

Cross region replication in Percona K8S Operators

Ege Günes

Software Engineer, Cloud

Data on Kubernetes Talks #107

2021-12-09



- Ege Günes
- Istanbul, Turkey
- GitHub: [@egegunes](https://github.com/egegunes)
- <https://ege.dev>



About Percona



PERCONA DISTRIBUTION
FOR **MYSQL OPERATOR**



PERCONA DISTRIBUTION
FOR **MONGODB OPERATOR**



PERCONA DISTRIBUTION
FOR **POSTGRESQL OPERATOR**



PERCONA
Distribution for MySQL



PERCONA
Distribution for MongoDB

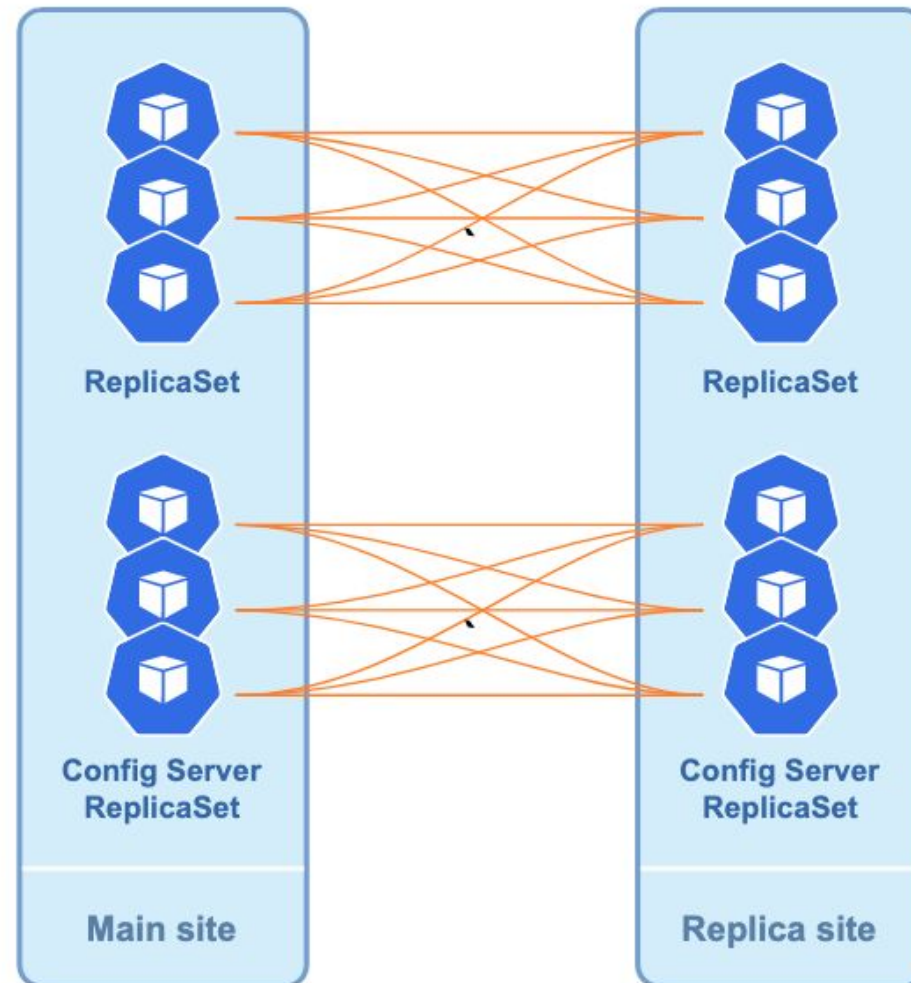
PERCONA
Distribution for PostgreSQL

Cross region replication in MongoDB Operator

Quick Primer on MongoDB Replication

- **Replica Set (Don't confuse with k8s ReplicaSet)**
- **Asynchronous replication**
- **Automatic failover**

Cross region replication



What happens when you deploy a new cluster?

