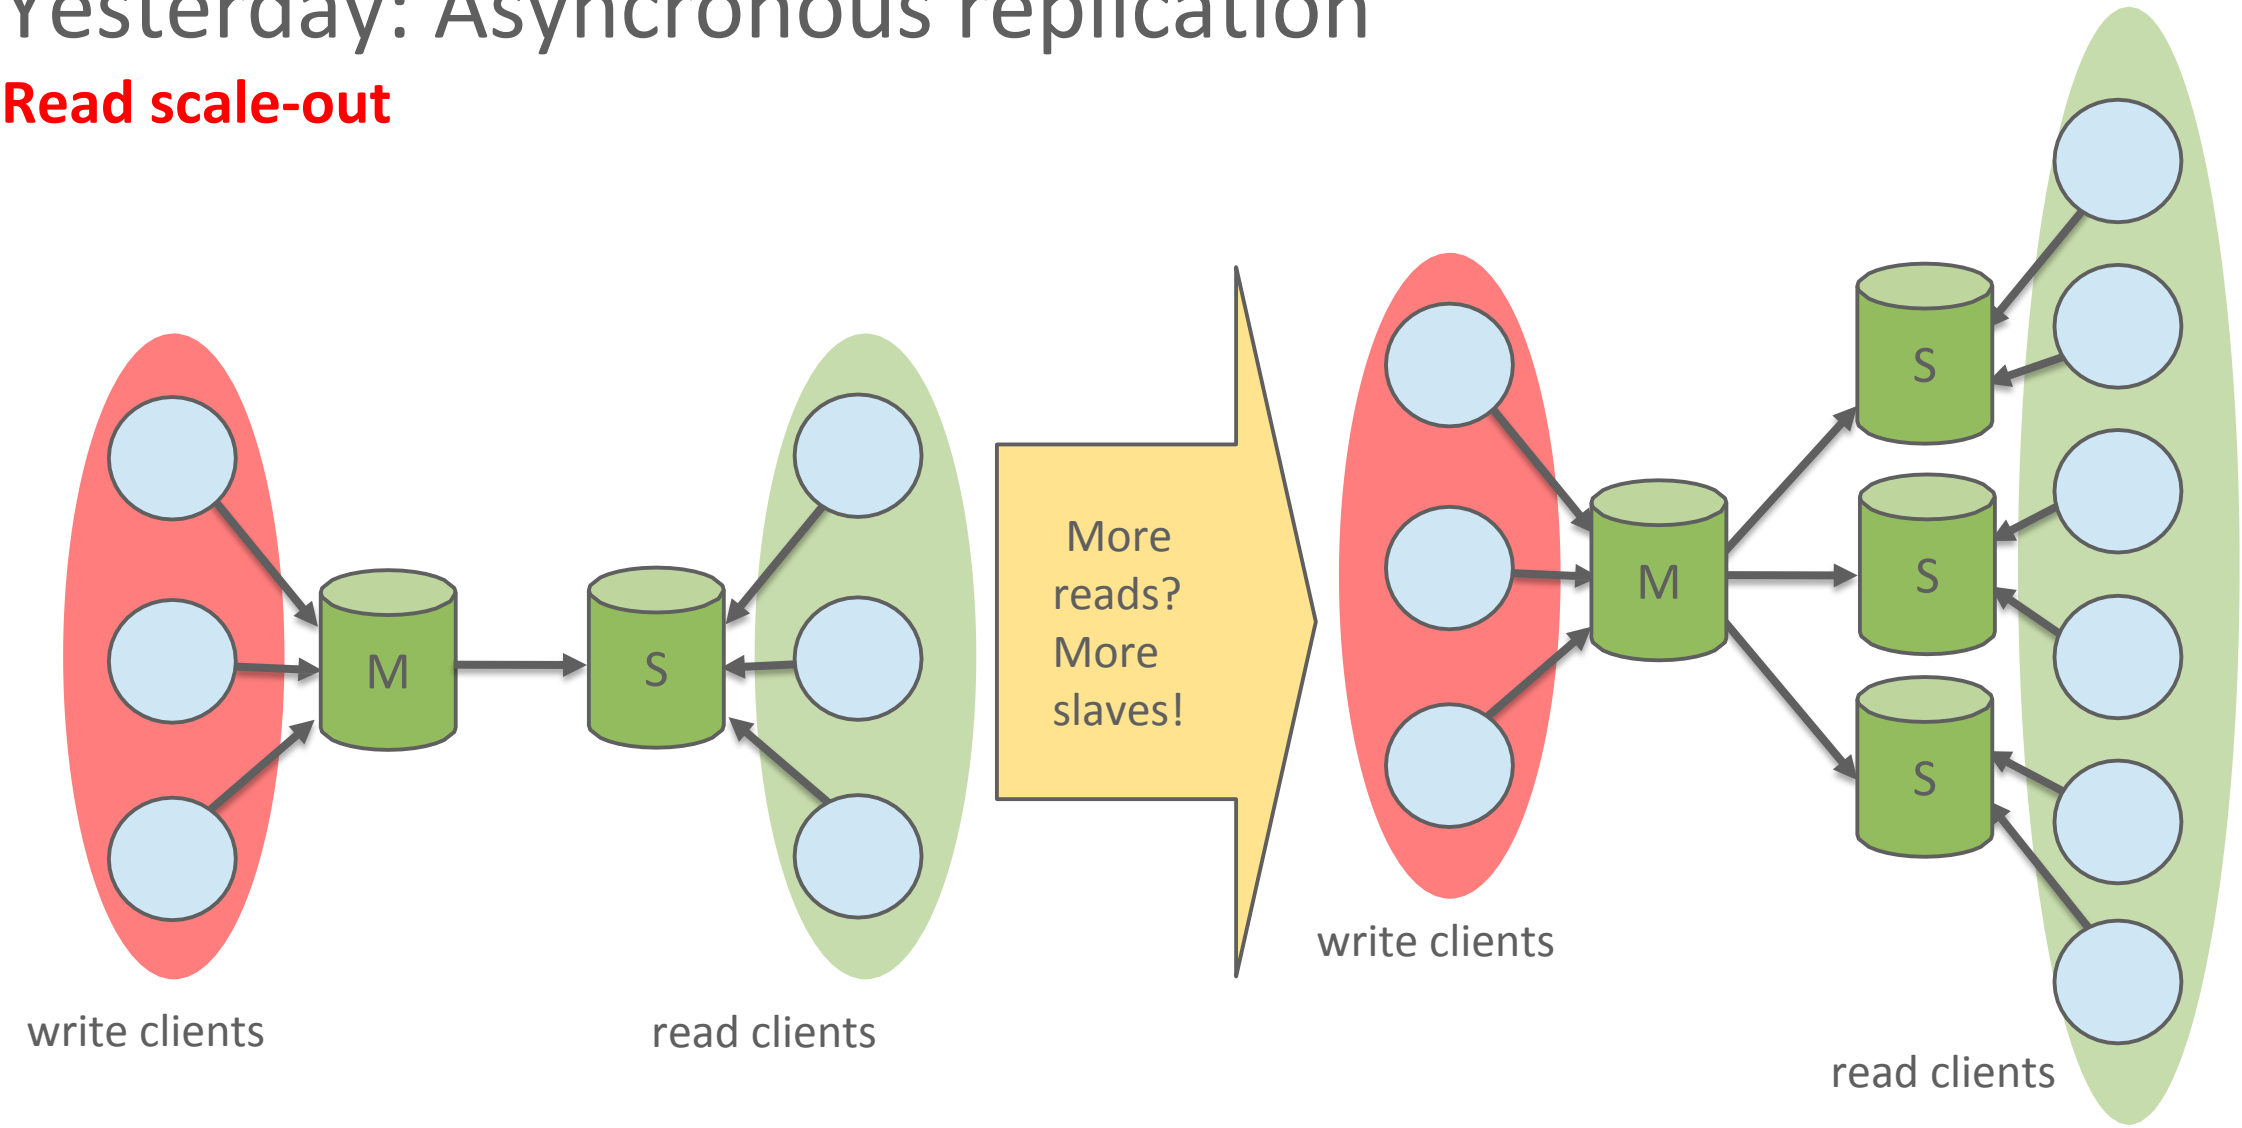

Yesterday: Asynchronous Replication Challenges

Redundancy as a major building block for high availability: If master crashes, promote slave to master

