ORACLE

99.99% SLAs with MySQL HA

Name
Luis Soares
Senior Software Development Director
May 1, 2024

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, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Forward-looking Statements

Statements in this presentation relating to Oracle's future plans, expectations, beliefs, intentions, and prospects are "forward-looking statements" and are subject to material risks and uncertainties. A detailed discussion of these factors and other risks that affect our business is contained in Oracle's Securities and Exchange Commission (SEC) filings, including our most recent reports on Form 10-K and Form 10-Q under the heading "Risk Factors." These filings are available on the SEC's website or on Oracle's website at http://www.oracle.com/investor. All information in this presentation is current as of May 2024 and Oracle undertakes no duty to update any statement in light of new information or future events."

Who Am I ?

Luis Soares

- Born and raised in Portugal
- Sports: Football, Basket, Karate, Running, Biking
- Physics, Astronomy
- Fault-Tolerance, High Availability, Computers
- Read, Travel, Engage People
- Long time MySQLer (15+ years)

Agenda

- Background
- MySQL InnoDB Cluster, ReplicaSets and ClusterSets
- MySQL Heatwave Service – High Availability
- MySQL Heatwave Service – Replication
- MySQL Heatwave Service – Scale Out
- Considerations
- Conclusion

Background

Subhead goes here

Faults, Failures and Errors: RPO, RTO and SLAs

Concepts

RTO: Recovery Time Objective

- How long does it take to recover from a single failure

RPO: Recovery Point Objective

- How much data can be lost when a failure occurs

SLA: Service Level Agreement

- RTO and RPO heavily influence SLAs

Types of Failure

High Availability:

- Single Server Failure, Network Partition

Disaster Recovery:

- Full Region/Network Failure

Human Error:

- Little Bobby Tables

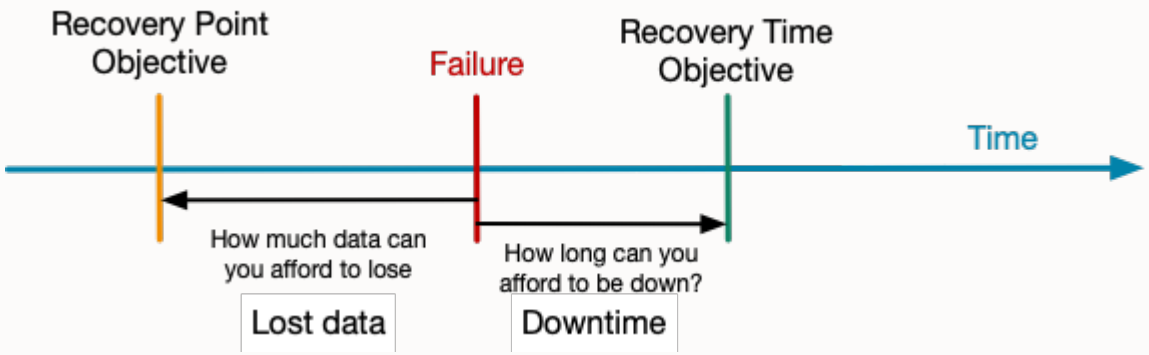

Well-Estabilished Foundation

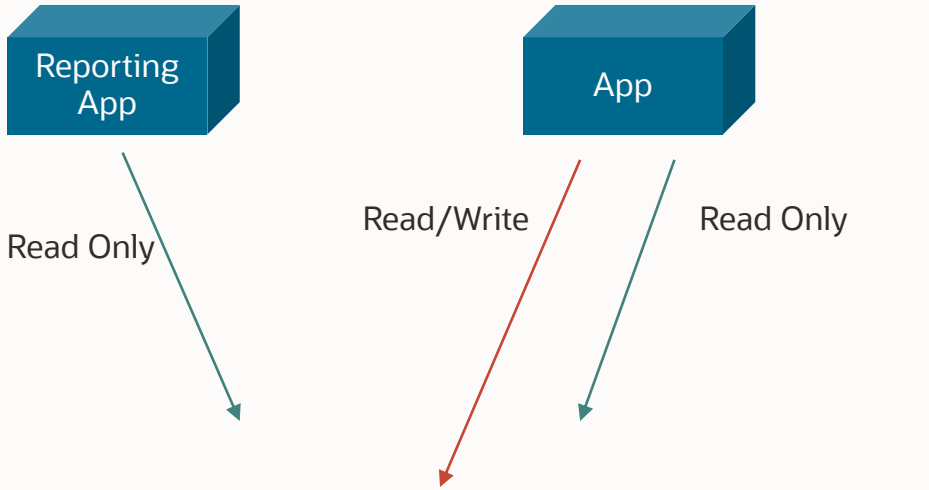

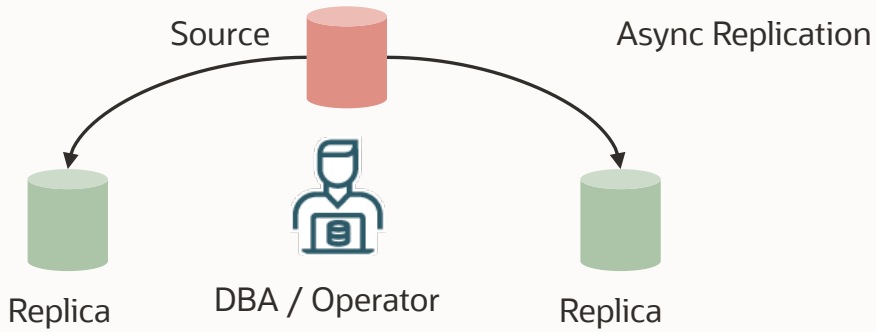

MySQL 3.23.15 (Released in May 2000)

- MySQL Replication is released.

Proven Replication Technology

- Logical events describing changes
- Global Transaction Identifiers make any topology change super-easy and automatic
- Full flexibility

Not only Replication

- Capture Data Change
- Data Integration
- Distributed Recovery

MySQL Everywhere

Observe, Automate, Operate

Observe

- Instrument, emit
- Learn, understand, diagnose
- Trends and historical data
- Robots first, then Humans

Automate

- Balance
- Predict
- Self-heal
- Stabilize

Operate

- Plan
- Troubleshoot
- Press buttons, turn nods, flip switches

MySQL Everywhere, Solutions for Everyone

Observe, Automate, Operate

- MySQL is used everywhere and handles different workload patterns

- The toolset makes it remarkably easier to run it yourself
  - InnoDB Cluster (HA, resilient, fault-tolerant)
  - InnoDB ReplicaSet (Asynchronous)
  - InnoDB ClusterSet (Across clusters, Across regions)

MySQL Server

- Durability
- Storage
- Clone
- Async Repl.