The operator:

1. **Generates SSL certificates**
2. **Generates random passwords for system users**
3. **Creates a StatefulSet to run mongod pods**
4. **Initializes the Replica Set**
5. **Creates the users in Replica Set**

Deploying a Replica Cluster

```
apiVersion: psmdb.percona.com/v1-11-0
kind: PerconaServerMongoDB
metadata:
  name: replica-cluster
spec:
  unmanaged: true
  replsets:
  - name: rs0
    expose:
      enabled: true
      exposeType: LoadBalancer
  ...
```

- **You need to use same secrets in two clusters.**
- **Operator won't configure the replica set.**
- **Cluster needs to be exposed.**

Adding Replica Nodes to Primary Cluster

```
apiVersion: psmdb.percona.com/v1-11-0
kind: PerconaServerMongoDB
metadata:
  name: primary-cluster
spec:
  unmanaged: false
  replsets:
  - name: rs0
    expose:
      enabled: true
      exposeType: LoadBalancer
  externalNodes:
  - host: replica-rs0-0.ege.dev
  - host: replica-rs0-1.ege.dev
  - host: replica-rs0-2.ege.dev
  ...
```


- **Primary cluster needs to be exposed too.**
- **The operator will add external nodes to Replica Set configuration.**
- **External nodes have no votes by default.**

What if you want sharding?

```
apiVersion: psmdb.percona.com/v1-11-0
kind: PerconaServerMongoDB
metadata:
  name: primary-cluster
spec:
  unmanaged: false
  replsets:
  - name: rs0
    externalNodes:
    - host: replica-rs0-0.ege.dev
    - host: replica-rs0-1.ege.dev
    - host: replica-rs0-2.ege.dev
    ...
  sharding:
    configsvrReplSet:
      size: 3
      expose:
        enabled: true
        exposeType: LoadBalancer
      externalNodes:
      - host: replica-cfg-0.ege.dev
      - host: replica-cfg-1.ege.dev
      - host: replica-cfg-2.ege.dev
      ..
```

- **Config server replica set needs to be exposed on both clusters.**
- **You can connect to mongos on either side and they'll route the operations to correct shard.**

What if you want sharding?