Limitations: failover and conflict detection should be handled manually or handled at application level

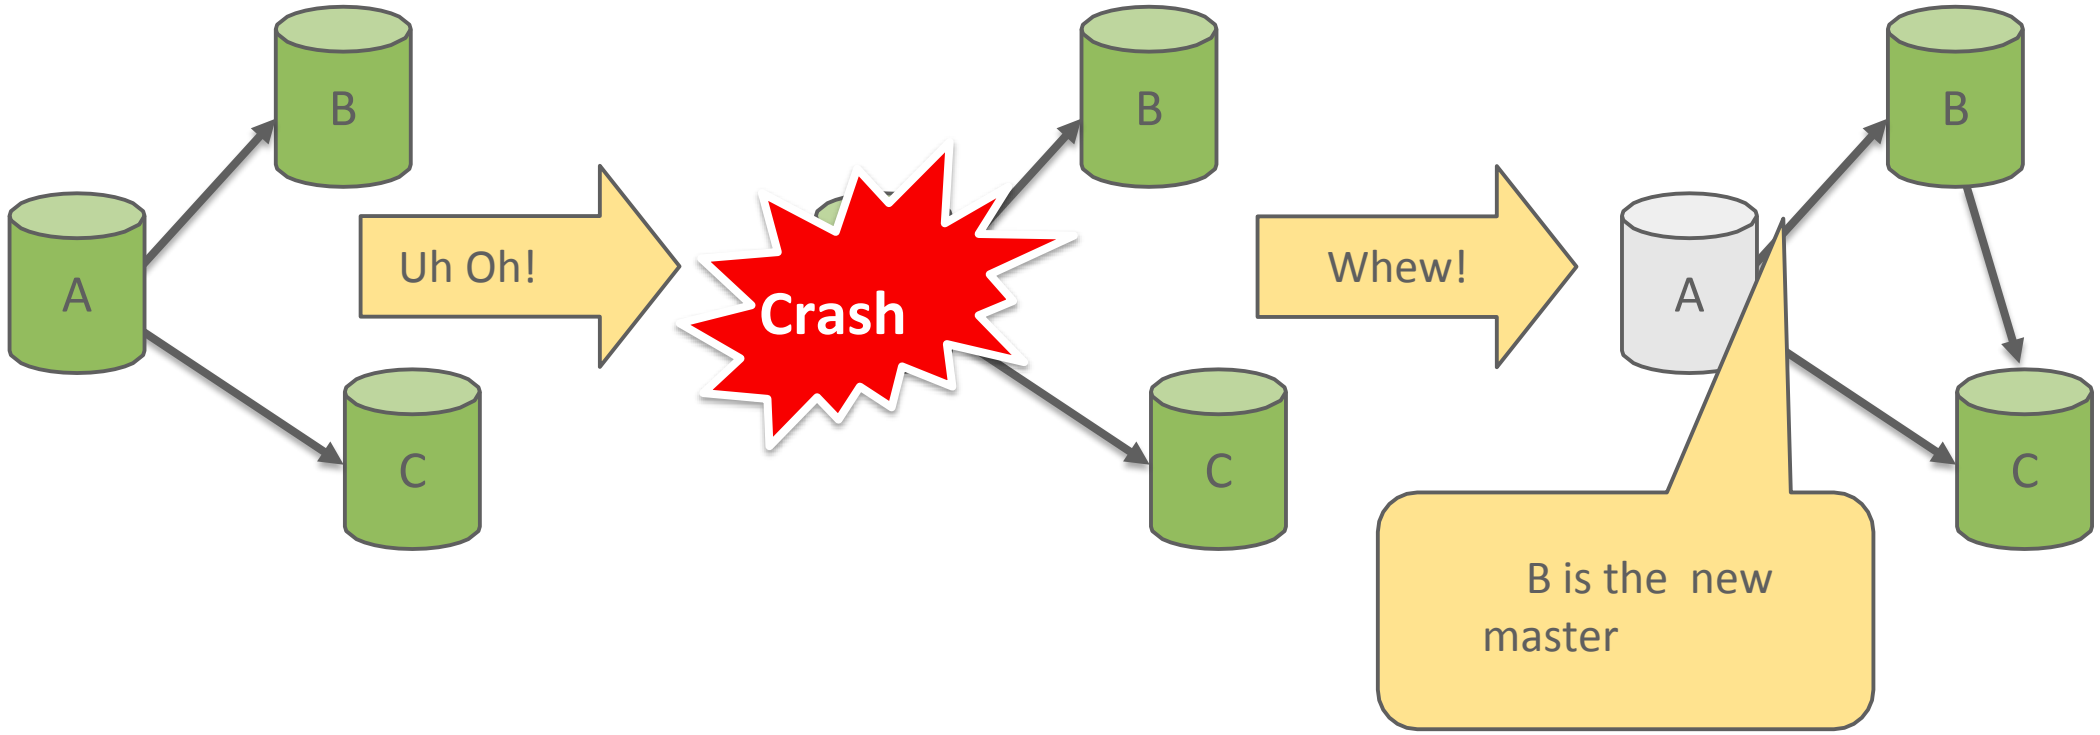

Today: MySQL InnoDB Cluster

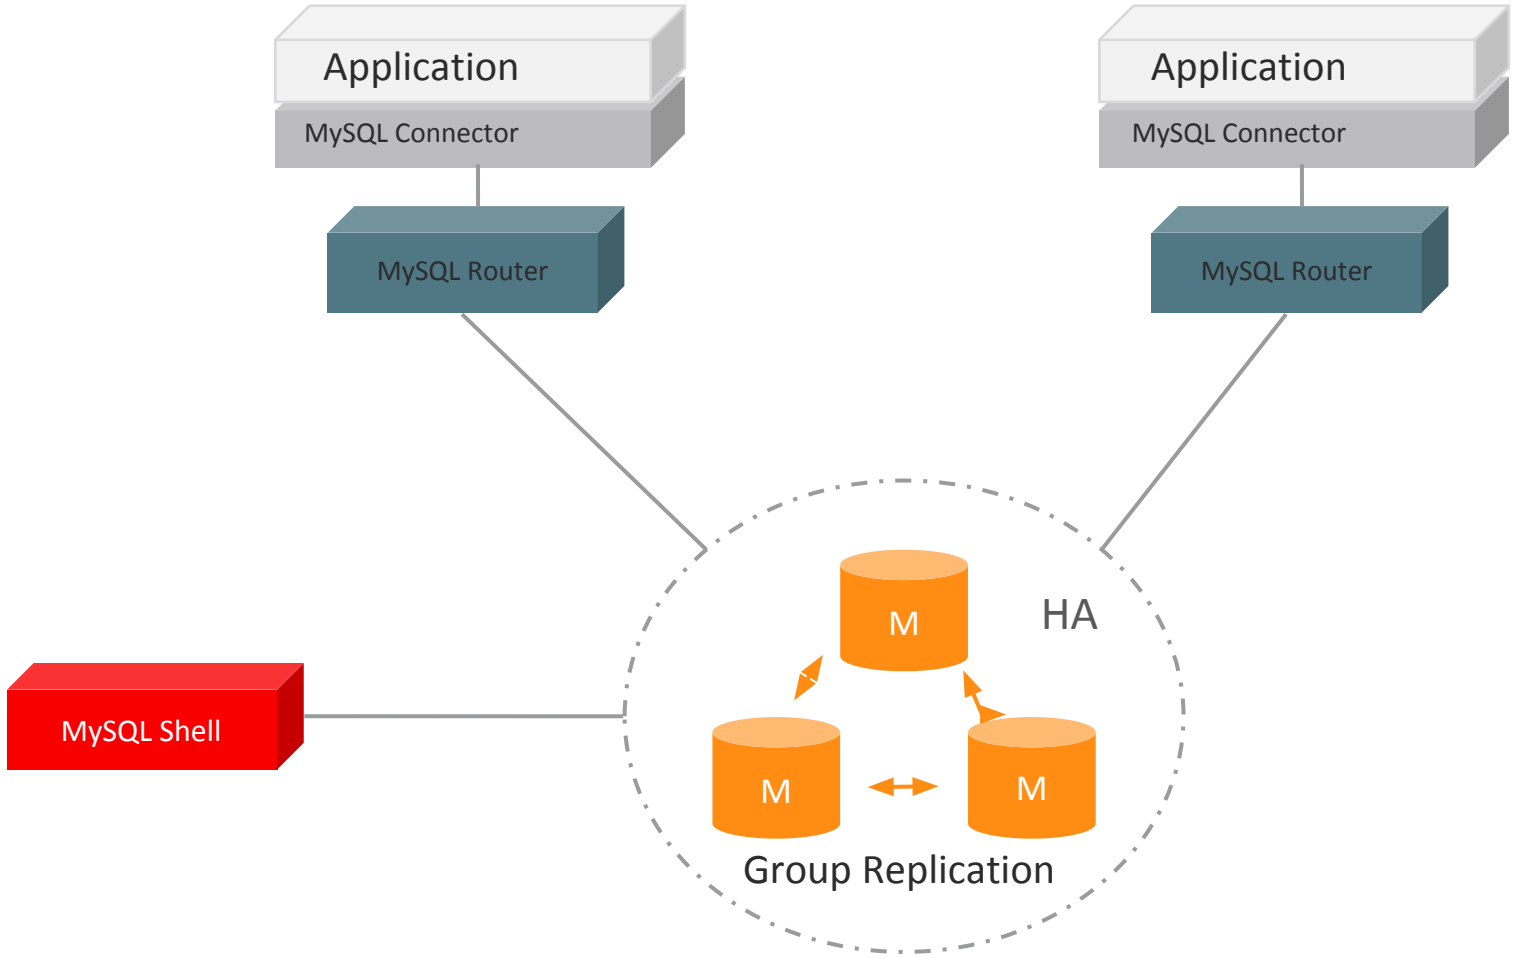

MySQL Router

- Native support for MySQL InnoDB clusters
  - Understands Group Replication topology
  - Utilizes metadata schema stored on each member
    - Bootstraps itself and sets up client routing for the GR cluster
    - Allows for intelligent client routing into the GR cluster
    - Supports multi-master and single primary modes

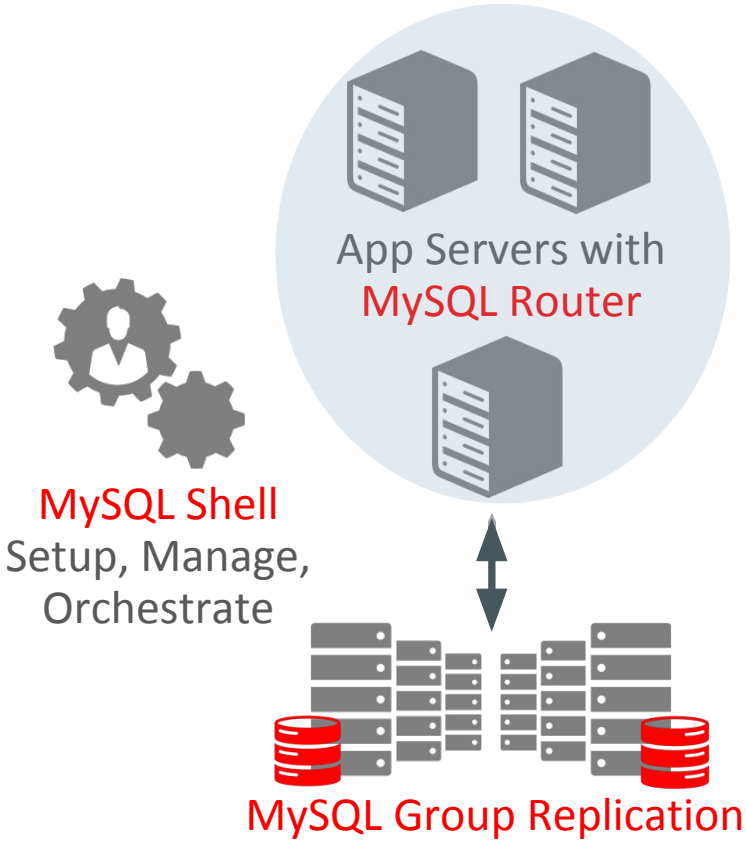

MySQL Shell

A single unified client for all administrative and operations tasks

- Advanced command-line client and code editor for the MySQL Server
  - Supports **development** & **administration** for the MySQL Server
  - Can be used to perform data queries/update & administration operations

- Interactive multi-language: **JavaScript**, **Python**, and **SQL**
  - Naturally scriptable (with development & administrative **APIs**)
  - Both interactive and batch operations

- Exposes full Admin API to manage InnoDB Cluster:
  - create, configure, modify, validate, monitor... and script!!

The core component: MySQL Group Replication

- **What is MySQL Group Replication?**

  "**Update everywhere** replication plugin for MySQL with built-in **automatic distributed recovery**, **conflict handling**, **group membership** and **distributed agreement**."

- **What does the MySQL Group Replication plugin do for the user?**
  - Removes the need for handling server fail-over.
  - Provides fault tolerance.
  - Enables update everywhere setups.
  - Automatic distributed coordination (protects against split-brain and message loss).
  - Less admin overhead, means more fun time!

