Online OLTP computing and traffic governance as a service for true digital transformation

Trista Pan
panjuan@apache.org

Trista Pan

SphereEx Co-Founder & CTO

Apache Member

AWS Data Hero

Tencent Cloud TVP

Apache ShardingSphere PMC

Apache brpc (Incubating) & Apache AGE

& Apache HugeGraph (Incubating) mentor

China Mulan Community Mentor

Bio: https://tristazero.github.io

LinkedIn: https://www.linkedin.com/in/panjuan

GitHub: https://github.com/tristaZero

Twitter: @tristaZero

Project Twitter: @ShardingSphere

Content

- ✓ Background
- ✓ The needs for a database on the cloud
- ✓ Idea & architecture
- ✓ Handling SQL
- ✓ Demo

Digital transformation

- ✓ Leveraging novel technology
- ✓ Fundamentally change the delivery paradigm
- ✓ A continuous cultural change

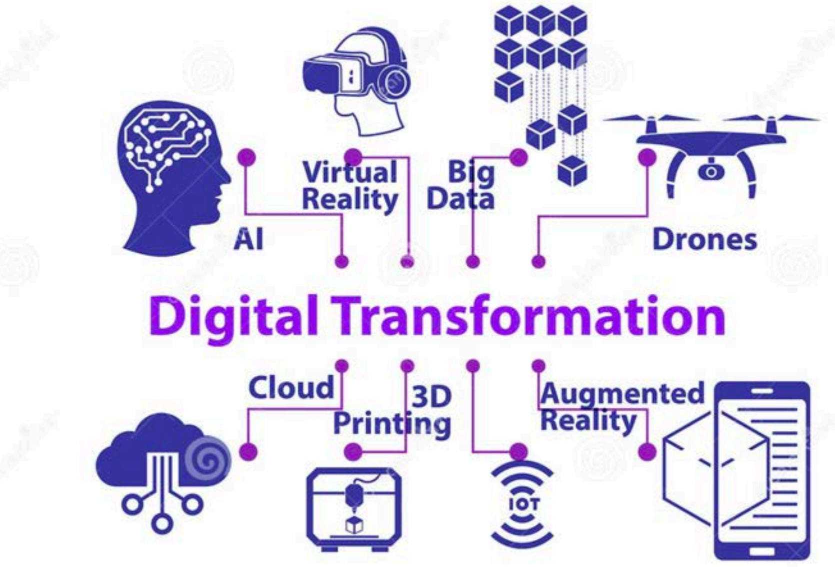

The needs for a database on the cloud

✓ Large data to manage

✓ Efficient queries

→ Data Sharding

✓ Traffic governance → HA & read/write splitting & traffic strategy

✓ Elastic scaling → Reshard for computing nodes and storage nodes

✓ Out-of-the-box deployment → Helm & Operator on Kubernetes

Database system

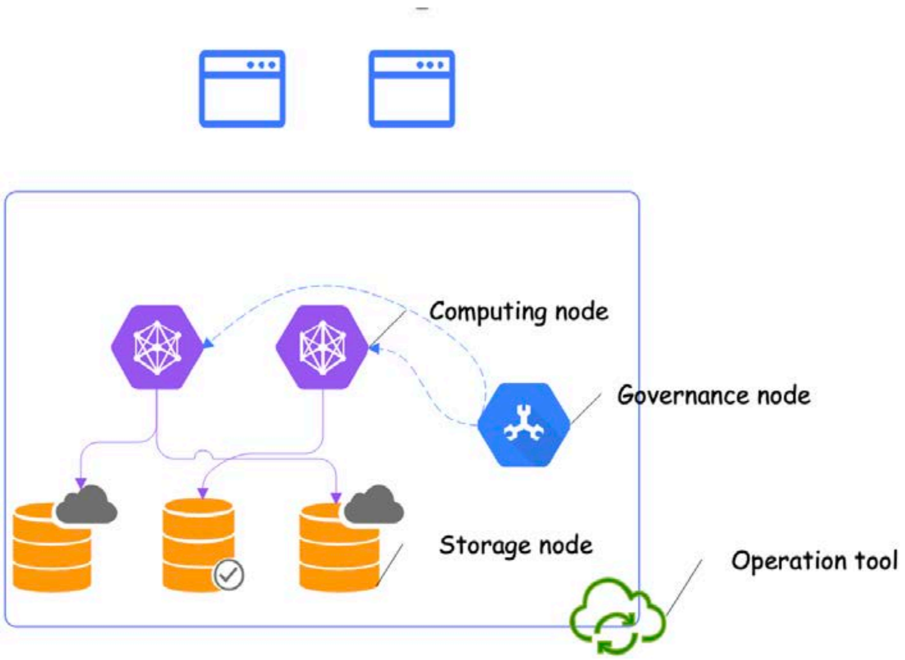

Apache ShardingSphere

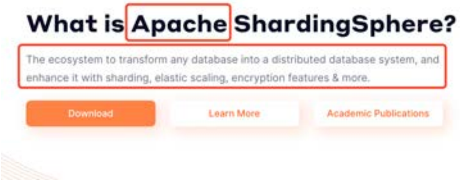

What is Apache ShardingSphere?

The ecosystem to transform any database into a distributed database system, and enhance it with sharding, elastic scaling, encryption features & more.

ShardingSphere

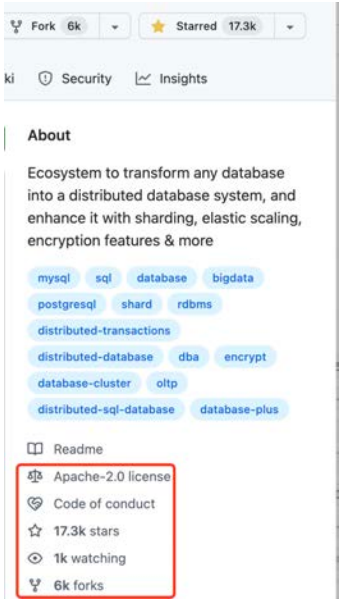

About

Ecosystem to transform any database into a distributed database system, and enhance it with sharding, elastic scaling, encryption features & more