MySQL Router

- Application failover

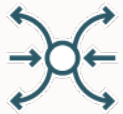

MySQL Shell

- Automation
- Improved UX
- Standard recipes

MySQL Group Replication

- Automation
- Self-healing
- Distributed recovery
- Auto Membership
- Automated server failover

MySQL Everywhere, Solutions for Everyone

Observe, Automate, Operate

- MySQL is used everywhere and handles different workload patterns

- The toolset makes it remarkably easier to run it yourself
  - InnoDB Cluster (HA, resilient, fault-tolerant)
  - InnoDB ReplicaSet (Asynchronous)
  - InnoDB ClusterSet (Across clusters, Across regions)

- **The need to run, monitor, operate and manage:**
  - At scale
  - Heterogeneous workloads
  - Under transiently unreliable networks
  - With multiple versions of the world (split-brain)
  - While undergoing maintenance

MySQL Everywhere, Solutions for Everyone

Observed, Automated, Operated, **Managed**

- MySQL is used everywhere and handles different workload patterns

- MySQL Heatwave Service
  - HA, resilient, fault-tolerant
  - Replicate between DB Systems (even across regions)
  - Backup and Restore (to a point in time)

- **The need to run, monitor, operate and manage:**
  - At scale
  - Heterogeneous workloads
  - Under transiently unreliable networks
  - With multiple versions of the world (split-brain)
  - While undergoing maintenance

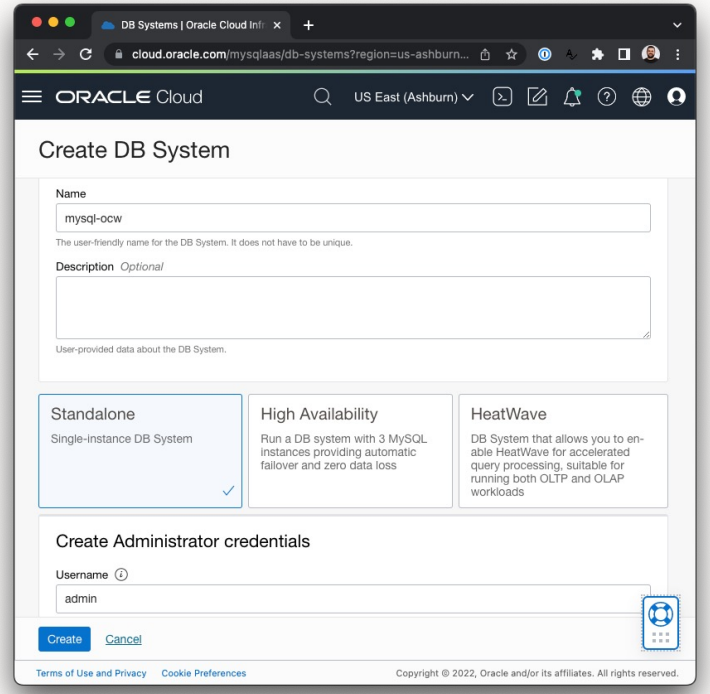

Heatwave Service

MySQL Everywhere, Solutions for Everyone

Observed, Automated, Operated, **Managed**

- MySQL is used everywhere and handles different workload patterns

- MySQL Heatwave Service
  - HA, resilient, fault-tolerant
  - Replicate between DB Systems (even across regions)
  - Backup and Restore (to a point in time)

MySQL Server

- Enterprise Edition
- Durability
- Storage

User Experience

- Console
- REST APIs
- Terraform
- SDKs

Oracle Cloud Infrastructure

- Block Volume
- VCN
- Compute
- Endpoints
- Workflows
- Redundancy

MySQL Group Replication

- Automation
- Self-healing
- Distributed recovery
- Auto Membership
- Automated server failover

MySQL InnoDB Cluster, ReplicaSets and ClusterSets

Deploy and Run your MySQL High Availability Setup

MySQL Solutions

One Product: MySQL

- Technologies coming together
  - MySQL Server
    - Async Replication
    - Group Replication
  - MySQL Shell
  - MySQL Router

- Full stack testing

Easy to use!

- One client: MySQL Shell
- Transparent Access to Database Architecture: MySQL Router

User deploys and runs MySQL infrastructure.

- The technology helps.

Solution: MySQL InnoDB ReplicaSet

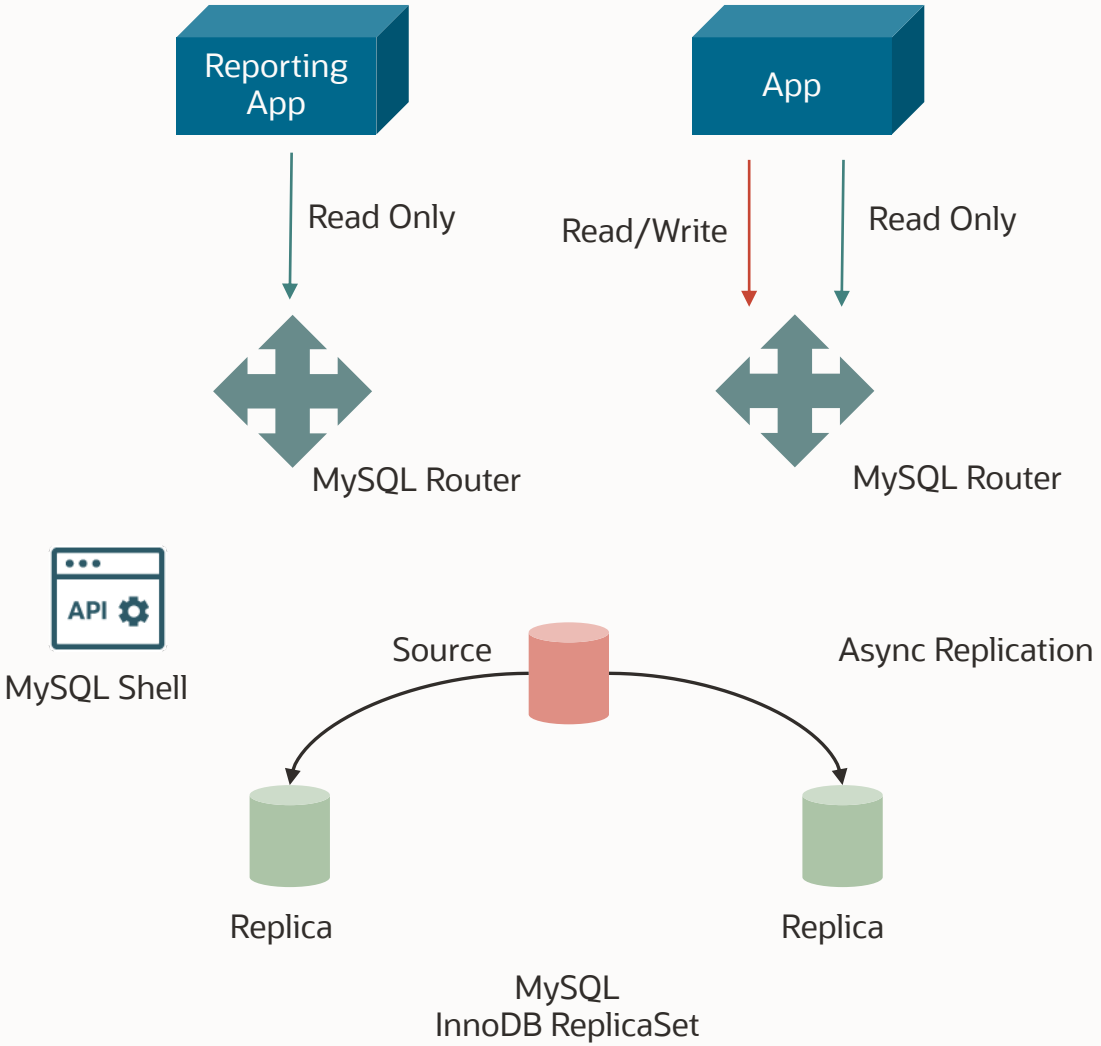

'classic', 'asynchronous' Replication based Solution, fully integrated

- MySQL Shell
- MySQL Router
- MySQL Server
- MySQL Async Replication

- RPO != 0
- RTO = minutes or more (manual failover)