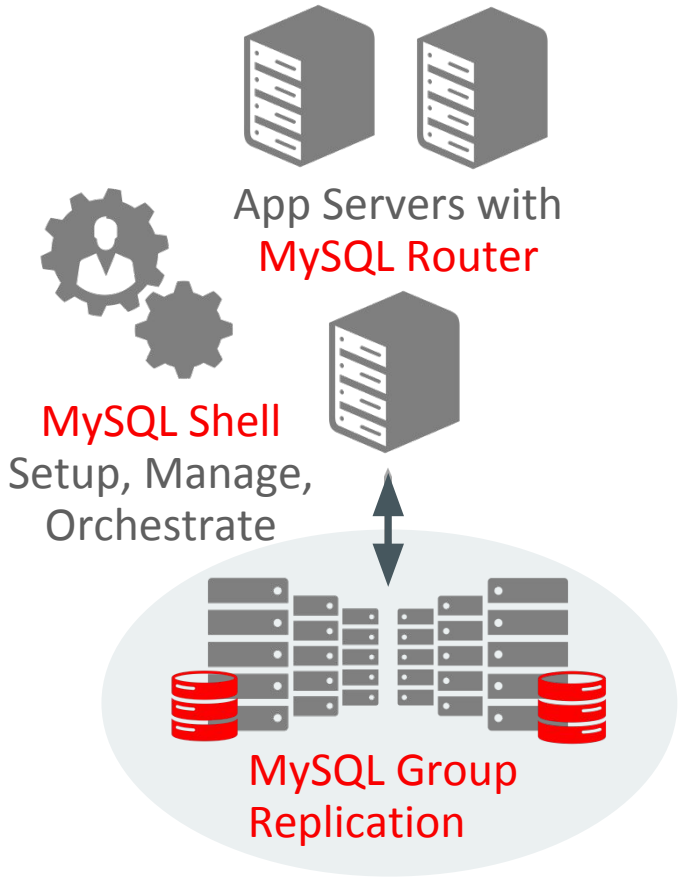

MySQL Group Replication: Architecture

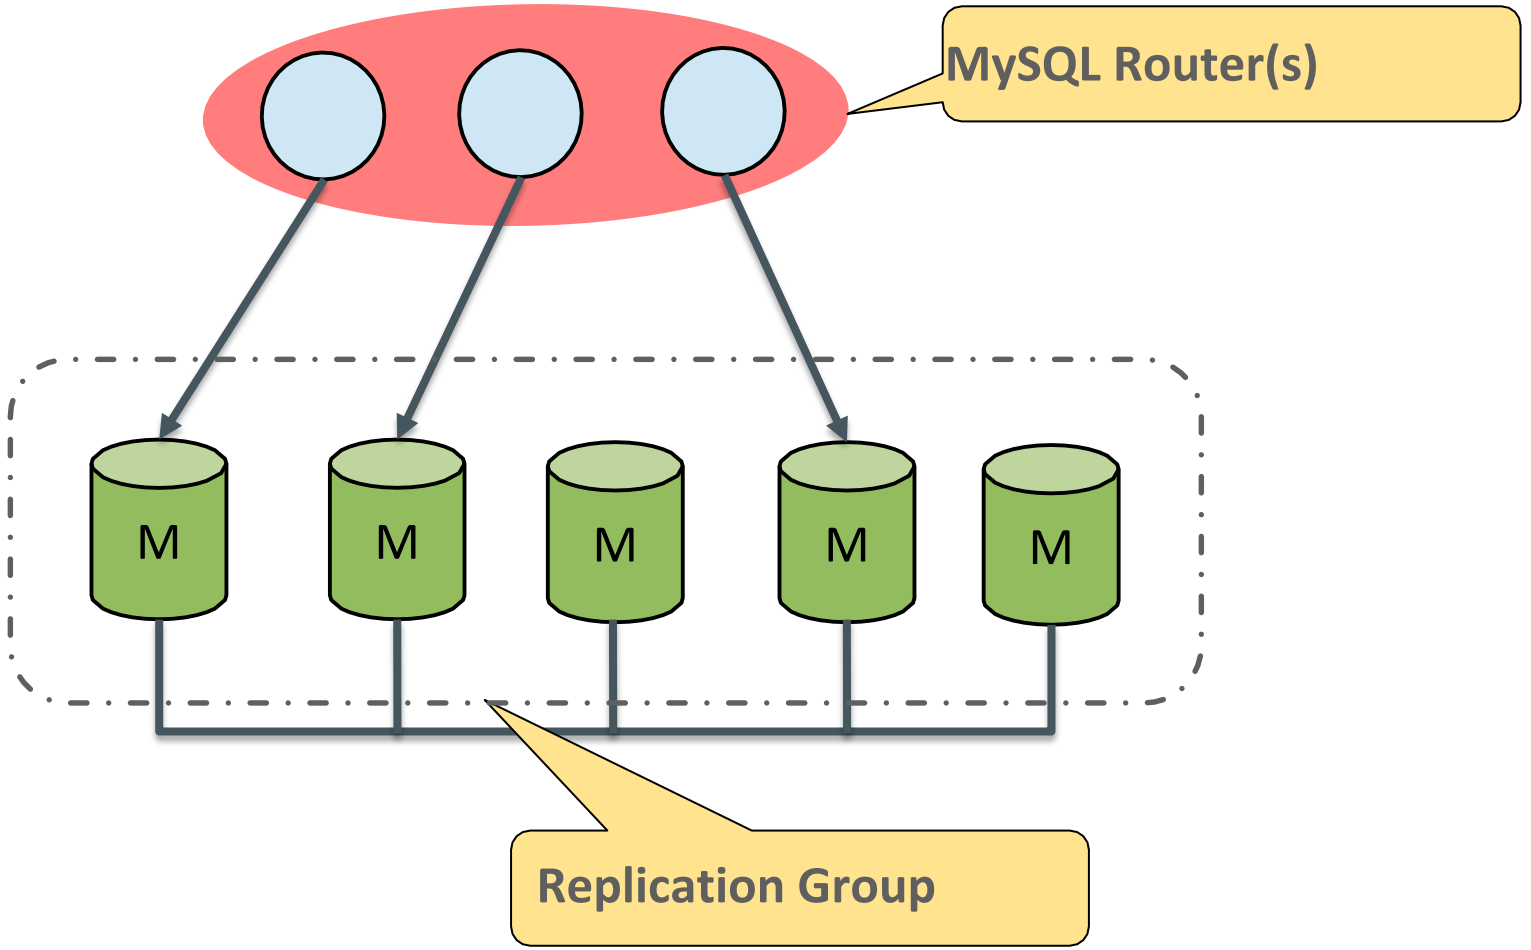

MySQL Group Replication: Some Theory Behind It...

- Implementation based in Replicated Database State Machines
  - Group Communication Primitives resemble properties of Databases.

- Deferred update replication: **propagate atomically**, **check conflicts**, **eventually apply**
  - Distributed state machine requires agreed delivery – implies total order;
  - Deterministic certification requires total order delivery.

2 Working modes: **SINGLE Primary** and **MULTI Primary**

MySQL Group Replication: Single Primary Mode

Configuration mode that makes a single member act as a writable master (PRIMARY) and the rest of the members act as hot-standbys (SECONDARY).

- **Failover**: the group itself coordinates automatically to figure out which is the member that will act as the PRIMARY, through a leader election mechanism.

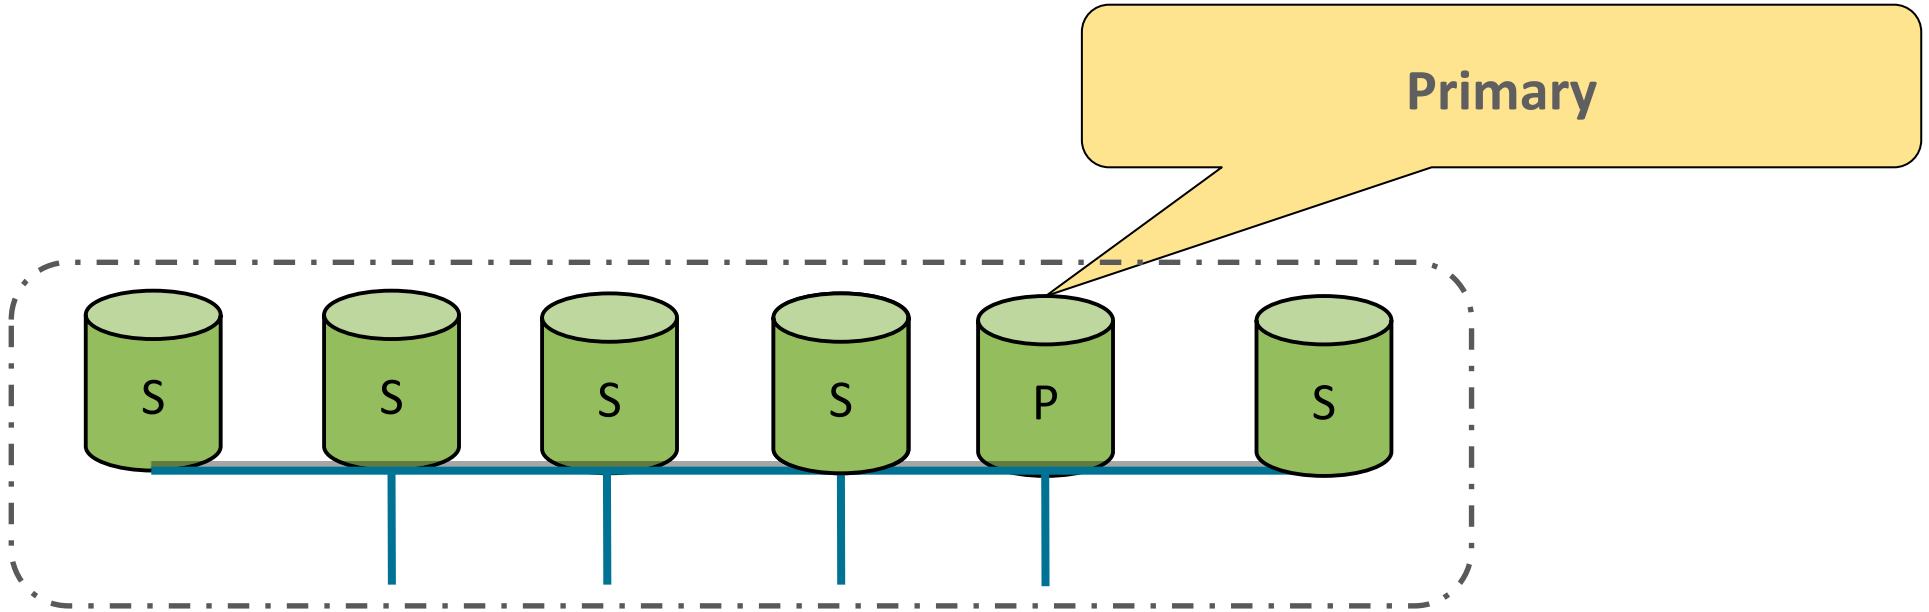

MySQL Group Replication: Single Primary Mode

Configuration mode that makes a single member act as a writable master (PRIMARY) and the rest of the members act as hot-standbys (SECONDARY).