- Apache-2.0 license
- Code of conduct
- 17.3k stars
- 1k watching
- 6k forks

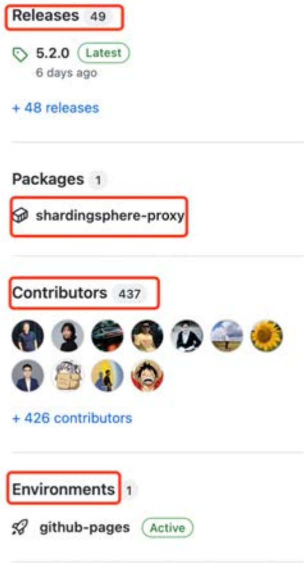

Releases 49

5.2.0 Latest
6 days ago

+ 48 releases

Packages 1

shardingsphere-proxy

Contributors 437

+ 426 contributors

Environments 1

github-pages Active

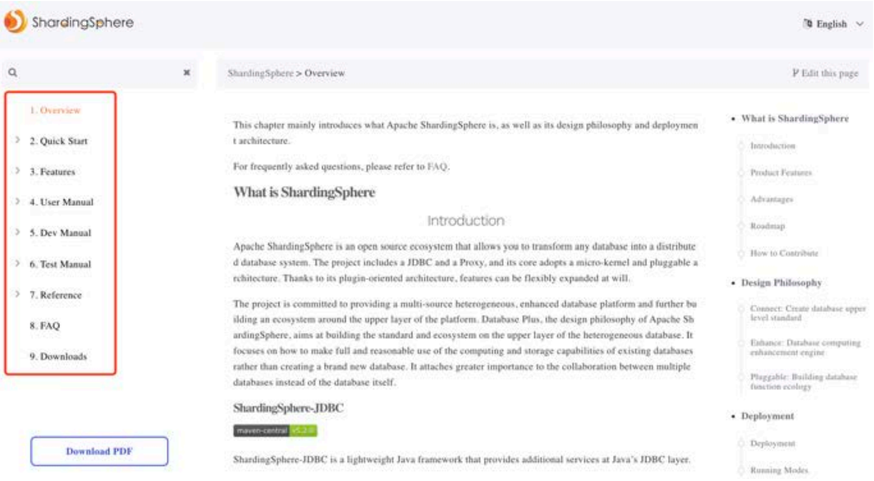

1. Overview
2. Quick Start
3. Features
4. User Manual
5. Dev Manual
6. Test Manual
7. Reference
8. FAQ
9. Downloads

This chapter mainly introduces what Apache ShardingSphere is, as well as its design philosophy and deployment architecture.

For frequently asked questions, please refer to FAQ.

What is ShardingSphere

Introduction

Apache ShardingSphere is an open source ecosystem that allows you to transform any database into a distributed database system. The project includes a JDBC and a Proxy, and its core adopts a micro-kernel and pluggable architecture. Thanks to its plugin-oriented architecture, features can be flexibly expanded at will.

The project is committed to providing a multi-source heterogeneous, enhanced database platform and further building an ecosystem around the upper layer of the platform. Database Plus, the design philosophy of Apache ShardingSphere, aims at building the standard and ecosystem on the upper layer of the heterogeneous database. It focuses on how to make full and reasonable use of the computing and storage capabilities of existing databases rather than creating a brand new database. It attaches greater importance to the collaboration between multiple databases instead of the database itself.

ShardingSphere-JDBC

ShardingSphere-JDBC is a lightweight Java framework that provides additional services at Java's JDBC layer.

- What is ShardingSphere
  - Introduction
  - Product Features
  - Advantages
  - Roadmap
  - How to Contribute
- Design Philosophy
  - Connect: Create database upper level standard
  - Enhance: Database computing enhancement engine
  - Pluggable: Building database function ecology
- Deployment
  - Deployment
  - Running Modes

ShardingSphere clients

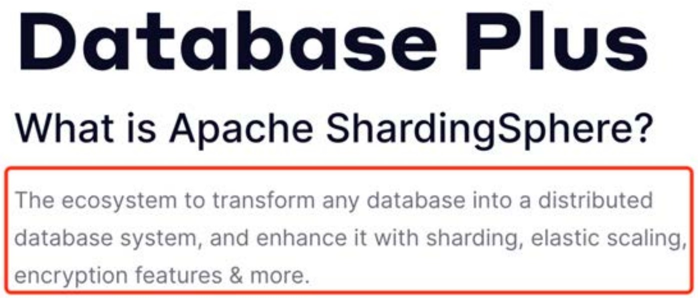

Database Plus

What is Apache ShardingSphere?

The ecosystem to transform any database into a distributed database system, and enhance it with sharding, elastic scaling, encryption features & more.