Solution: MySQL InnoDB Cluster

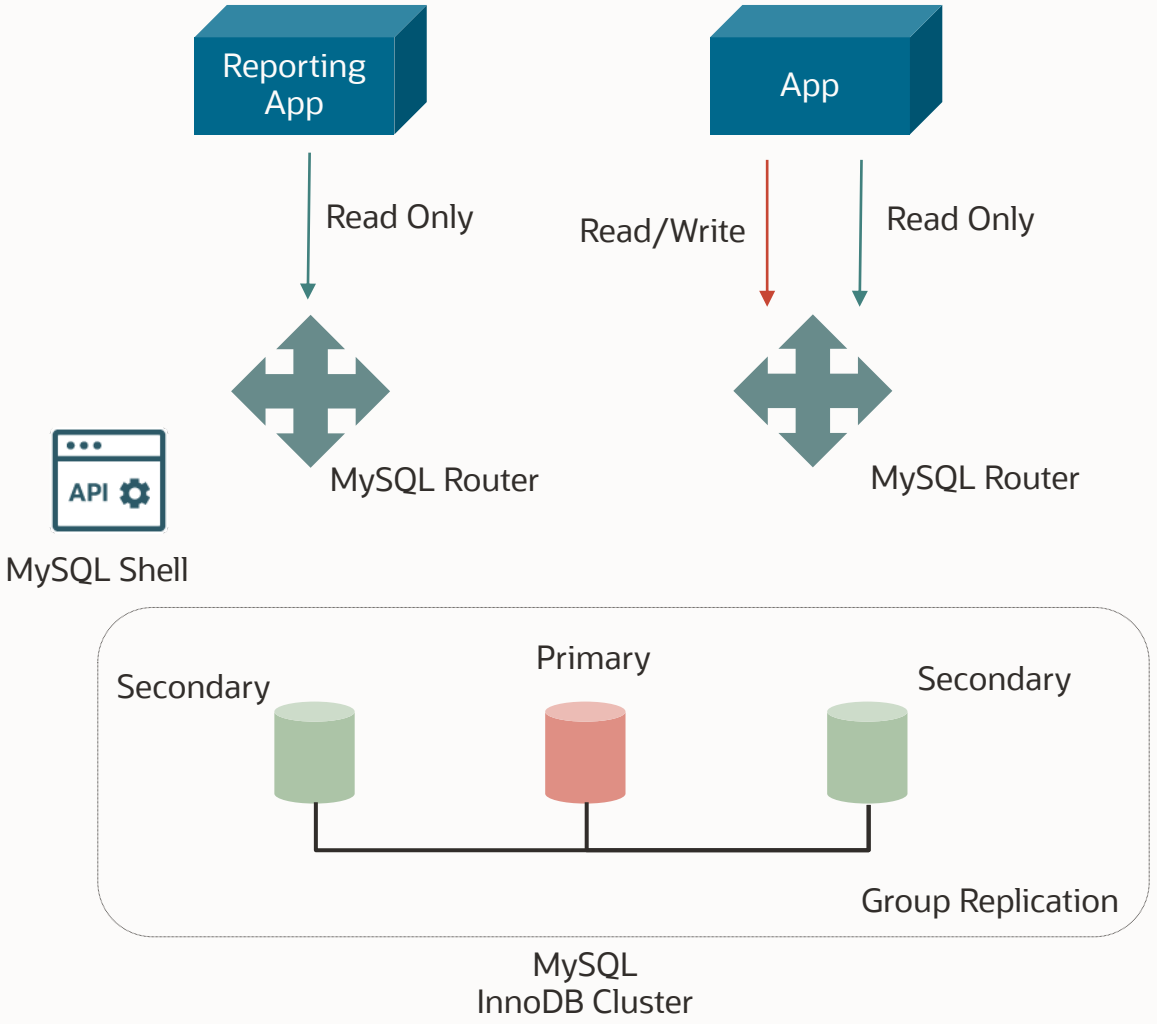

High Availability solution based on Group Replication, fully integrated

- MySQL Shell
- MySQL Router
  - Application failover
- MySQL Server
  - Group Replication providing:
    - Automatic Server Failover
    - Automatic membership changes

- RPO = 0
- RTO = seconds (automatic failover)