- **Failover**: the group itself coordinates automatically to figure out which is the member that will act as the PRIMARY, through a leader election mechanism.

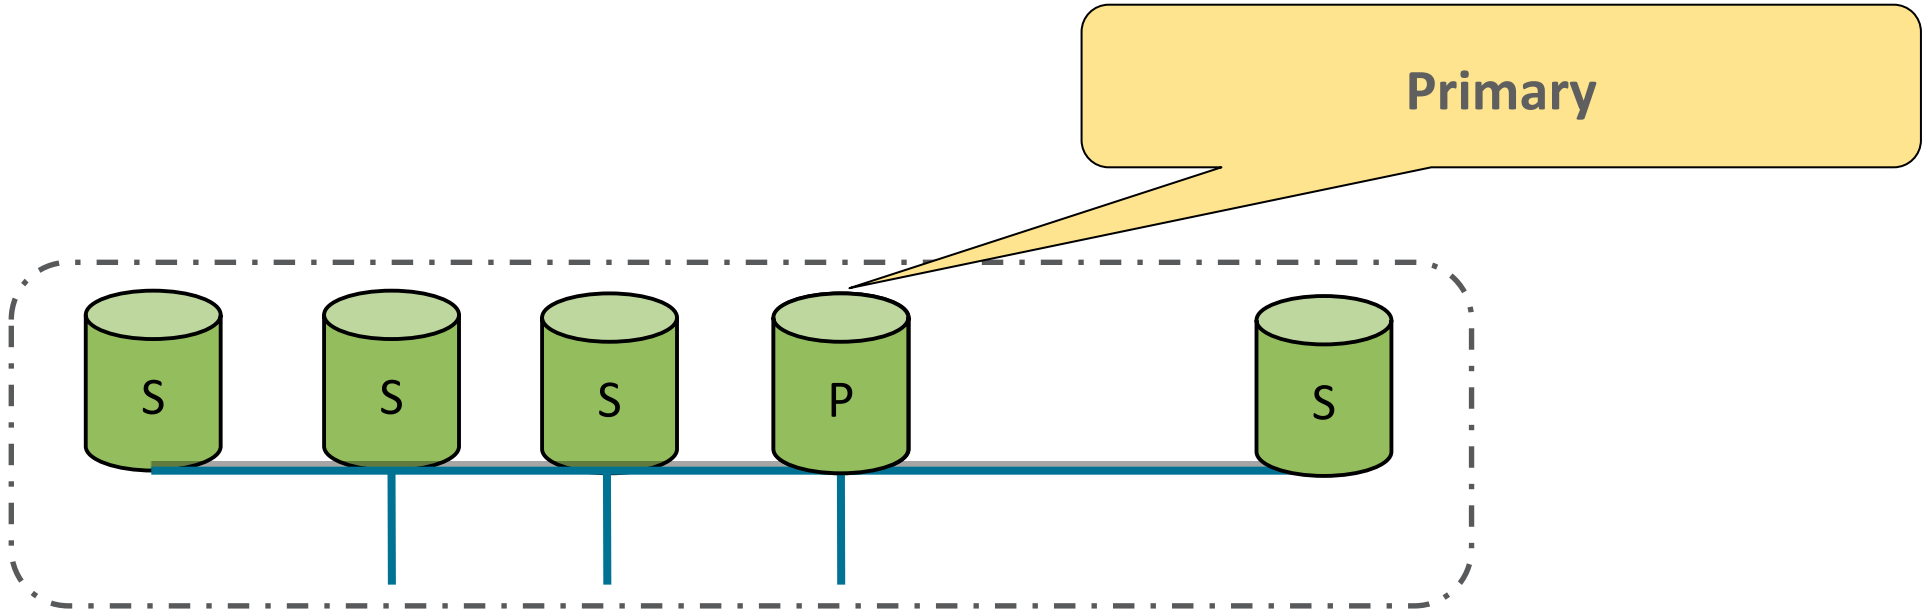

MySQL Group Replication: Automatic distributed recovery!

- Server that (re)joins the group will automatically synchronize with the others.
- If a server leaves the group, the others will automatically be informed.

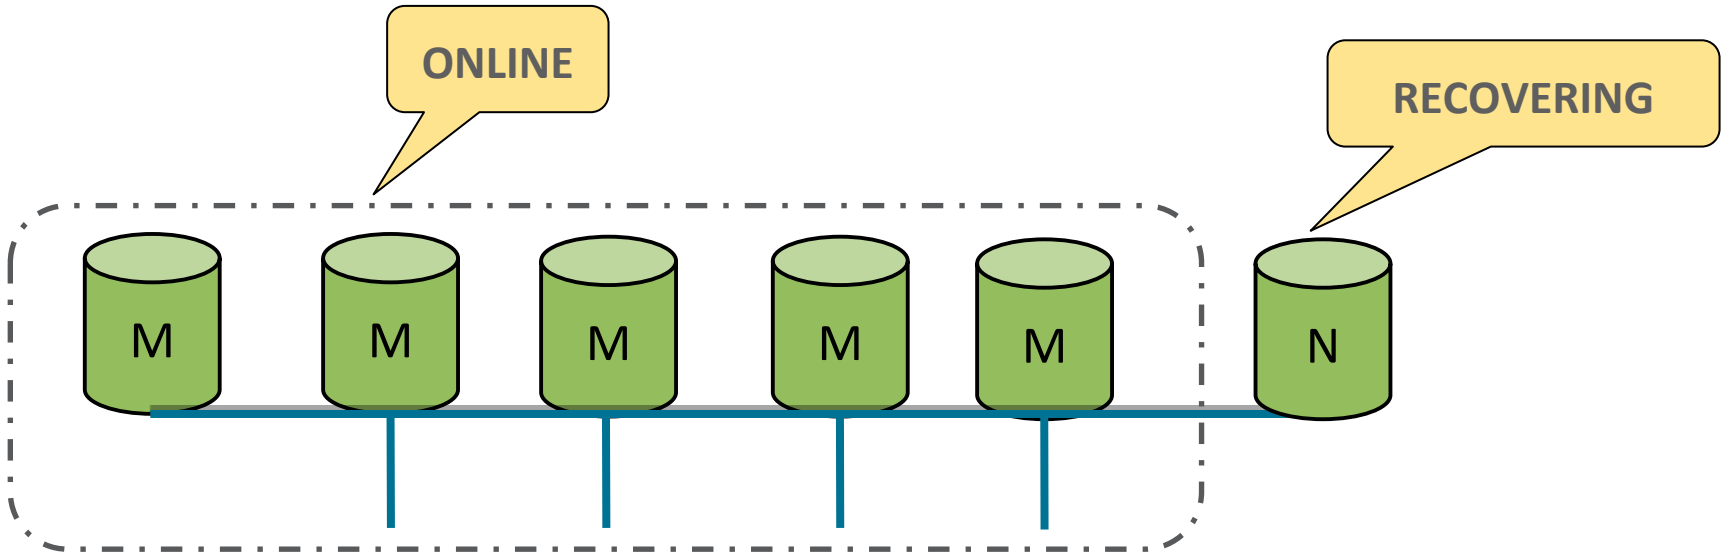

MySQL Group Replication: Multi Primary Automatic Conflict Detection!