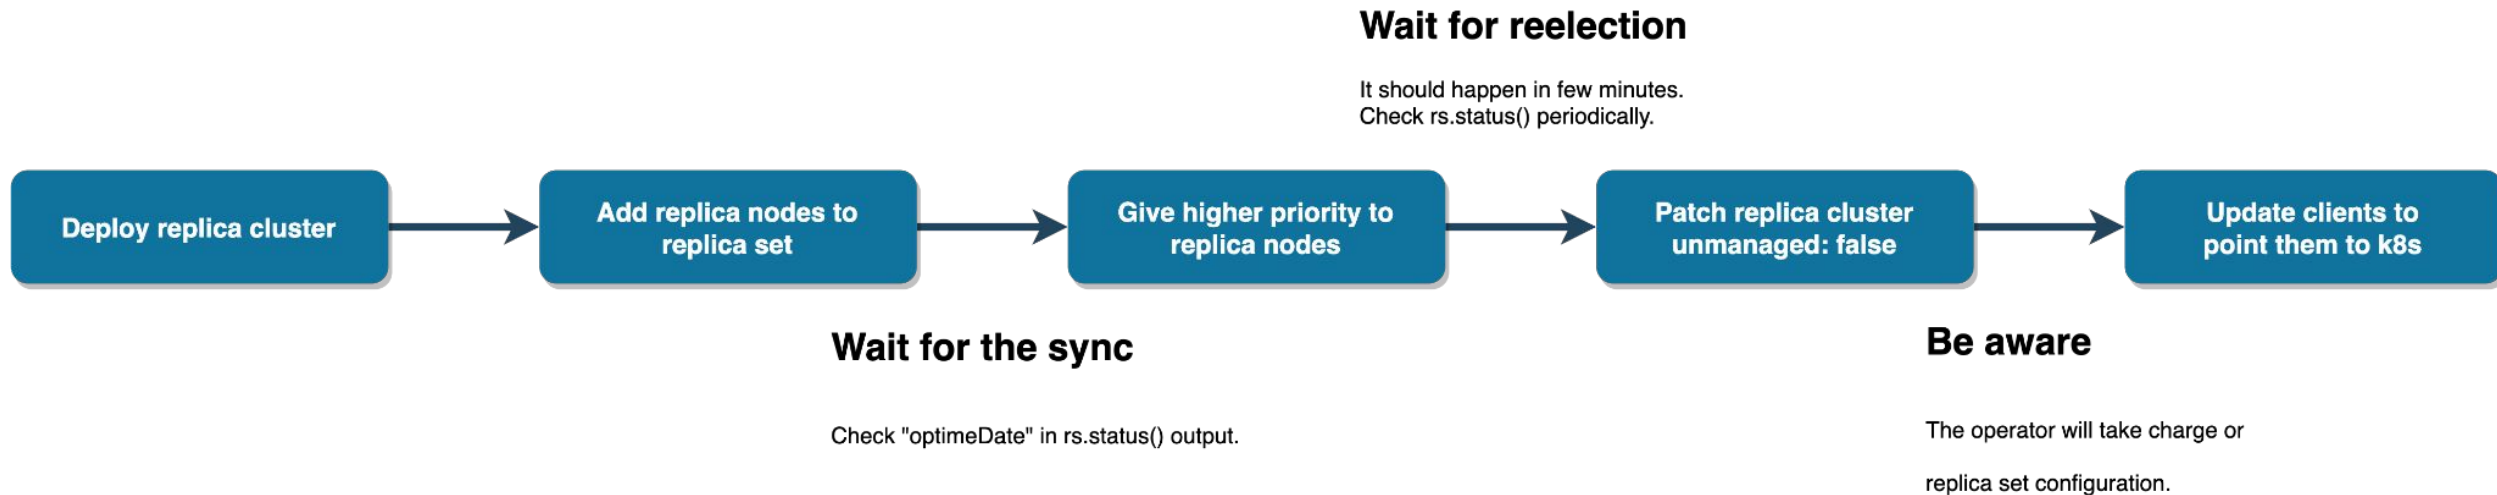
Feedback Wanted



- The feature is a Tech Preview at the moment.
- What you can do:
 - Write to us in [Forums](#)
 - Report bugs in [Jira](#)
 - Submit PRs in [GitHub](#)

Migrate to Kubernetes

How to migrate your DB to k8s?



How to migrate your DB to k8s?

Complete walkthrough: [Migrating MongoDB to Kubernetes](#)