Solution: MySQL InnoDB ClusterSet

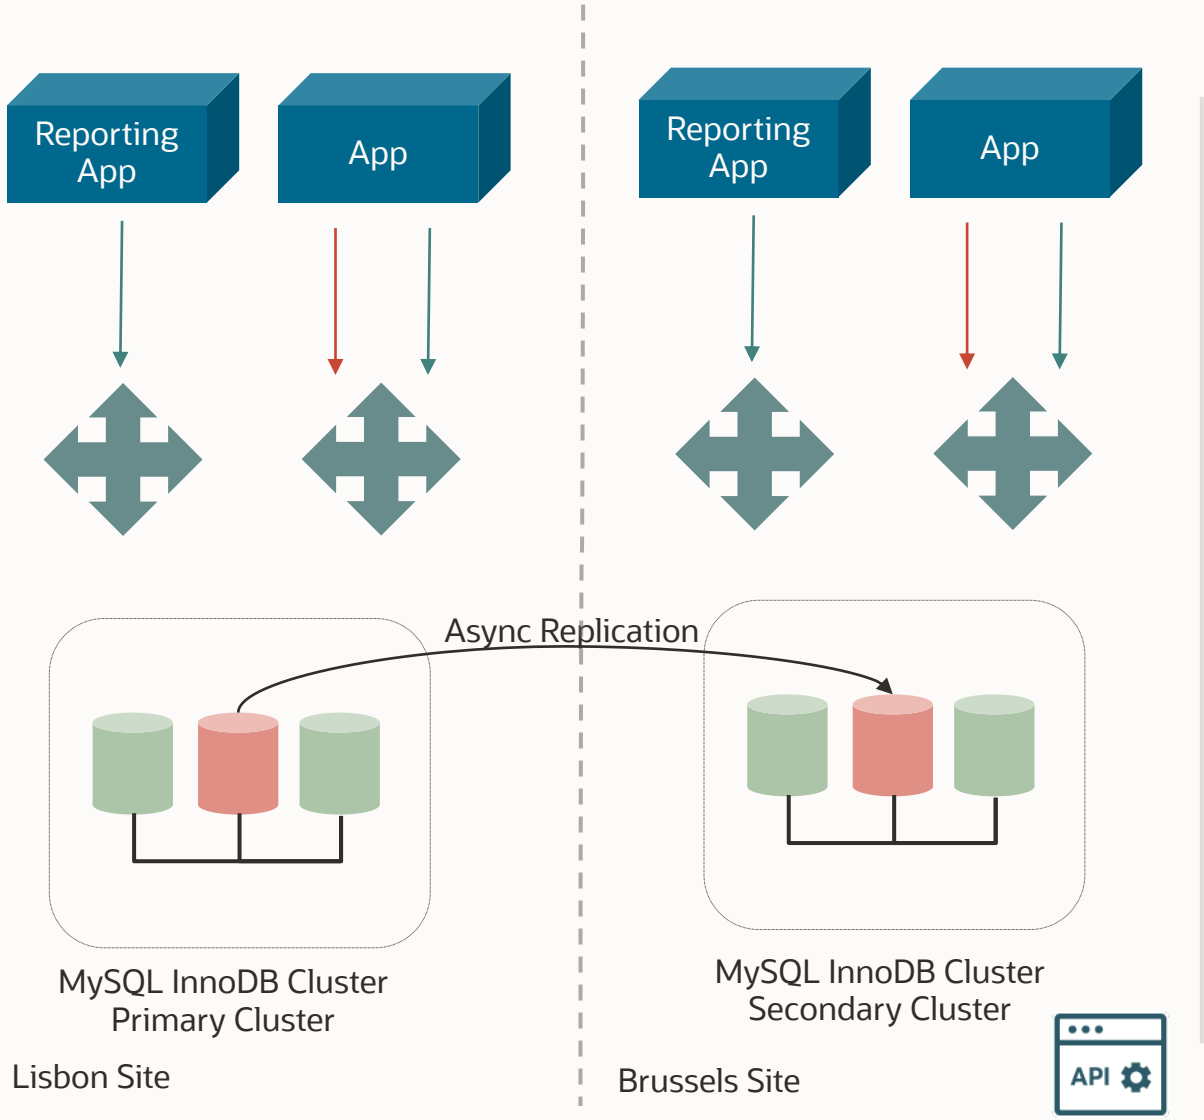

Disaster Tolerance Solution for InnoDB Clusters

- High Availability (Failure within a site)
  - RPO = 0
  - RTO = seconds (automatic failover)
- Disaster Recovery (Site Failure)
  - RPO != 0
  - RTO = minutes or more (manual failover)
- Easy to use: MySQL Shell
- Full integration of MySQL Router

MySQL Heatwave Service: High Availability

High Availability in the MySQL Heatwave Service on Oracle Cloud Infrastructure

MySQL Heatwave Service

Observed, Automated, Operated, Managed

MySQL Replication powers key features:

- High Availability
- Inbound Replication
- Outbound Replication
- Managed Read Replicas
- Point-in-time Recovery

The need to run, monitor and operate:

- At **scale**
- Exposed to **heterogeneous** workloads
- Coping with network **bursts** or packet **delays**
- Dealing with the “world” splitting
- Through **maintenance**

Simple, intuitive, one-click operations:

- Create DB Systems
- Create Read Replica
- Create Inbound Channel

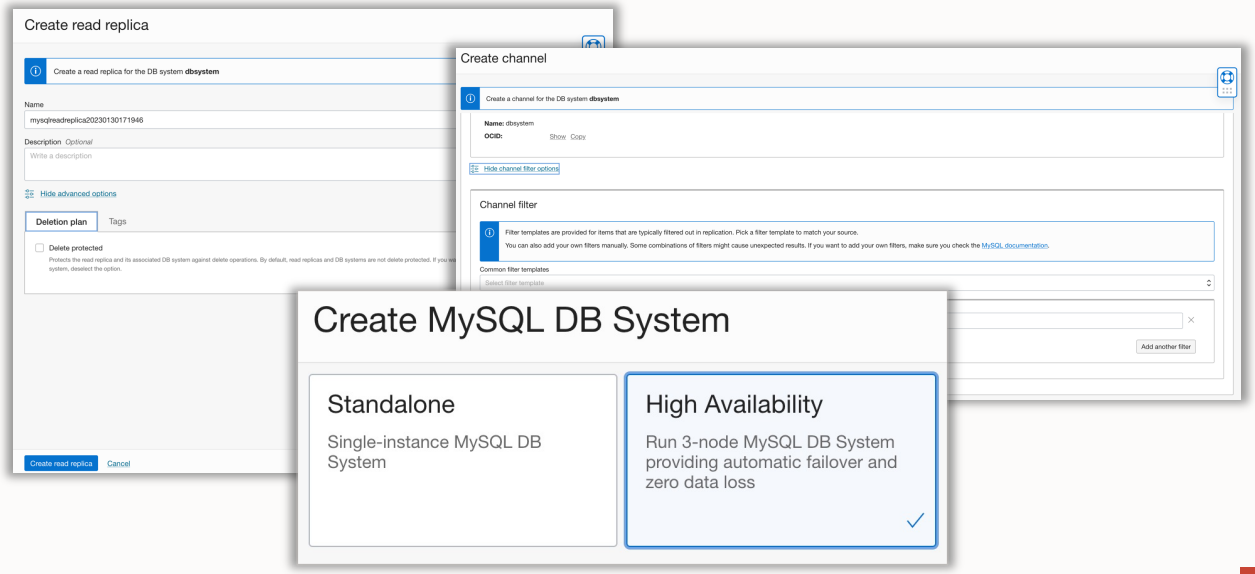

Features

- Latest releases of MySQL Enterprise Servers.
- Easy provisioning & SLA
- Automation via Terraform, CLI, SDK, API
- One click:
  - High Availability
  - Manual, Automatic Backups
  - Point-in-time Recovery
  - Inbound and Outbound Replication
  - OLAP and OLTP together in one database