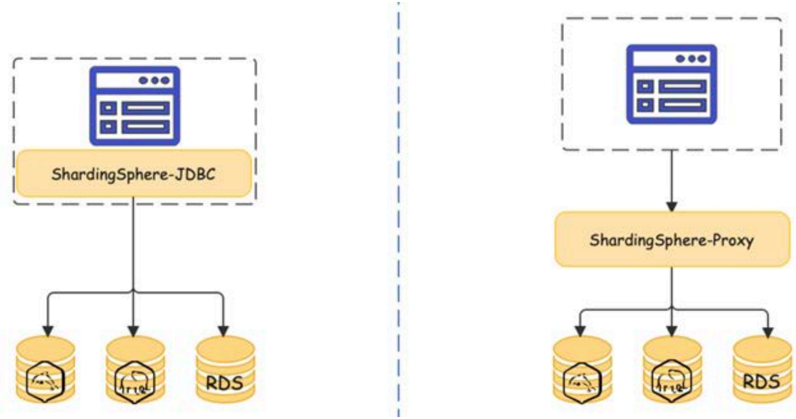

ShardingSphere features

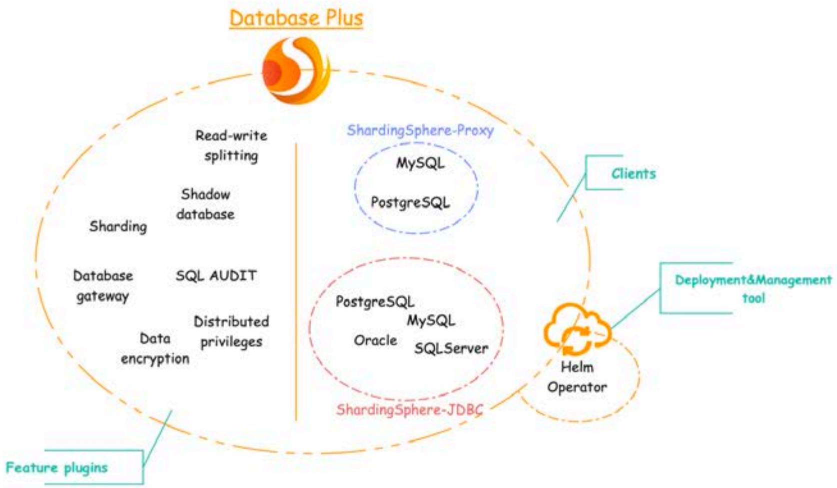

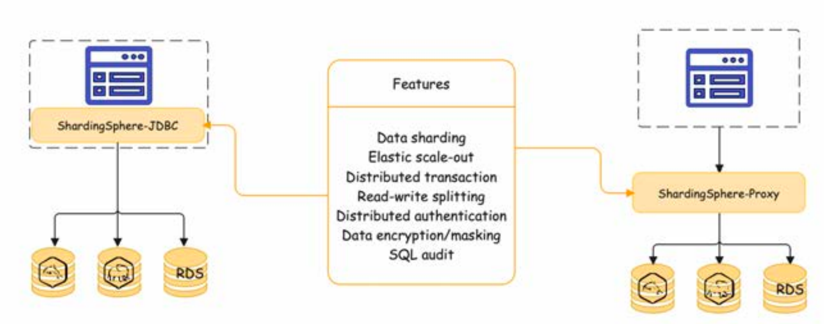

Solution

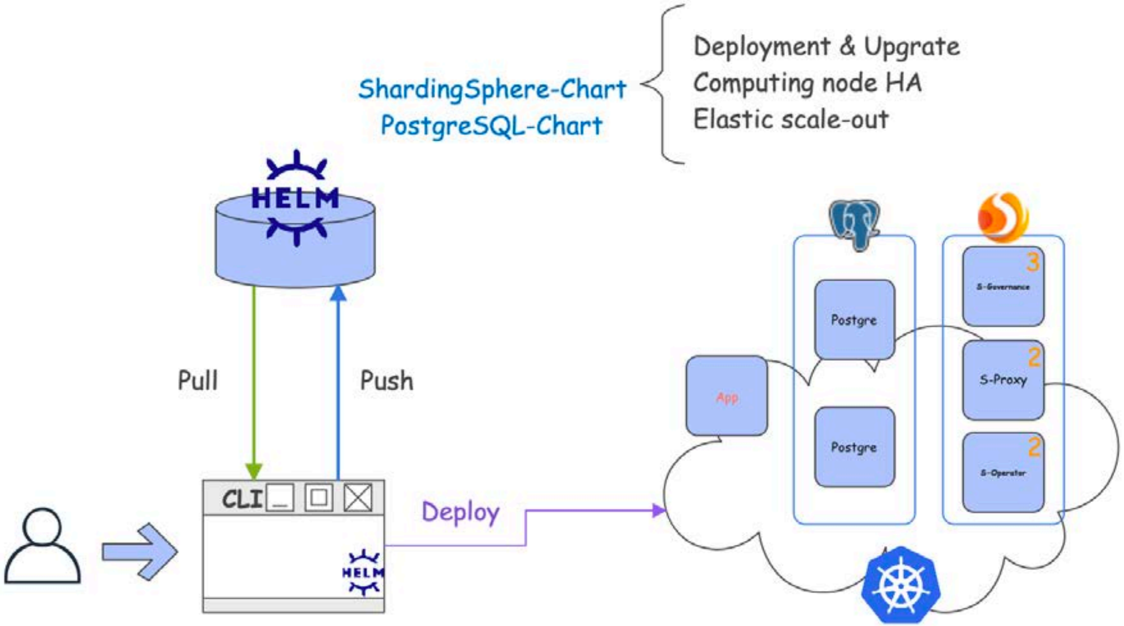

ShardingSphere-Chart
PostgreSQL-Chart

Deployment & Upgrate
Computing node HA
Elastic scale-out

Solution

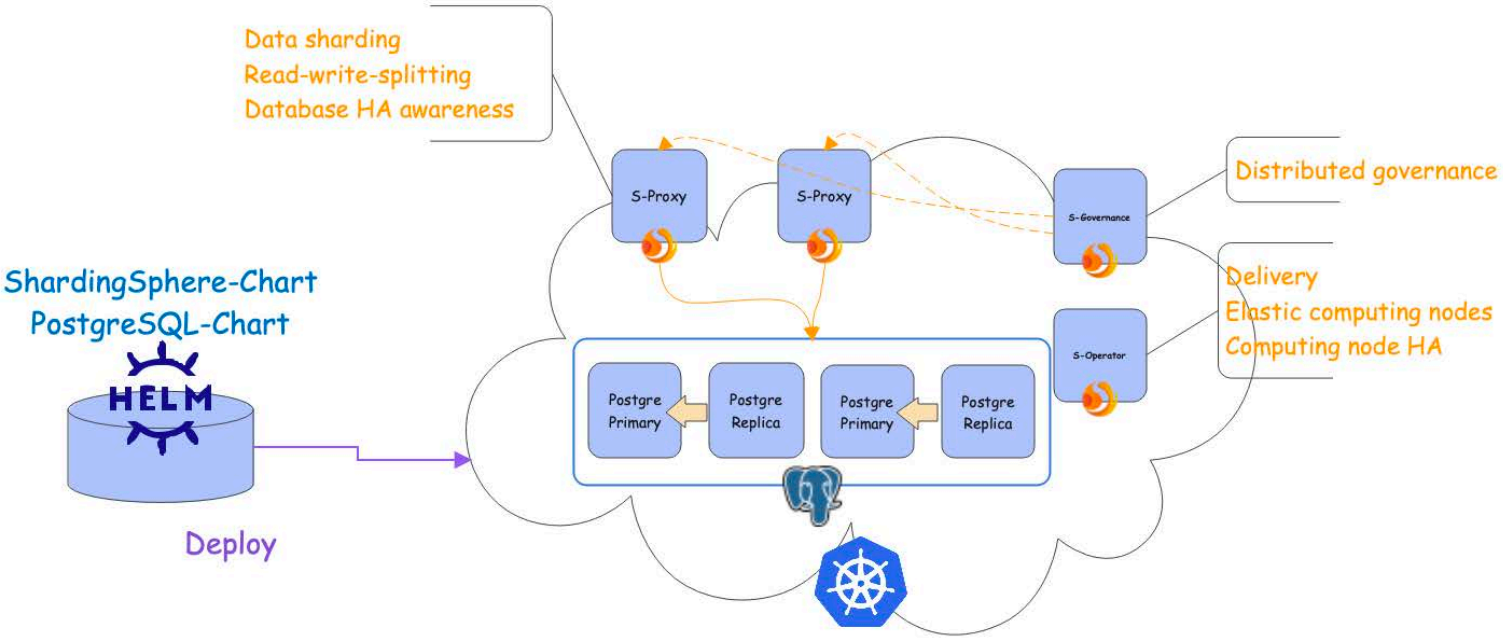

Data sharding
Read-write-splitting
Database HA awareness

S-Proxy

S-Proxy

S-Governance

Distributed governance

S-Operator

Delivery
Elastic computing nodes
Computing node HA

ShardingSphere-Chart
PostgreSQL-Chart

Deploy

The handling process of one SQL

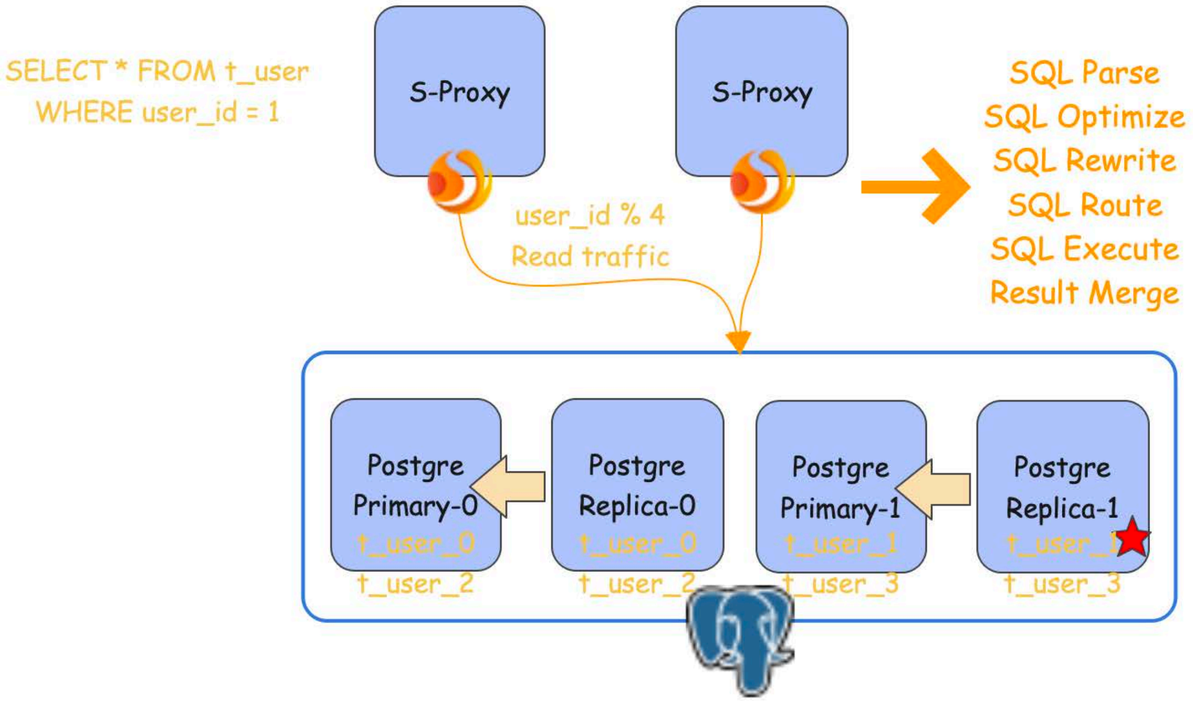

SphereEx

The demo show

1. Deploy two PostgreSQL (Storage node) clusters made of a primary node and a replica
2. Deploy ShardingSphere-Operator
3. Deploy two ShardingSphere-Proxy (Computing node) and ShardingSphere-governance