- Any two transactions on different servers can write to the same tuple.
- Conflicts will be detected and dealt with.
  - First committer wins rule (based on db-versions in writesets)

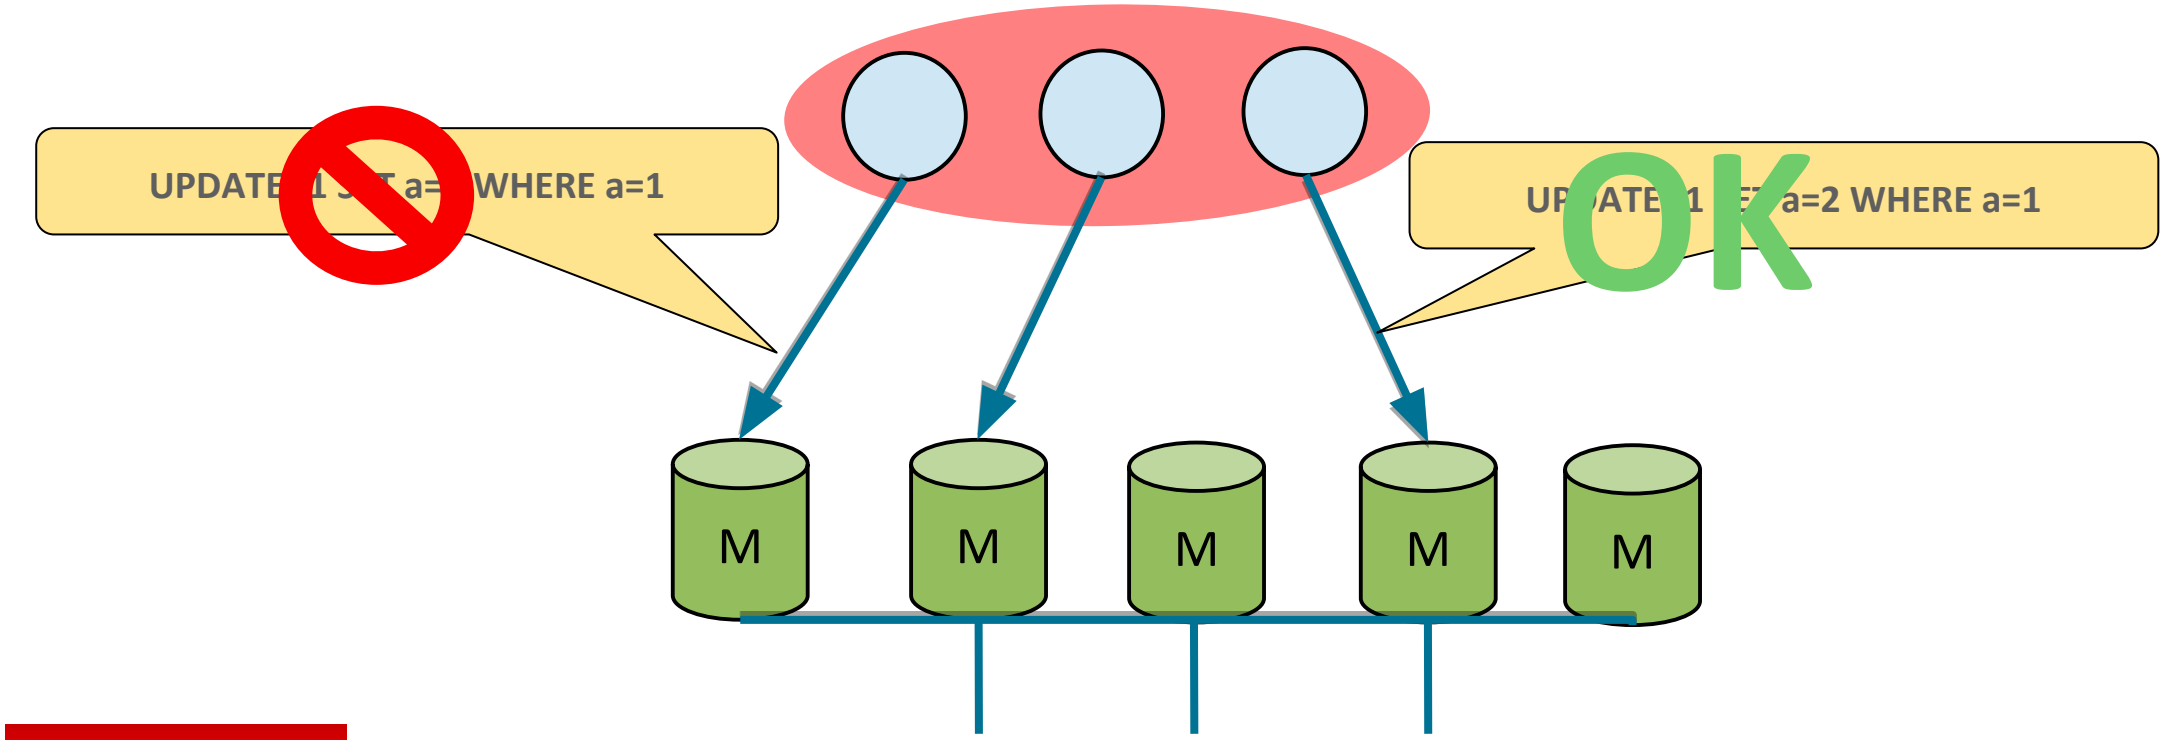

MySQL Group Replication: Major Building Blocks

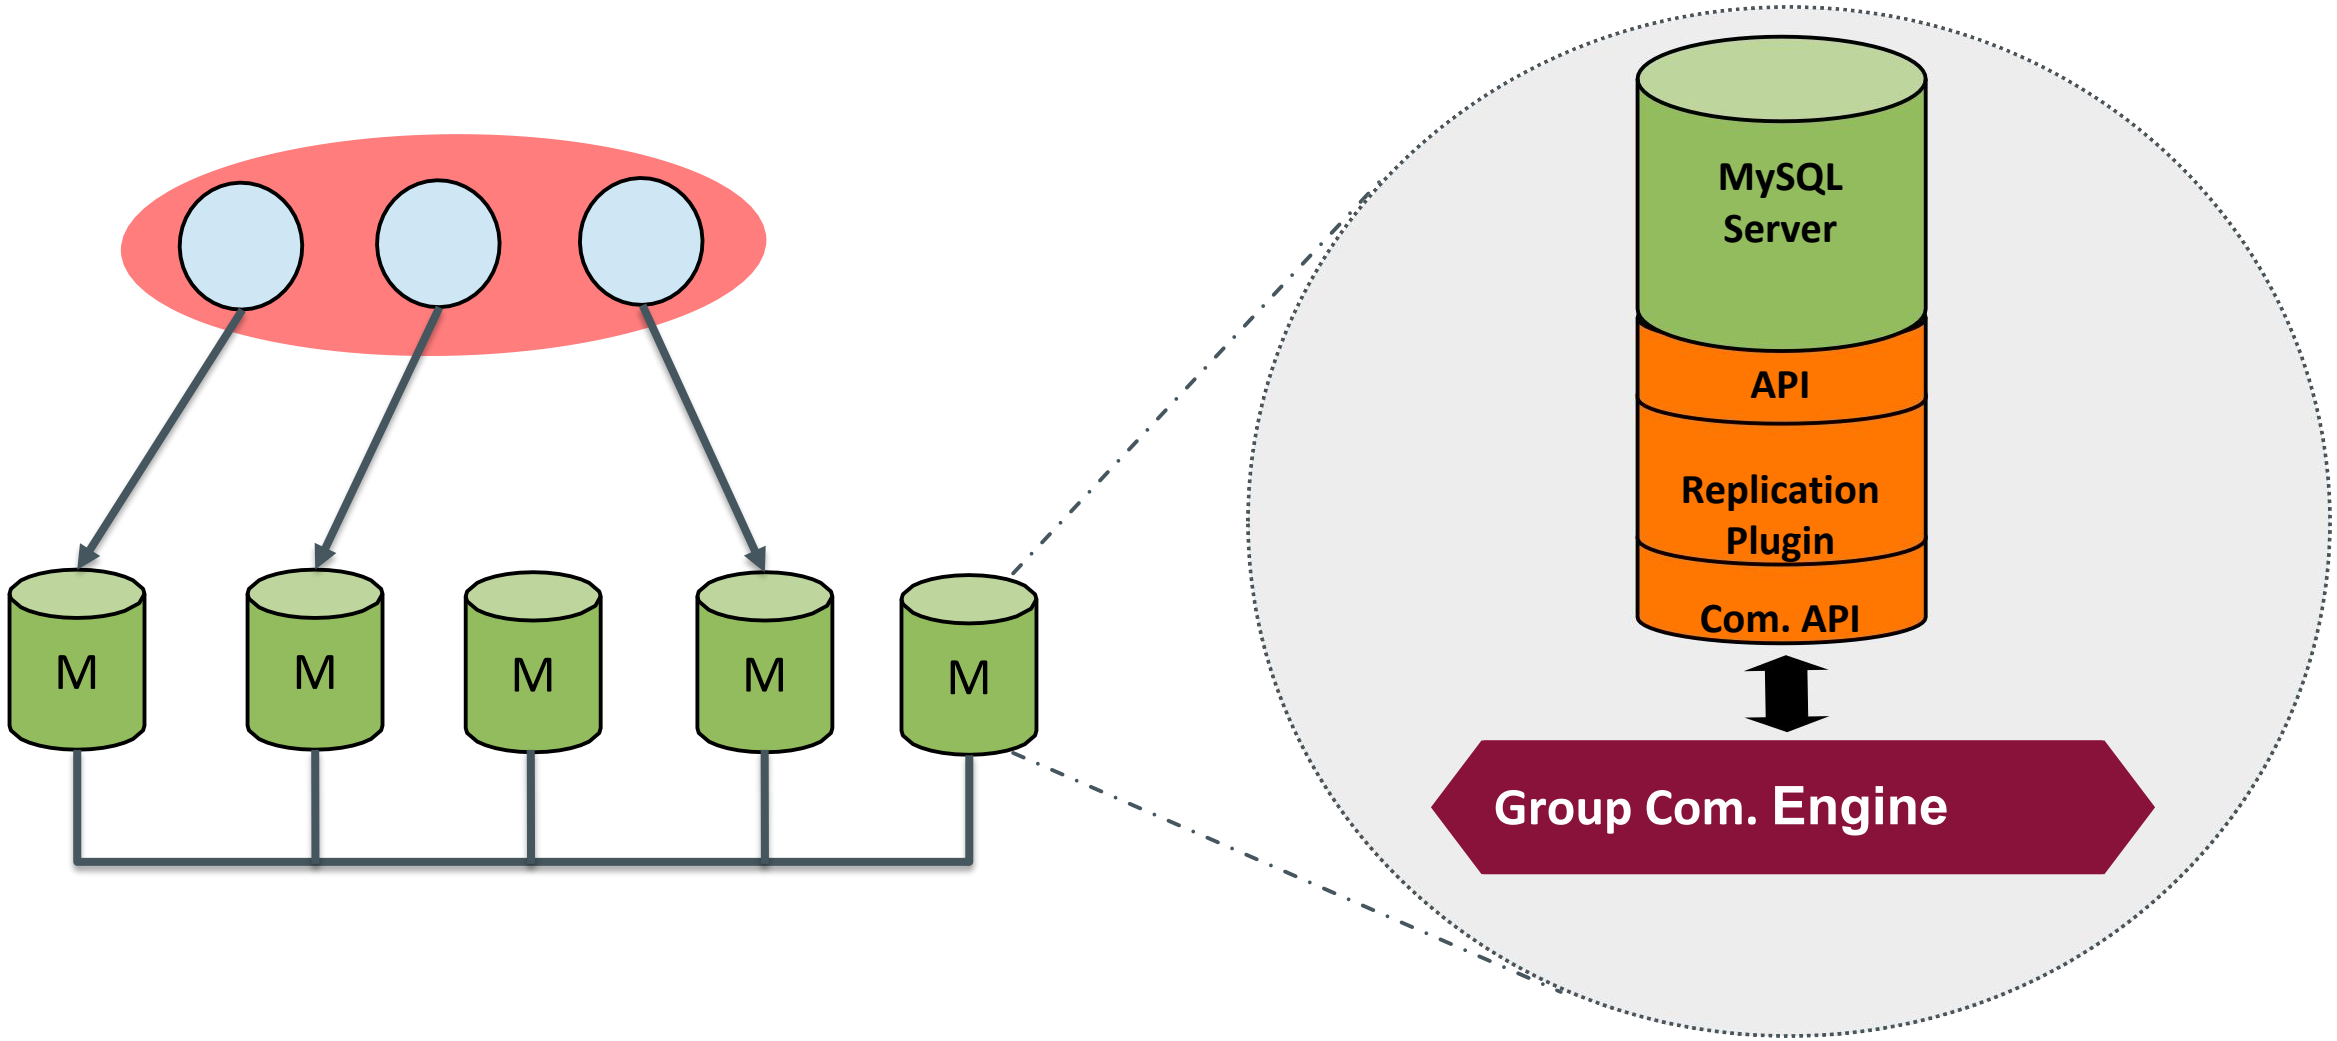

MySQL Group Replication

Full Transaction life cycle

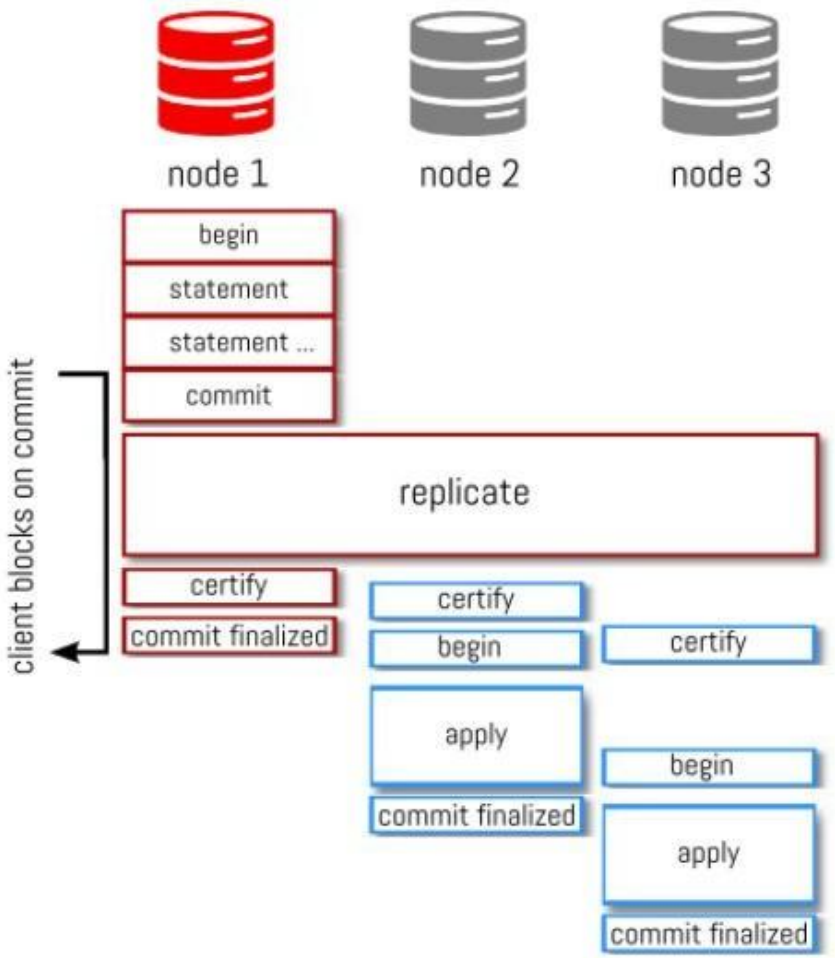

MySQL Group Replication: Performance Comparison

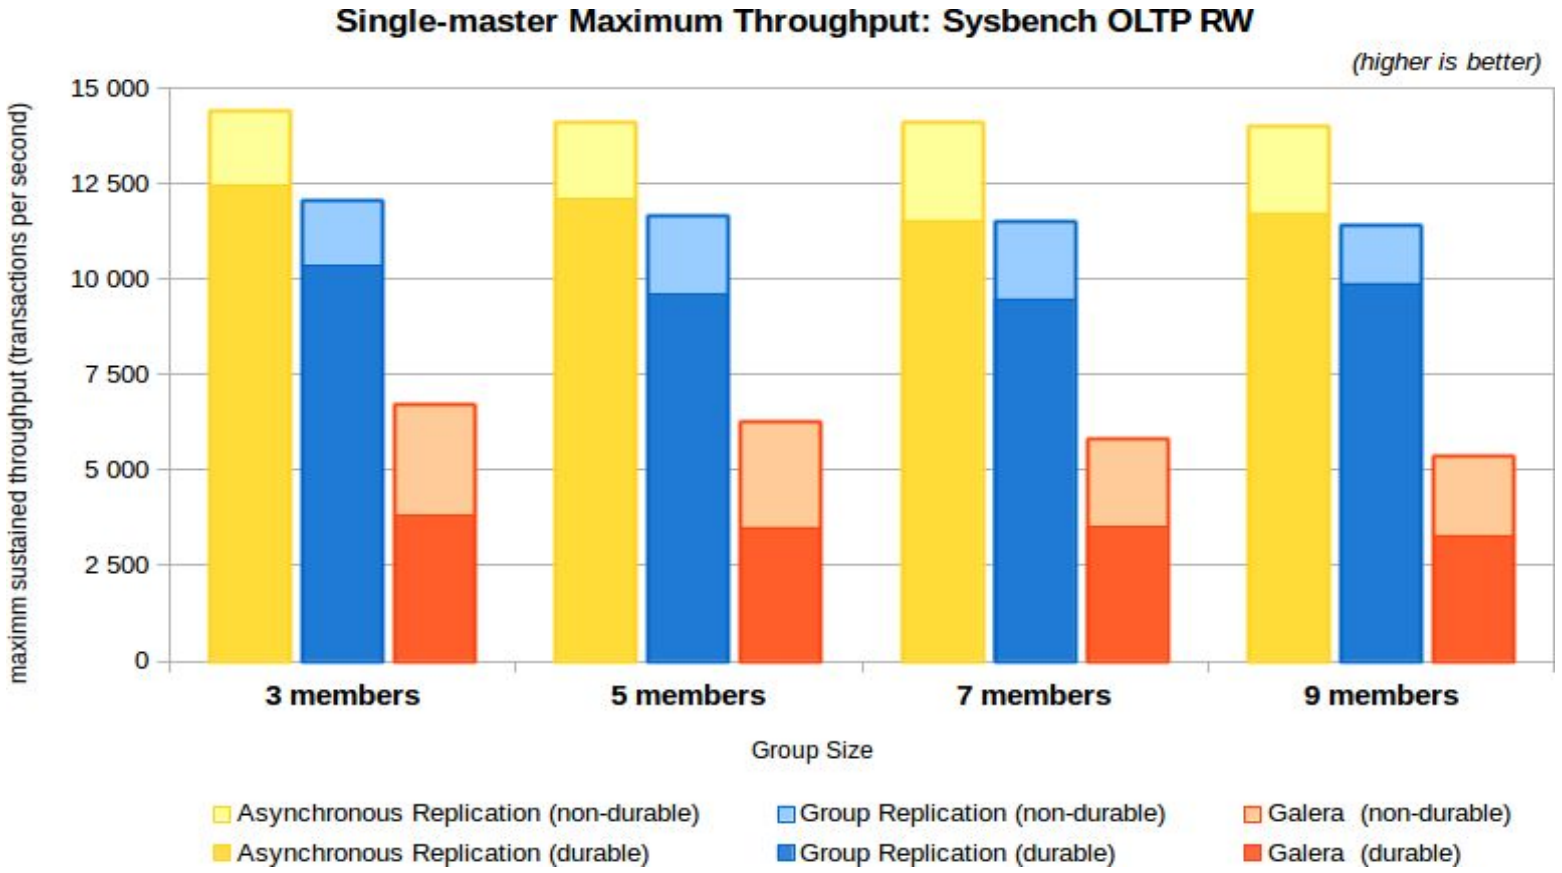

Source: https://mysqlhighavailability.com/performance-evaluation-mysql-5-7-group-replication/

Durability is the **D** in ACID: https://dev.mysql.com/doc/refman/8.0/en/mysql-acid.html

Zabbix HA Deployment with MySQL InnoDB Cluster

ORACLE

InnoDB Cluster Requirements (by design)