Database
DB System

One-click **DB System** creation

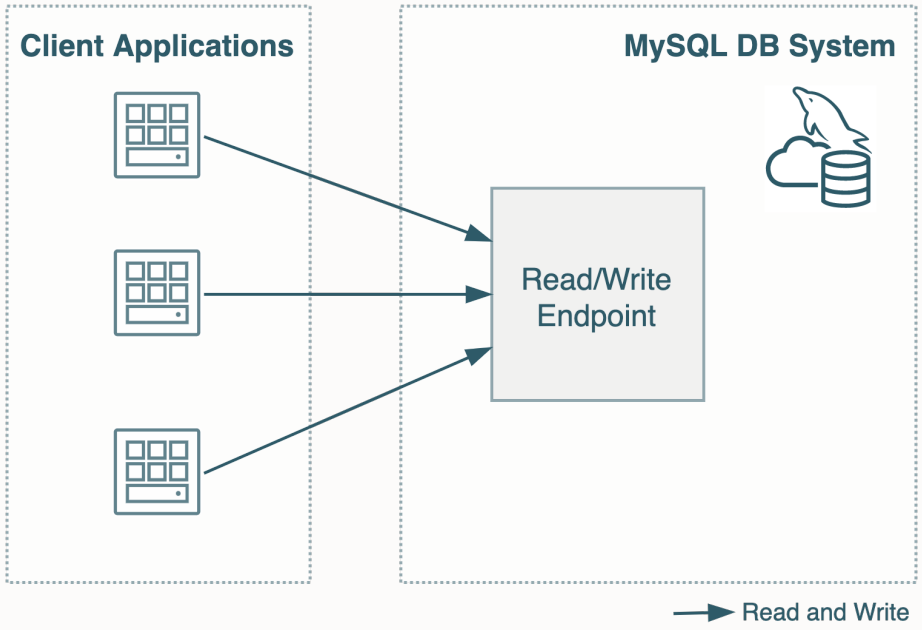

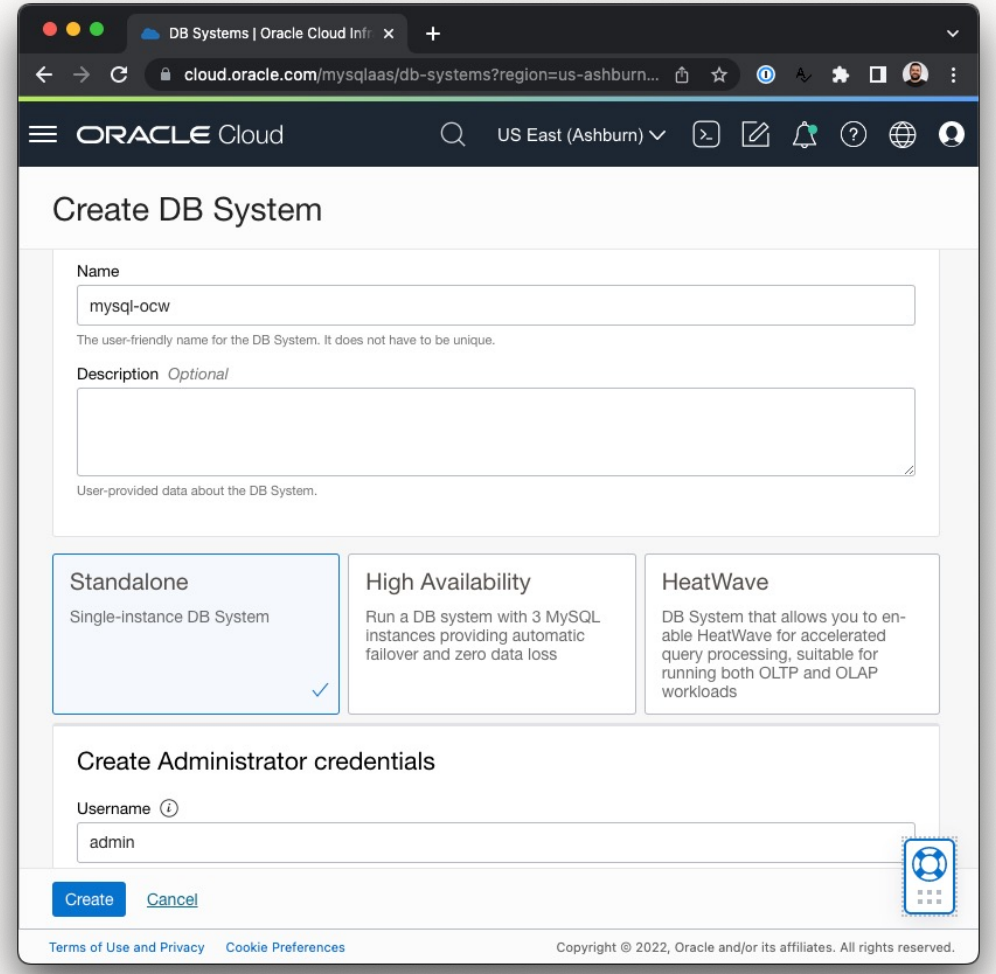

High Availability
Automatic Failover

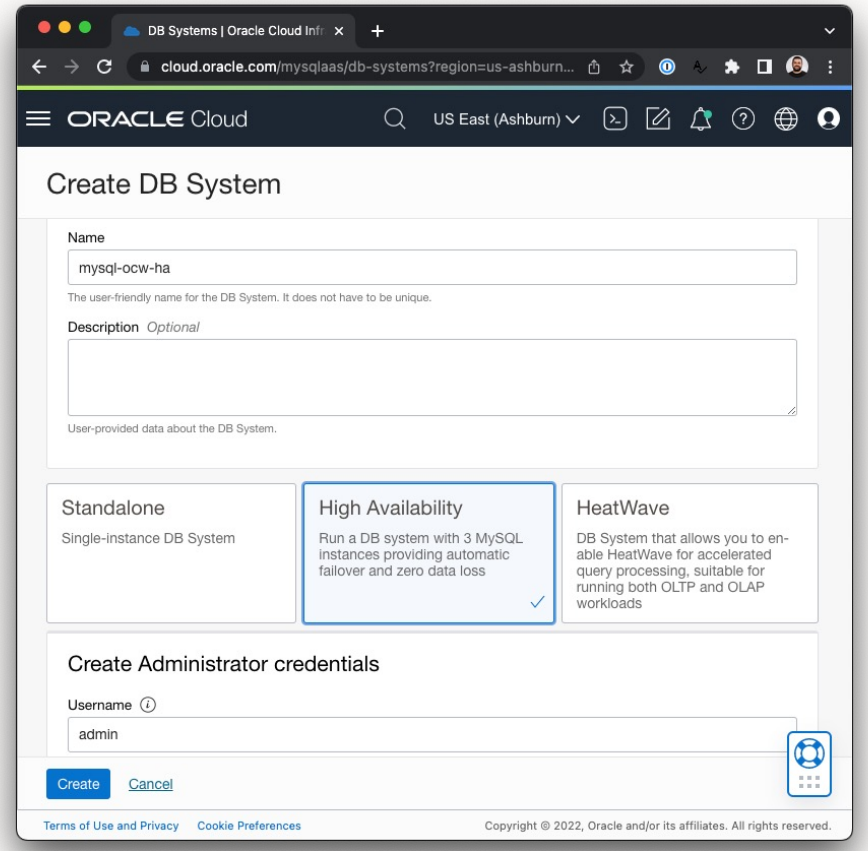

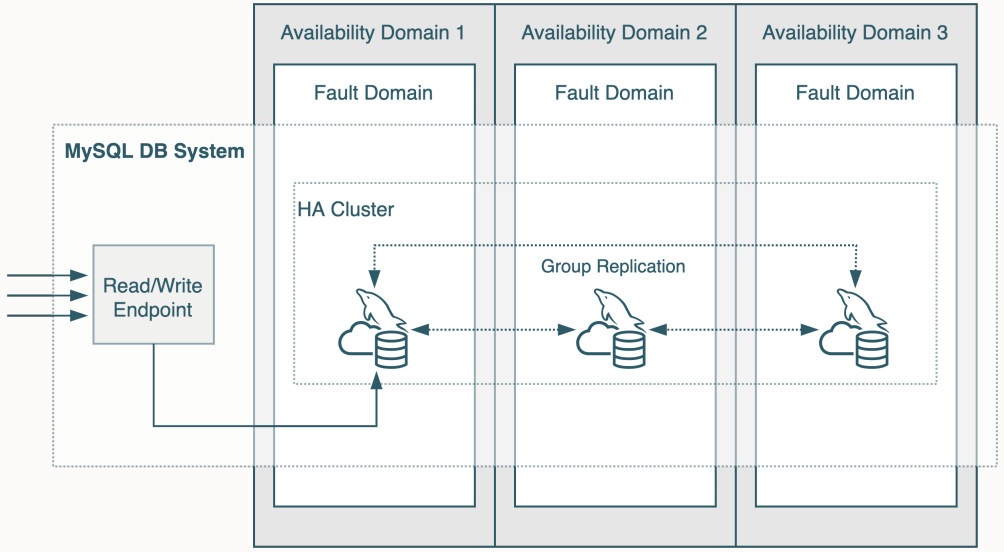

High Availability
Automatic Failover

- SLA 99.99%
- Automatic failover
- Manual switchover
- Rolling upgrades, shape and config changes