4. Test ShardingSphere-Operator to guarantee the HA of ShardingSphere-Proxy
5. Add PostgreSQL resources and their relationship into ShardingSphere-Proxy
6. Create sharding table t_user on ShardingSphere-Proxy
7. Show the metadata of this distributed database system
8. INSERT data for test on ShardingSphere-Proxy
9. Preview SELECT routing result
10. Execute SELECT query

Step 1, 2, 3

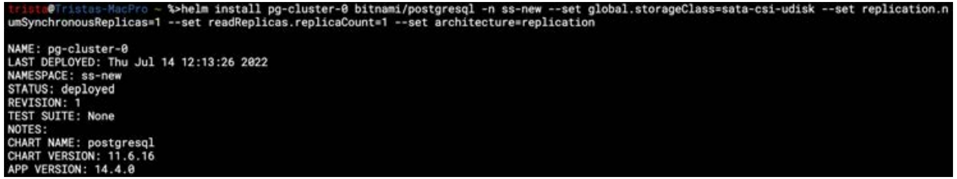

```
trista@Tristas-MacPro ~ %>helm install pg-cluster-0 bitnami/postgresql -n ss-new --set global.storageClass=sata-csi-udisk --set replication.n
umSynchronousReplicas=1 --set readReplicas.replicaCount=1 --set architecture=replication

NAME: pg-cluster-0
LAST DEPLOYED: Thu Jul 14 12:13:26 2022
NAMESPACE: ss-new
STATUS: deployed
REVISION: 1
TEST SUITE: None
NOTES:
CHART NAME: postgresql
CHART VERSION: 11.6.16
APP VERSION: 14.4.0
```