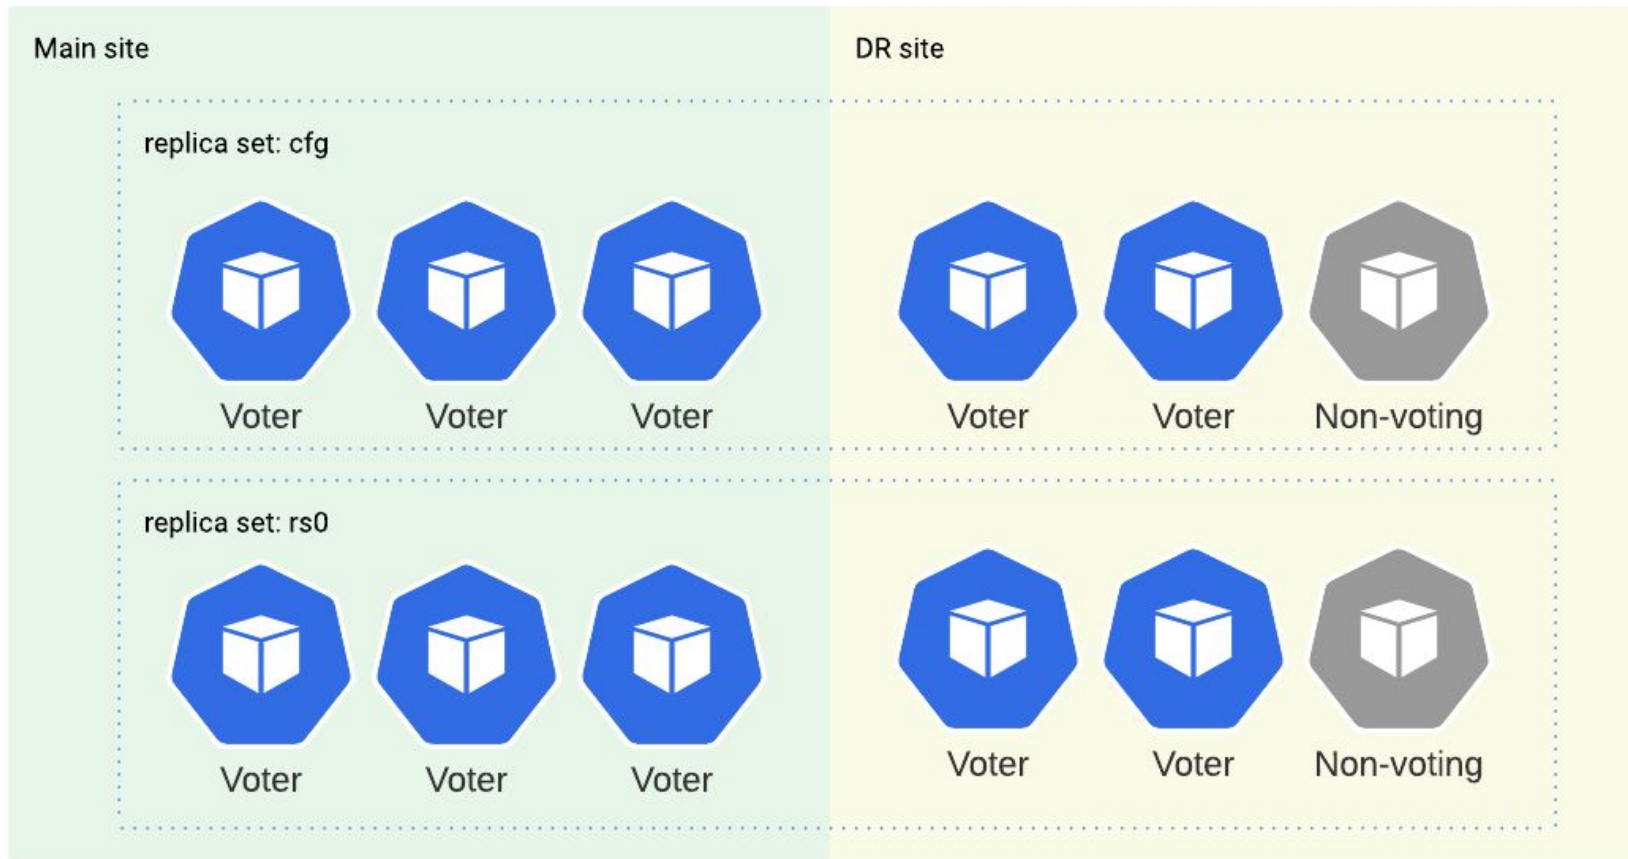
Set up a Disaster Recovery Site on Kubernetes

How to have a DR site on k8s?