  - Just a **few seconds** impact

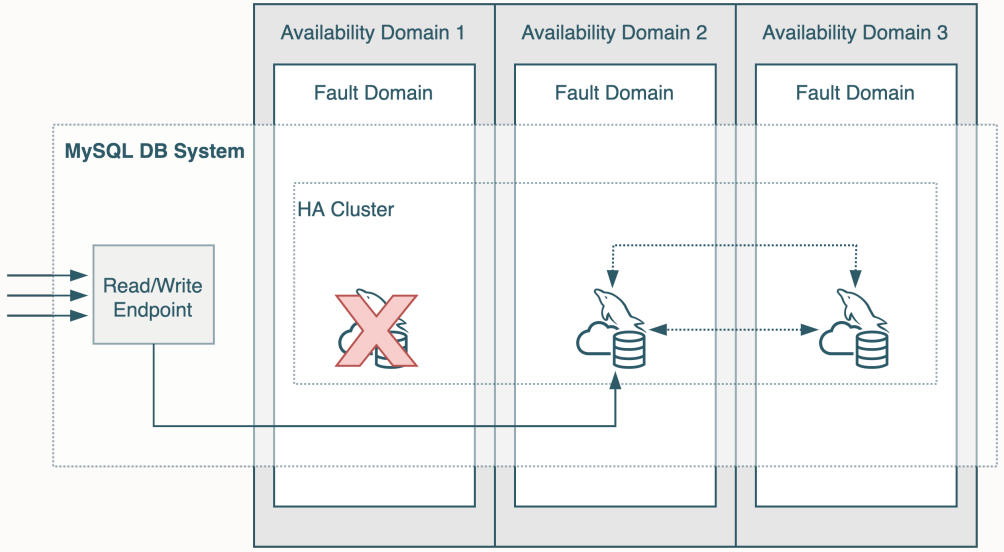

MySQL Heatwave Service: Replication

Multi-cloud and On-premise Integration

Inbound and Outbound Replication

Hybrid Deployments and Migrations

Hybrid deployments

- On-premise and multi-cloud
- OCI as your main site
- OCI as your Disaster Recovery site
- OCI for capacity bursting
- HeatWave for Analytics