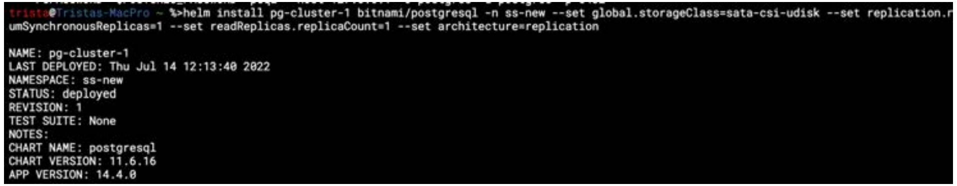

```
trista@Tristas-MacPro ~ %>helm install pg-cluster-1 bitnami/postgresql -n ss-new --set global.storageClass=sata-csi-udisk --set replication.r
umSynchronousReplicas=1 --set readReplicas.replicaCount=1 --set architecture=replication

NAME: pg-cluster-1
LAST DEPLOYED: Thu Jul 14 12:13:40 2022
NAMESPACE: ss-new
STATUS: deployed
REVISION: 1
TEST SUITE: None
NOTES:
CHART NAME: postgresql
CHART VERSION: 11.6.16
APP VERSION: 14.4.0
```

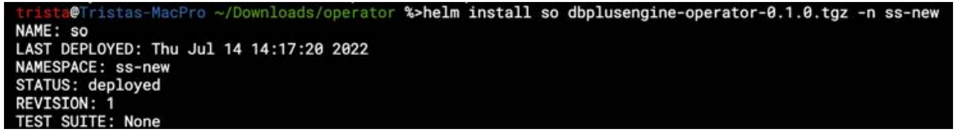

```
trista@Tristas-MacPro ~/Downloads/operator %>helm install so dbplusengine-operator-0.1.0.tgz -n ss-new
NAME: so
LAST DEPLOYED: Thu Jul 14 14:17:20 2022
NAMESPACE: ss-new
STATUS: deployed
REVISION: 1
TEST SUITE: None
```

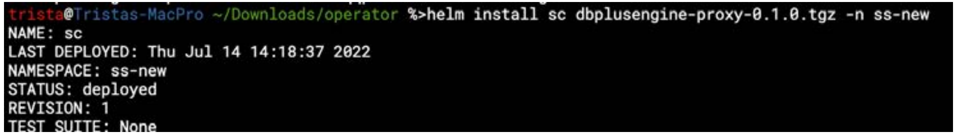

```
trista@Tristas-MacPro ~/Downloads/operator %>helm install sc dbplusengine-proxy-0.1.0.tgz -n ss-new
NAME: sc
LAST DEPLOYED: Thu Jul 14 14:18:37 2022
NAMESPACE: ss-new
STATUS: deployed
REVISION: 1
TEST SUITE: None
```

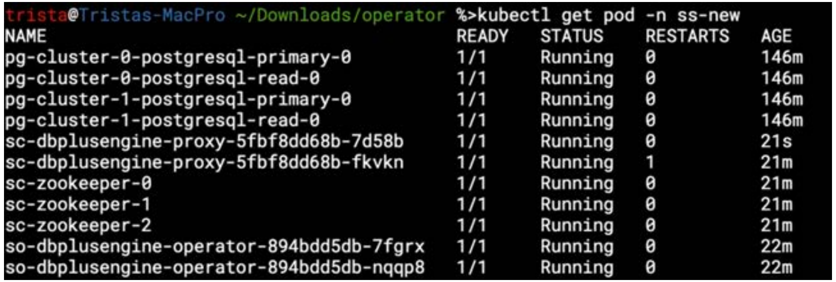

```
trista@Tristas-MacPro ~/Downloads/operator %>kubectl get pod -n ss-new
NAME                                         READY   STATUS    RESTARTS   AGE
pg-cluster-0-postgresql-primary-0            1/1     Running   0          146m
pg-cluster-0-postgresql-read-0               1/1     Running   0          146m
pg-cluster-1-postgresql-primary-0            1/1     Running   0          146m
pg-cluster-1-postgresql-read-0               1/1     Running   0          146m
sc-dbplusengine-proxy-5fbf8dd68b-7d58b       1/1     Running   0          21s
sc-dbplusengine-proxy-5fbf8dd68b-fkvkn       1/1     Running   1          21m
sc-zookeeper-0                               1/1     Running   0          21m
sc-zookeeper-1                               1/1     Running   0          21m
sc-zookeeper-2                               1/1     Running   0          21m
so-dbplusengine-operator-894bdd5db-7fgrx     1/1     Running   0          22m
so-dbplusengine-operator-894bdd5db-nqqp8     1/1     Running   0          22m
```

SphereEx

Step 4

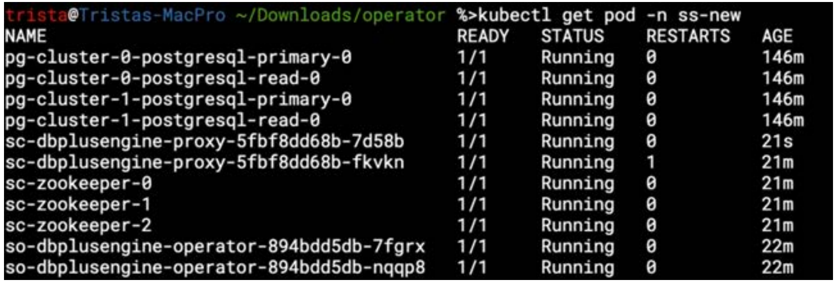

```
trista@Tristas-MacPro ~/Downloads/operator %>kubectl get pod -n ss-new
NAME                                       READY   STATUS    RESTARTS   AGE
pg-cluster-0-postgresql-primary-0          1/1     Running   0          146m
pg-cluster-0-postgresql-read-0             1/1     Running   0          146m
pg-cluster-1-postgresql-primary-0          1/1     Running   0          146m
pg-cluster-1-postgresql-read-0             1/1     Running   0          146m
sc-dbplusengine-proxy-5fbf8dd68b-7d58b     1/1     Running   0          21s
sc-dbplusengine-proxy-5fbf8dd68b-fkvkn     1/1     Running   1          21m
sc-zookeeper-0                             1/1     Running   0          21m
sc-zookeeper-1                             1/1     Running   0          21m
sc-zookeeper-2                             1/1     Running   0          21m
so-dbplusengine-operator-894bdd5db-7fgrx   1/1     Running   0          22m
so-dbplusengine-operator-894bdd5db-nqqp8   1/1     Running   0          22m
```