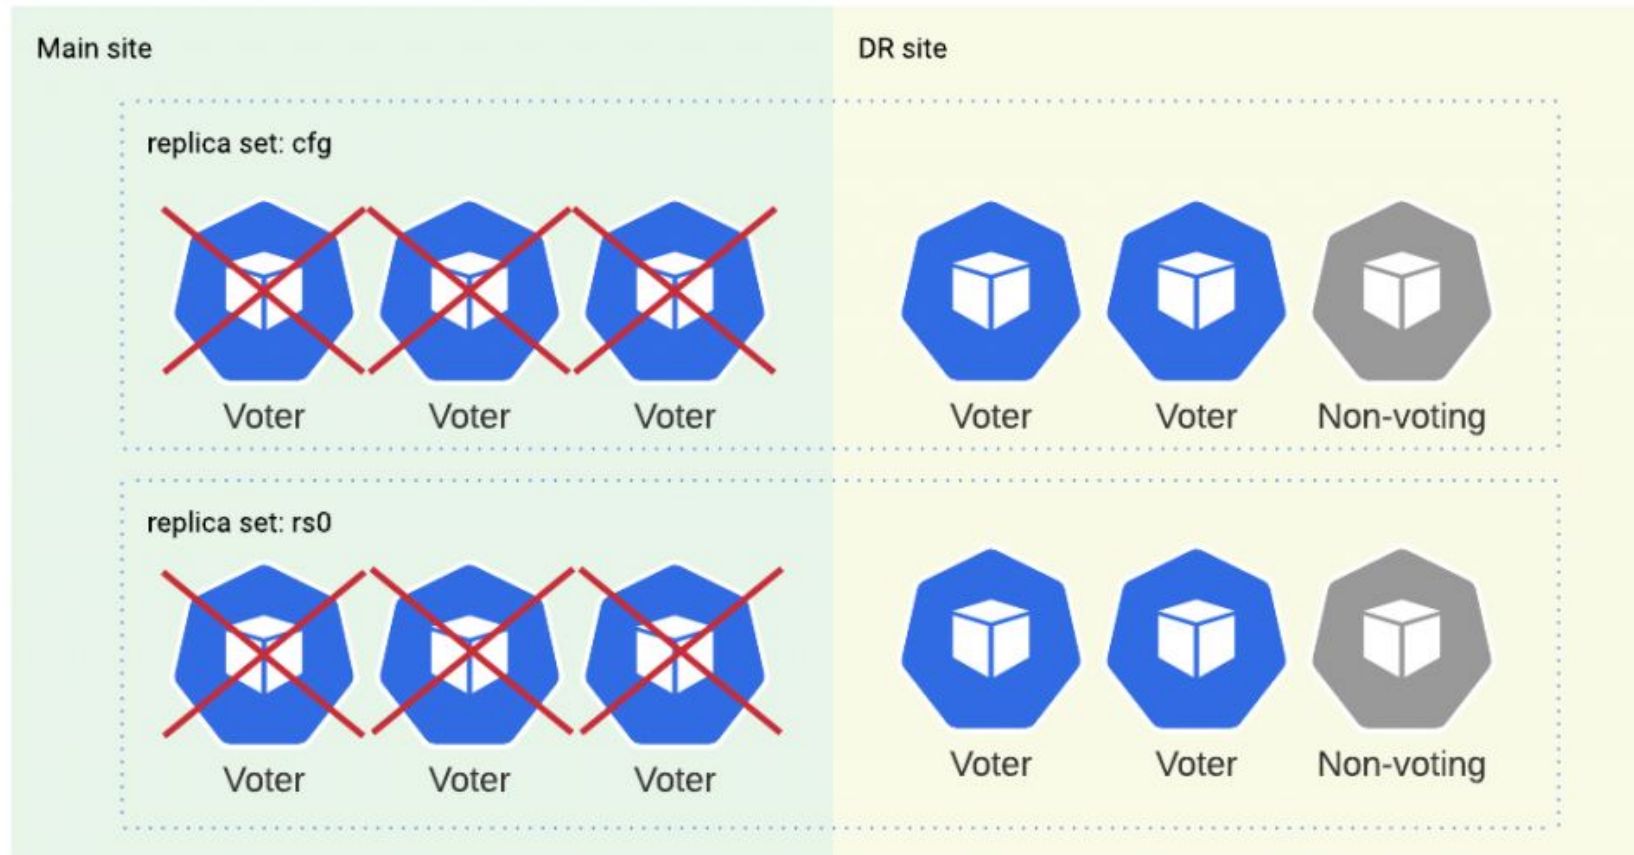
```
apiVersion: psmdb.percona.com/v1-11-0
kind: PerconaServerMongoDB
metadata:
  name: primary-cluster
spec:
  unmanaged: false
  replsets:
  - name: rs0
    externalNodes:
    - host: replica-rs0-0.ege.dev
      priority: 1
      votes: 1
    - host: replica-rs0-1.ege.dev
      priority: 1
      votes: 1
    - host: replica-rs0-2.ege.dev
      ...
```

- **We're giving priority 1 to replica nodes to prevent selecting a primary on DR site.**
- **We're giving votes to only 2 nodes to keep the number of voters odd.**