Live Migrations

- Minimize downtime

Cross-region replication

- DB System to DB System

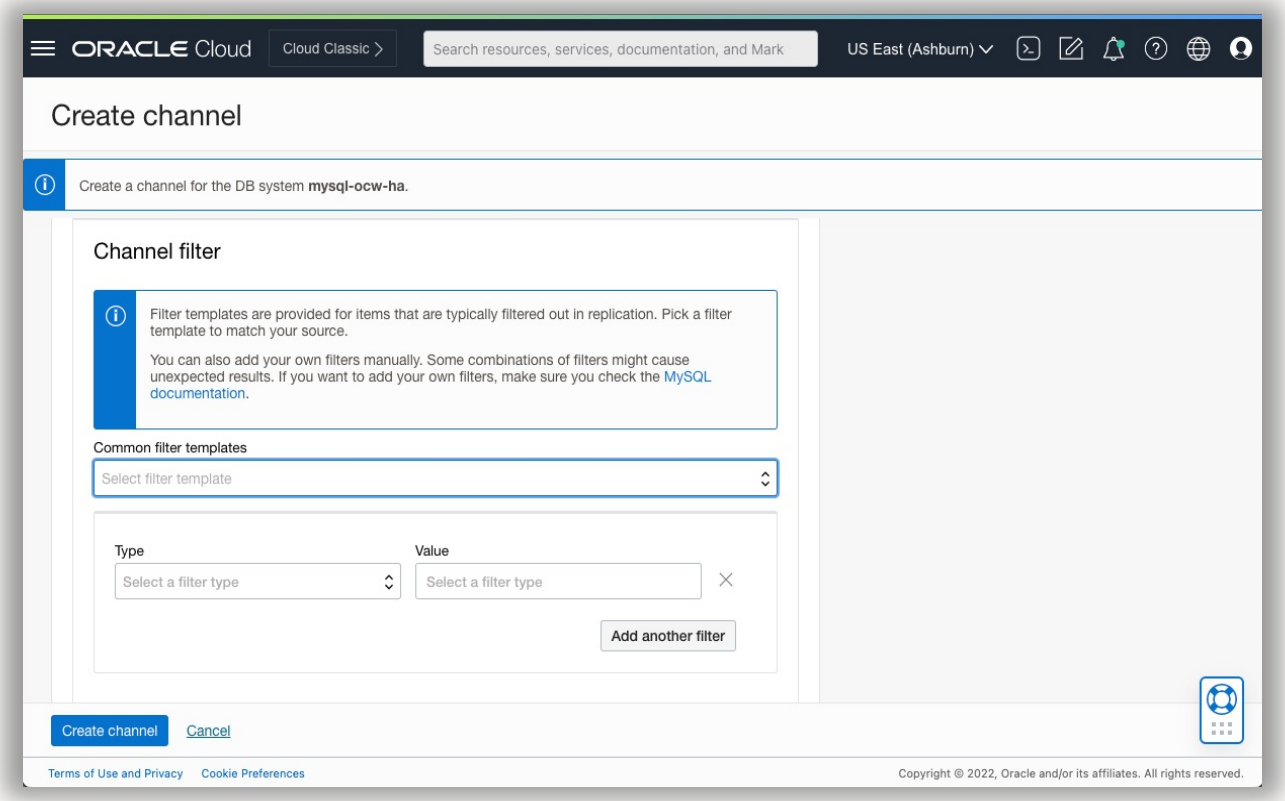

Inbound and Outbound Replication

Disaster Recovery in Hybrid Environments

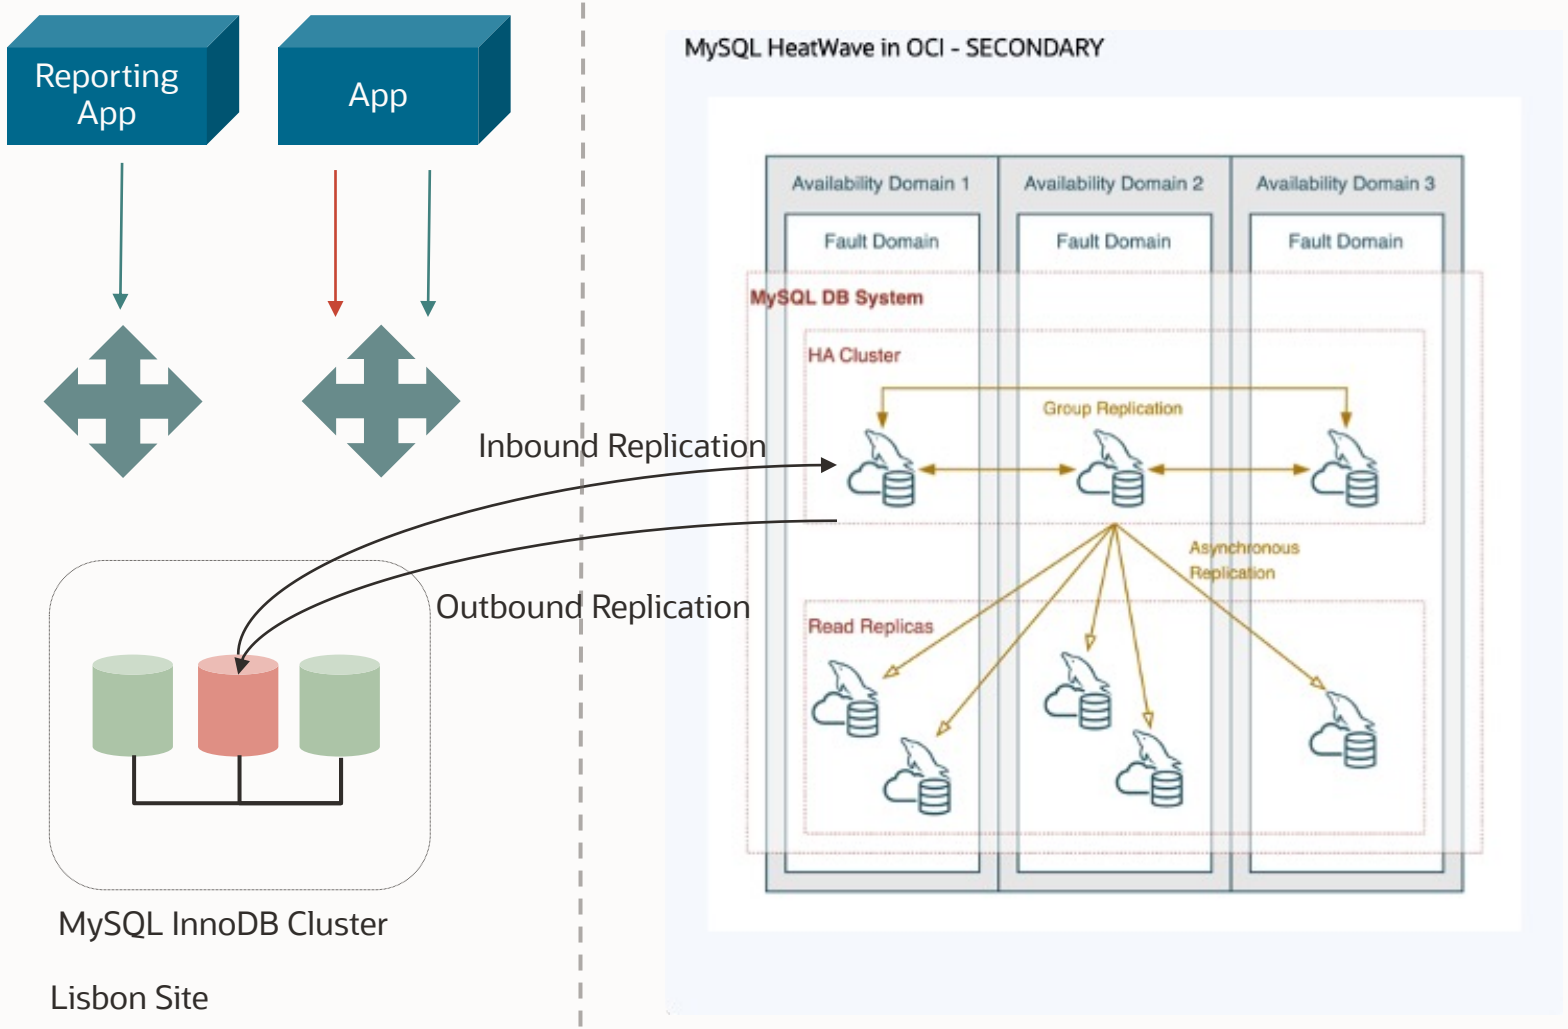

- Use OCI as DR or vice-versa
- Leverage MySQL HeatWave Service in OCI
- MySQL HeatWave Service Channels to do inbound replication, based on built-in asynchronous replication

MySQL Heatwave Service: Scale Out

Managed Read Replicas

Grow and Shrink Read Capacity Seamlessly

High performance by scaling your reads. Offload RW endpoint minimize gray failures.

- A single click creates a Read Replica
  - Provision
  - Launch
  - Setup Replication
  - Monitor and Manage

- Read Replicas are associated with a DB System
  - RO endpoints in the DB System
  - Up to 18 max per DB System
  - Requires a shape of 4 OCPUs or larger
  - CLI, SDK and Terraform support