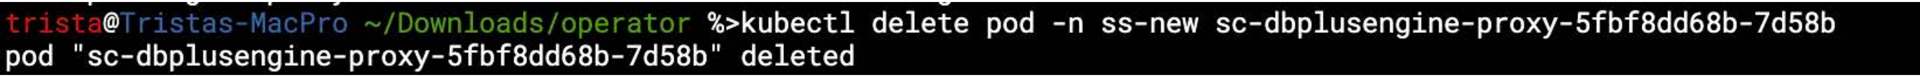

```
trista@Tristas-MacPro ~/Downloads/operator %>kubectl delete pod -n ss-new sc-dbplusengine-proxy-5fbf8dd68b-7d58b
pod "sc-dbplusengine-proxy-5fbf8dd68b-7d58b" deleted
```

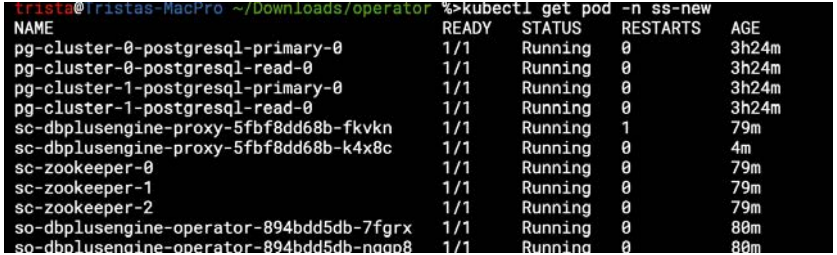

```
trista@Tristas-MacPro ~/Downloads/operator %>kubectl get pod -n ss-new
NAME                                       READY   STATUS    RESTARTS   AGE
pg-cluster-0-postgresql-primary-0          1/1     Running   0          3h24m
pg-cluster-0-postgresql-read-0             1/1     Running   0          3h24m
pg-cluster-1-postgresql-primary-0          1/1     Running   0          3h24m
pg-cluster-1-postgresql-read-0             1/1     Running   0          3h24m
sc-dbplusengine-proxy-5fbf8dd68b-fkvkn     1/1     Running   1          79m
sc-dbplusengine-proxy-5fbf8dd68b-k4x8c     1/1     Running   0          4m
sc-zookeeper-0                             1/1     Running   0          79m
sc-zookeeper-1                             1/1     Running   0          79m
sc-zookeeper-2                             1/1     Running   0          79m
so-dbplusengine-operator-894bdd5db-7fgrx   1/1     Running   0          80m
so-dbplusengine-operator-894bdd5db-nqqp8   1/1     Running   0          80m
```

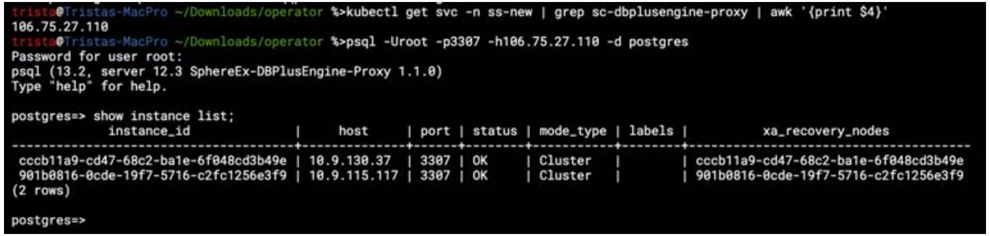

```
trista@Tristas-MacPro ~/Downloads/operator %>kubectl get svc -n ss-new | grep sc-dbplusengine-proxy | awk '{print $4}'
106.75.27.110
trista@Tristas-MacPro ~/Downloads/operator %>psql -Uroot -p3307 -h106.75.27.110 -d postgres
Password for user root:
psql (13.2, server 12.3 SphereEx-DBPlusEngine-Proxy 1.1.0)
Type "help" for help.

postgres=> show instance list;
             instance_id              |     host     | port | status | mode_type | labels |          xa_recovery_nodes
--------------------------------------+--------------+------+--------+-----------+--------+--------------------------------------
 cccb11a9-cd47-68c2-ba1e-6f048cd3b49e | 10.9.130.37  | 3307 | OK     | Cluster   |        | cccb11a9-cd47-68c2-ba1e-6f048cd3b49e
 901b0816-0cde-19f7-5716-c2fc1256e3f9 | 10.9.115.117 | 3307 | OK     | Cluster   |        | 901b0816-0cde-19f7-5716-c2fc1256e3f9
(2 rows)

postgres=>
```

SphereEx

Step 5, 6, 7

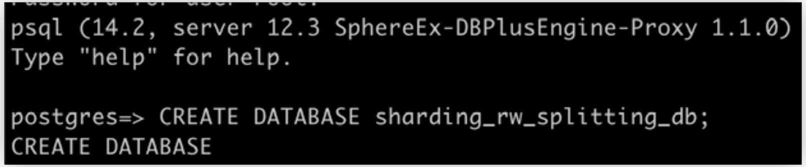

```
psql (14.2, server 12.3 SphereEx-DBPlusEngine-Proxy 1.1.0)
Type "help" for help.

postgres=> CREATE DATABASE sharding_rw_splitting_db;
CREATE DATABASE
```

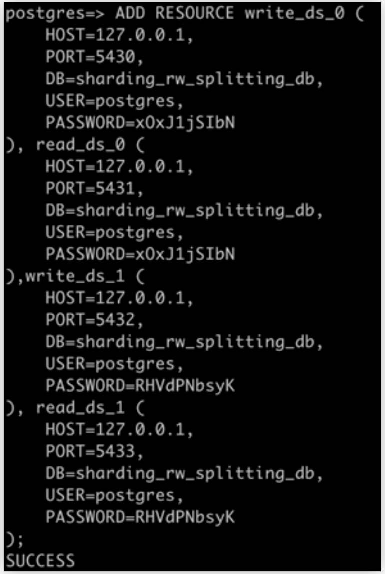

```
postgres=> ADD RESOURCE write_ds_0 (
    HOST=127.0.0.1,
    PORT=5430,
    DB=sharding_rw_splitting_db,
    USER=postgres,
    PASSWORD=xOxJ1jSIbN
), read_ds_0 (
    HOST=127.0.0.1,
    PORT=5431,
    DB=sharding_rw_splitting_db,
    USER=postgres,
    PASSWORD=xOxJ1jSIbN
),write_ds_1 (
    HOST=127.0.0.1,
    PORT=5432,
    DB=sharding_rw_splitting_db,
    USER=postgres,
    PASSWORD=RHVdPNbsyK
), read_ds_1 (
    HOST=127.0.0.1,
    PORT=5433,
    DB=sharding_rw_splitting_db,
    USER=postgres,
    PASSWORD=RHVdPNbsyK
);
SUCCESS
```