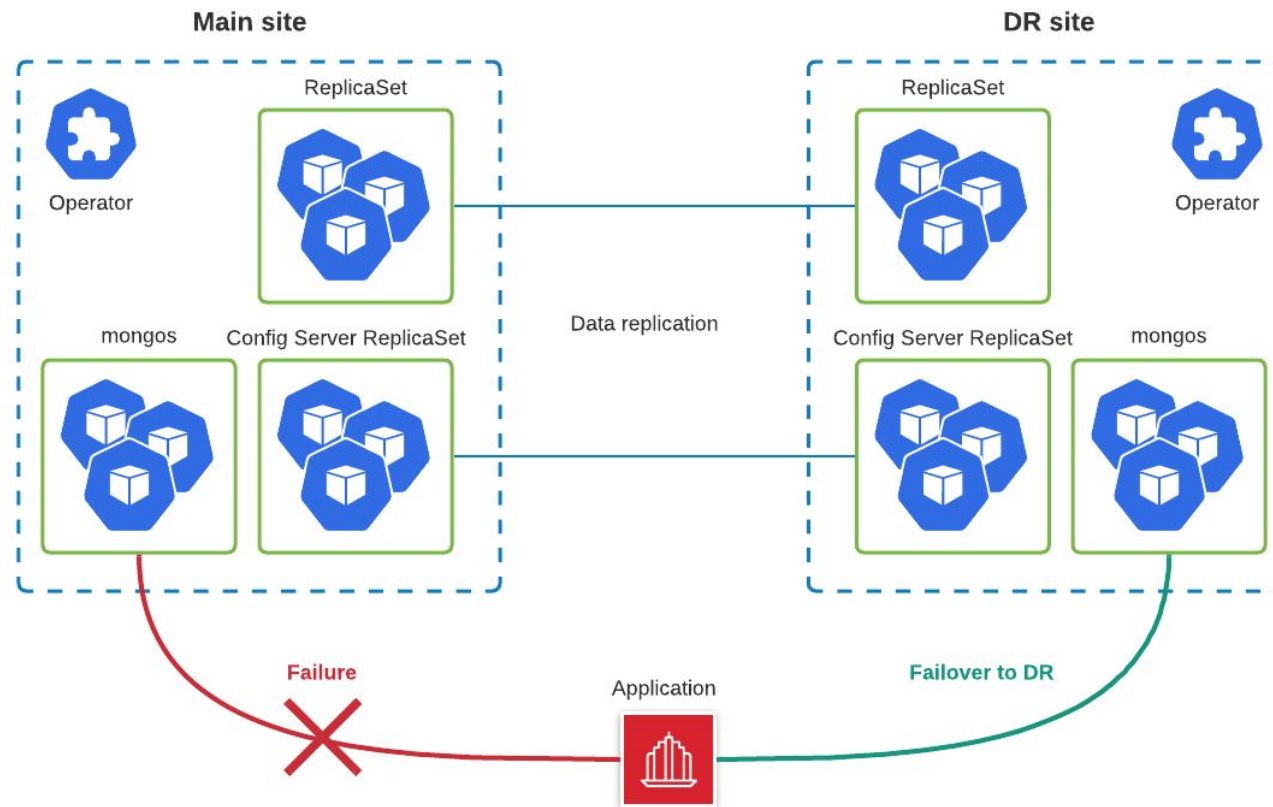
How to have a DR site on k8s?



How to have a DR site on k8s?



How to have a DR site on k8s?



How to have a DR site on k8s?

Complete walkthrough: [Disaster Recovery for MongoDB on Kubernetes](#)

What about the other operators?

- **Percona XtraDB Cluster Operator**

- Docs: [Set up Percona XtraDB Cluster cross-site replication](#)
- Blog: [Migration of a MySQL Database to a Kubernetes Cluster Using Asynchronous Replication](#)

- **Percona PostgreSQL Operator**

- Blog: [Migrating PostgreSQL to Kubernetes](#)

What if you need any help?

Write to us in Percona Community Forum:
<https://forums.percona.com/>

You can directly mention me: @egegunes

Thank you!



**Champions of Unbiased
Open Source Database Solutions**