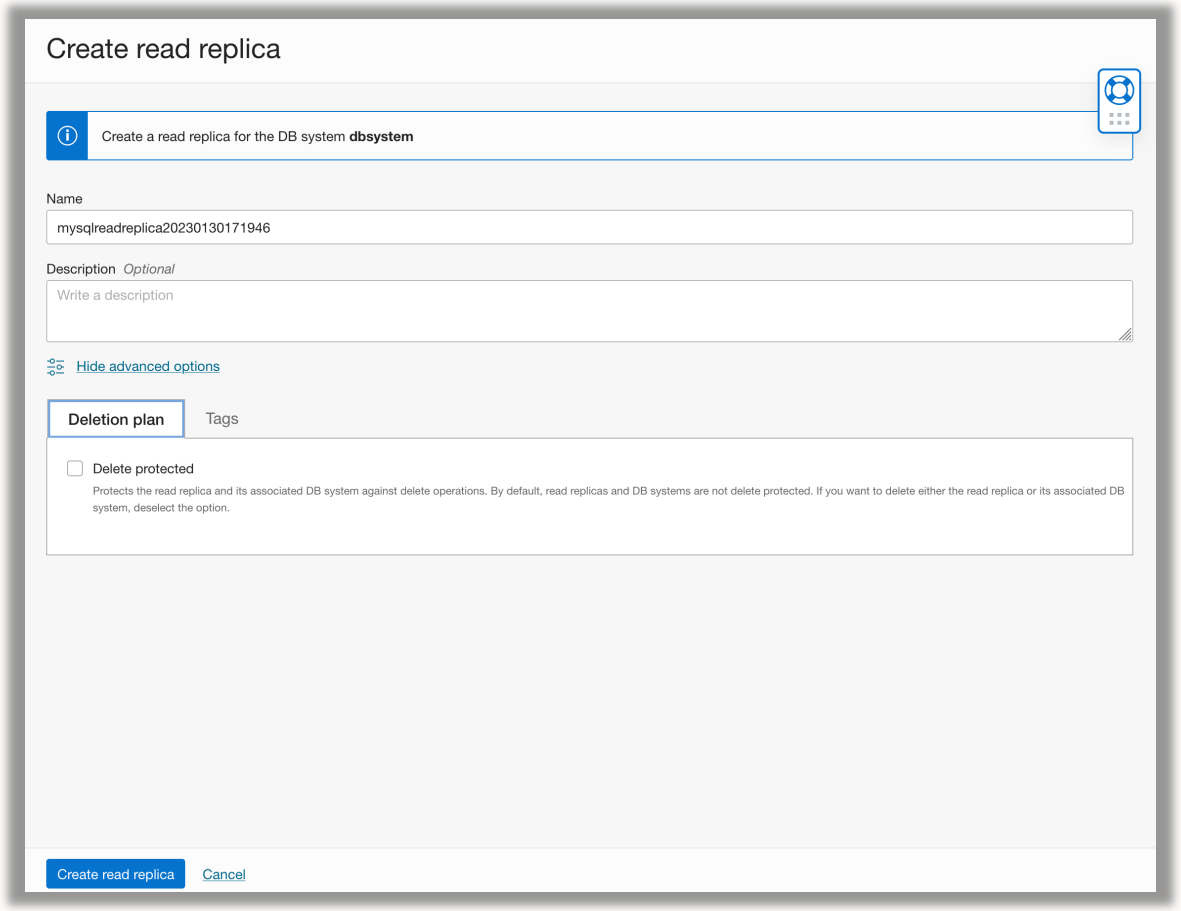

Load Balancer

Use Your Replicas Efficiently

When using Read Replicas a Load Balancer Endpoint is automatically provisioned in your DB System.

- Managed by the service
- Materializes as a Read-Only endpoint
- Round robins traffic across Read Replicas
- Manages Read Replica backends automatically

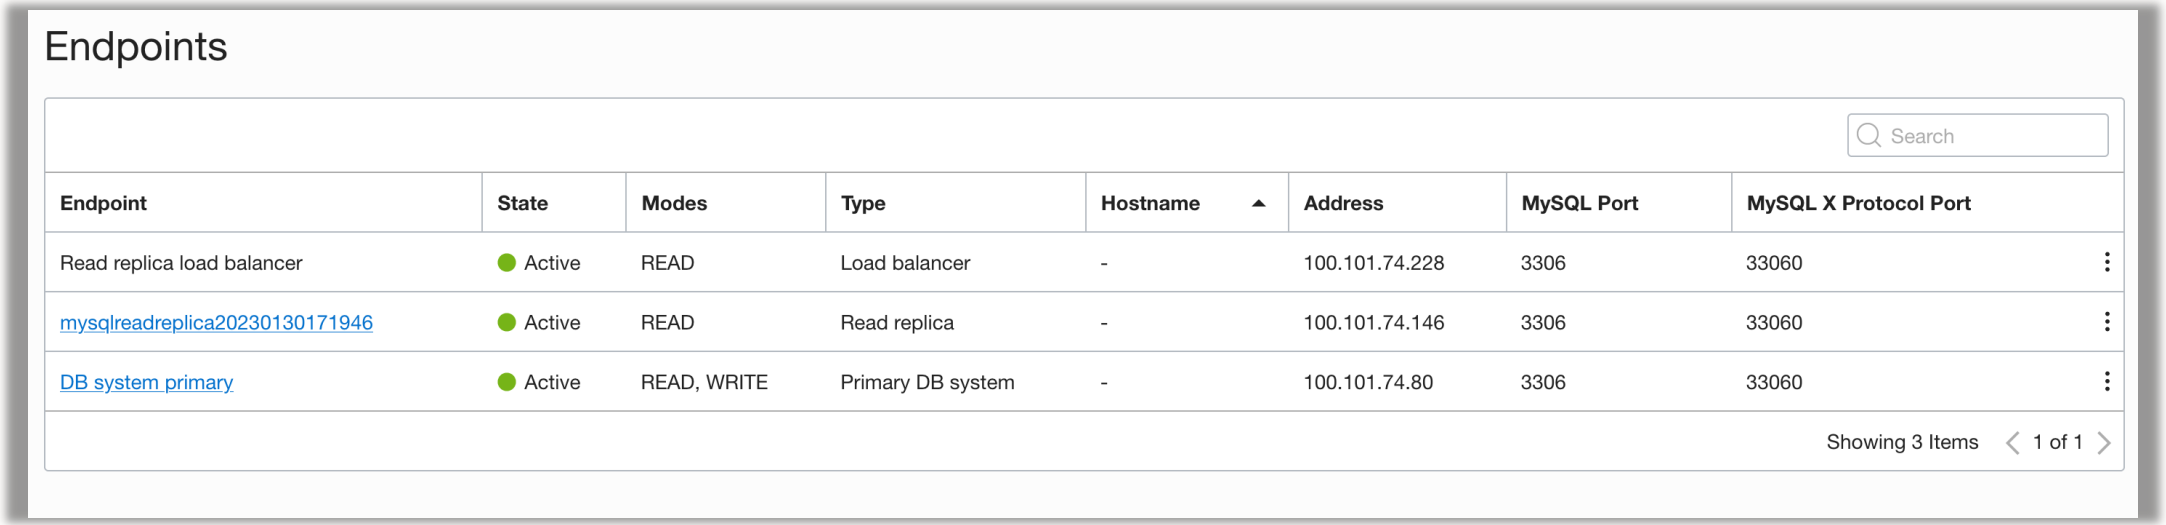

Considerations

Subhead goes here

Database User Advices

Some tips…

- Know the **workload**.
  - Peaks times, sustained load, anticipate problems.
- Do **capacity** planning.
  - Avoid triggering your own gray failures due to constrained resources.
- Make applications **aware of failures**.
  - Utilize application retry logic for network problems and high availability
- Understand the technology **limits**.
  - Read and understand the limits of the technology.
  - Make sure your application does not run into the documented limits.
- **Embrace** the technology.
  - Observe best practices and it will just work.
  - In the rare cases it does not, seek help, work with the experts to either mitigate or fix the problem.

Conclusion

Subhead goes here

Conclusion

- Solutions for everyone, everywhere:
  - On-prem
  - Cloud
  - Hybrid Deployments

- On-Premises, user managed, High Availability Solutions:
  - InnoDB Replica Sets
  - InnoDB Cluster
  - InnoDB Cluster Sets

- Cloud, Managed, Solution with well defined SLAs:
  - MySQL Heatwave Service


ORACLE

Our mission is to help people see data in new ways, discover insights, unlock endless possibilities.