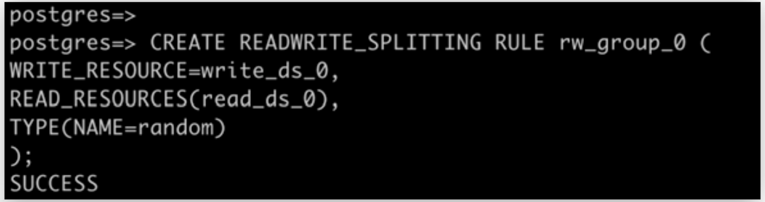

```
postgres=>
postgres=> CREATE READWRITE_SPLITTING RULE rw_group_0 (
WRITE_RESOURCE=write_ds_0,
READ_RESOURCES(read_ds_0),
TYPE(NAME=random)
);
SUCCESS
```

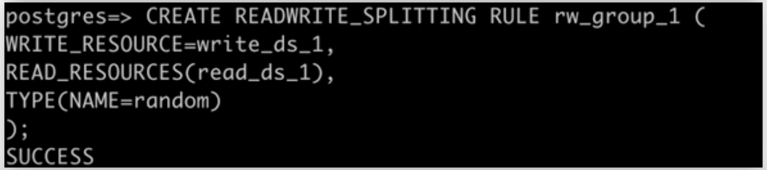

```
postgres=> CREATE READWRITE_SPLITTING RULE rw_group_1 (
WRITE_RESOURCE=write_ds_1,
READ_RESOURCES(read_ds_1),
TYPE(NAME=random)
);
SUCCESS
```

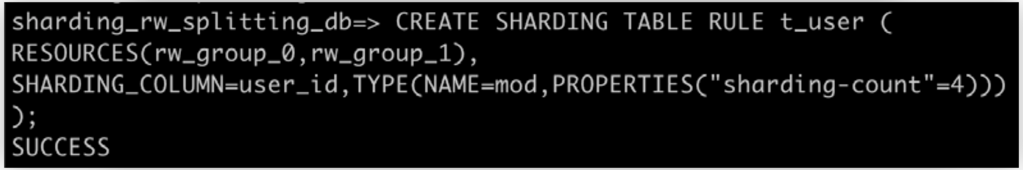

```
sharding_rw_splitting_db=> CREATE SHARDING TABLE RULE t_user (
RESOURCES(rw_group_0,rw_group_1),
SHARDING_COLUMN=user_id,TYPE(NAME=mod,PROPERTIES("sharding-count"=4)))
);
SUCCESS
```

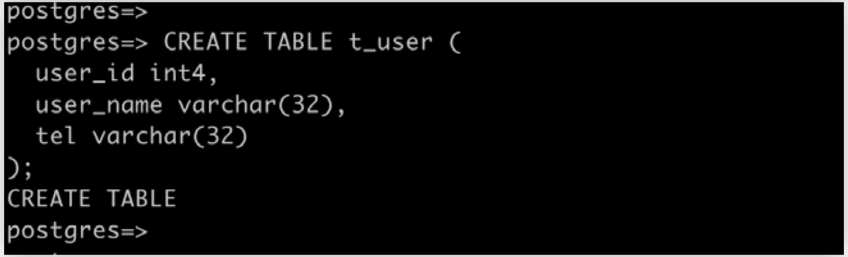

```
postgres=>
postgres=> CREATE TABLE t_user (
  user_id int4,
  user_name varchar(32),
  tel varchar(32)
);
CREATE TABLE
postgres=>
```

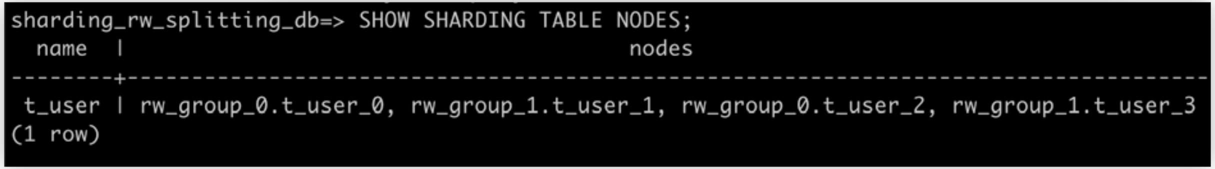

```
sharding_rw_splitting_db=> SHOW SHARDING TABLE NODES;
  name  |                                      nodes
--------+--------------------------------------------------------------------------------------
 t_user | rw_group_0.t_user_0, rw_group_1.t_user_1, rw_group_0.t_user_2, rw_group_1.t_user_3
(1 row)
```

Step 8, 9, 10

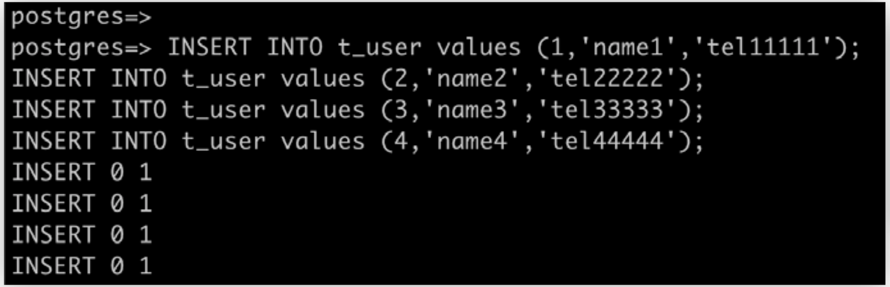

```
postgres=>
postgres=> INSERT INTO t_user values (1,'name1','tel11111');
INSERT INTO t_user values (2,'name2','tel22222');
INSERT INTO t_user values (3,'name3','tel33333');
INSERT INTO t_user values (4,'name4','tel44444');
INSERT 0 1
INSERT 0 1
INSERT 0 1
INSERT 0 1
```