- Requires InnoDB storage engine.
- Primary key/unique non-null key is required on every table.
- Requires global transaction identifiers turned on.
- No concurrent DDL
- No Transaction Isolation Mode “SERIALIZABLE”
- Restrictions on usage of Cascaded Foreign Keys

InnoDB Cluster Deployment: Hardware and Infrastructure Notes

- 3, 5, 7 or 9 machines per group
  - Isolate machine resources as much as possible
  - Limit virtualization layers
  - Machines configured for dedicated database server role
    - Recommended configuration
      - 32-64 vCPUs with fast CPU clock (2.5GHz+)
      - SSDs (for data and replication logs)
  - High quality network connection between each machine
    - Low latency, high throughput, reliable
    - Limit routers and hubs as much as possible
    - Isolated and dedicated network when possible

Architecture and Deployment overview

- InnoDB Cluster with 3 nodes
- MySQL Router on each app server
- Install Zabbix (if needed)
- Set/Clone the database instances
  (add primary keys where needed)
- Create a InnoDB Cluster
- Add a MySQL Router
- [if needed] Create Zabbix database
- Point the application to the router
  - Configure the database connection

what more? ...monitor? ...administer?
COME TO THE WORKSHOP!!

Conclusion

MySQL InnoDB Cluster

ZABBIX

InnoDB Cluster makes your Zabbix deployment highly available with a powerul, self-healing, easy to deploy high-availability solution, natively provided by MySQL.

Zabbix monitors your critical systems:

take care of with MySQL InnoDB Cluster!

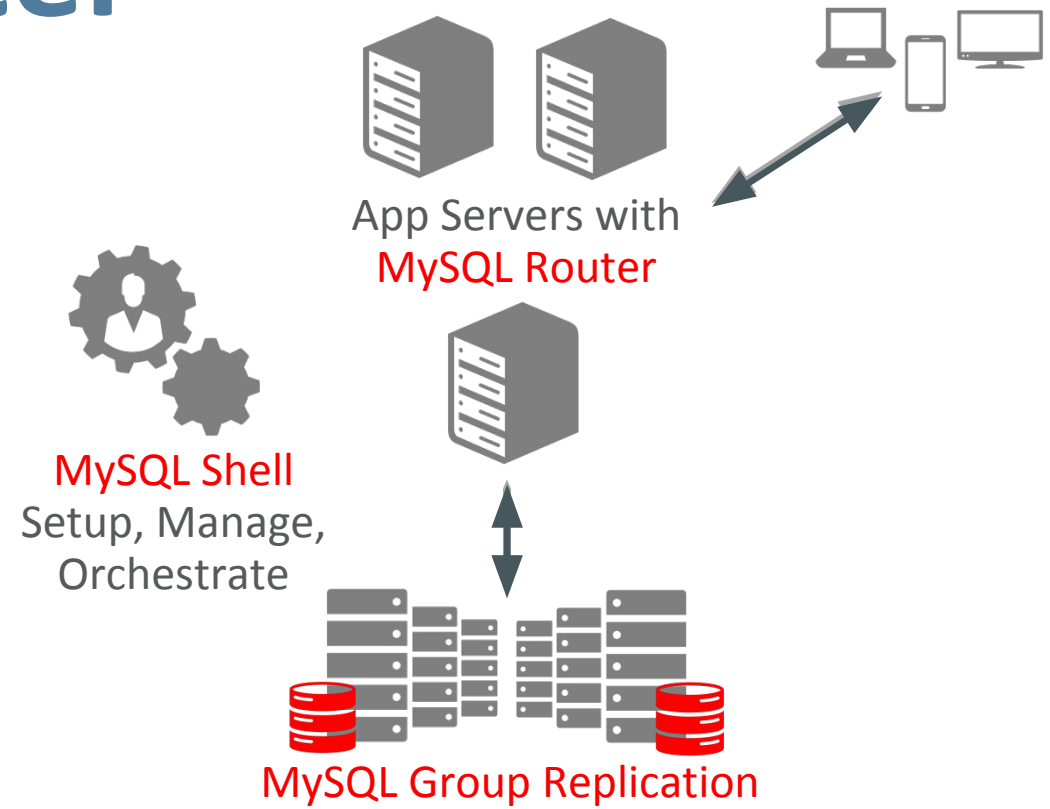

ORACLE