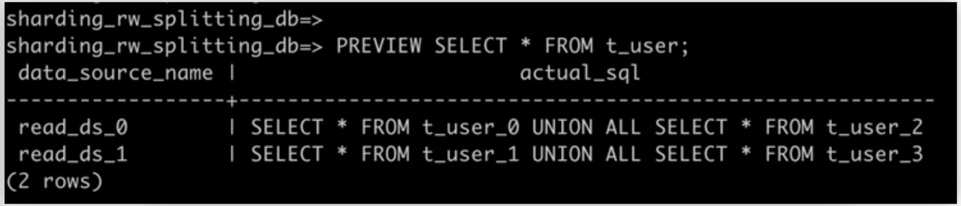

```
sharding_rw_splitting_db=>
sharding_rw_splitting_db=> PREVIEW SELECT * FROM t_user;
 data_source_name |                        actual_sql
------------------+----------------------------------------------------------
 read_ds_0        | SELECT * FROM t_user_0 UNION ALL SELECT * FROM t_user_2
 read_ds_1        | SELECT * FROM t_user_1 UNION ALL SELECT * FROM t_user_3
(2 rows)
```

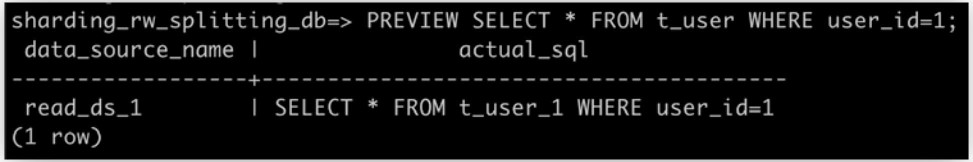

```
sharding_rw_splitting_db=> PREVIEW SELECT * FROM t_user WHERE user_id=1;
 data_source_name |              actual_sql
------------------+---------------------------------------
 read_ds_1        | SELECT * FROM t_user_1 WHERE user_id=1
(1 row)
```

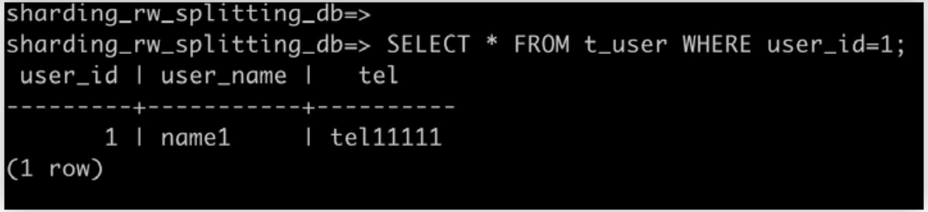

```
sharding_rw_splitting_db=>
sharding_rw_splitting_db=> SELECT * FROM t_user WHERE user_id=1;
 user_id | user_name |   tel
---------+-----------+----------
       1 | name1     | tel11111
(1 row)
```

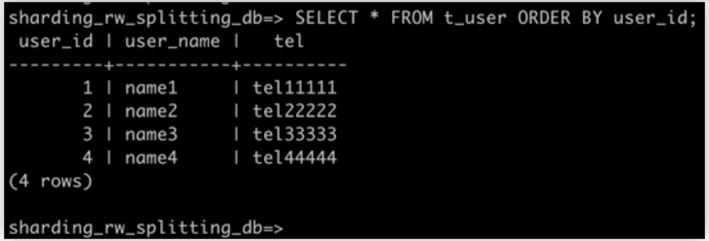

```
sharding_rw_splitting_db=> SELECT * FROM t_user ORDER BY user_id;
 user_id | user_name |   tel
---------+-----------+----------
       1 | name1     | tel11111
       2 | name2     | tel22222
       3 | name3     | tel33333
       4 | name4     | tel44444
(4 rows)

sharding_rw_splitting_db=>
```

Thanks!
Any questions?

Bio: https://tristazero.github.io

LinkedIn: https://www.linkedin.com/in/panjuan

GitHub: https://github.com/tristaZero

Twitter: @tristaZero

Project Twitter: @ShardingSphere

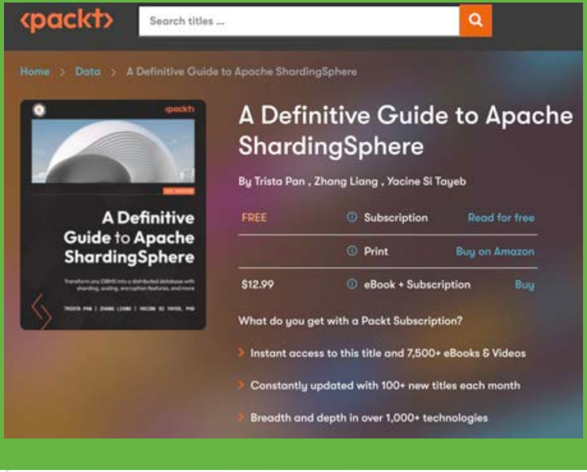

Home › Data › A Definitive Guide to Apache ShardingSphere

A Definitive Guide to Apache ShardingSphere

By Trista Pan , Zhang Liang , Yacine Si Tayeb

FREE	Subscription	Read for free
	Print	Buy on Amazon
$12.99	eBook + Subscription	Buy

What do you get with a Packt Subscription?

- Instant access to this title and 7,500+ eBooks & Videos
- Constantly updated with 100+ new titles each month
- Breadth and depth in over 1,000+ technologies

Apache ShardingSphere @ShardingSphere · Jul 20

We're celebrating the launch with a #Giveaway for this book. Get a chance to win a copy!

All you have to do is like and RT with your preference i.e. print book or ebook (print only for UK, US & Canada so give us a hint or DM us)

#opensource #database

If you would like to learn more about Apache ShardingSphere, Please contact it at Twitter to get an author copy as a giveaway!

SphereEx