

ShardingSphereProposal

Abstract

Sharding-Sphere is an ecosystem of transparent distributed database middleware, focusing on data sharding, distributed transaction and database orchestration. It provides maximum compatibility for applications through Sharding-JDBC (a driver to implement JDBC) or Sharding-Proxy (a proxy to implement database protocol).

Proposal

With a large number of end users, Sharding-Sphere has a fairly huge community in China. It is also widely adopted by many [companies and organizations](#) as a solution to process their massive amounts of data.

We believe that bringing Sharding-Sphere into Apache Software Foundation could advance development of a stronger and more diverse open source community.

Dangdang submits this proposal to donate Sharding-Sphere's source codes and all related documentations to Apache Software Foundation. The codes are already under Apache License Version 2.0.

- Code base: <https://github.com/sharding-sphere/sharding-sphere>
- Web site: <http://shardingsphere.io/>
- Documentations: <http://shardingsphere.io/document/current/>
- Community: <http://shardingsphere.io/community/>

Background

Relational database hardly supports such huge amounts of data any more which has increased rapidly in recent years, but for reason of technique maturity, developers and DBAs still want to use it to persist core data.

Sharding-Sphere was open sourced on Github in 2016. At the very beginning, Sharding-Sphere is just a JDBC driver for data sharding (name as Sharding-JDBC) at Dangdang internal framework; now it offers data sharding, distributed transaction and database orchestration. Besides JDBC, proxy to implement MySQL database protocol is also supported at present. Furthermore, our roadmap includes Proxy for PostgreSQL protocol, Sidecar model, data replica and elastic data scalability function as well.

Due to the extension of project, we provide proxy model and sidecar model in addition to JDBC model. Therefore, we rename it to Sharding-Sphere by [a public vote](#), which refers to a sharding ecosphere with Sharding-JDBC, Sharding-Proxy and Sharding Sidecar as its three sub-projects.

Sharding-JDBC has won the [TOP 20 most popular open source projects in China 2016](#).

Rationale

Relational database still plays a very important role on current application system. The maturity of products and surrounding ecosystem, the friendliness of its data query and developers' and DBAs' mastery of it, cannot be completely replaced with other types of database in the near future. However, current relational database cannot support cloud native very well and it is not friendly to distributed system.

It is the ultimate goal of Sharding-Sphere, which manages the databases scattering around the system, to make user use distributed databases as simply as using a single one.

Without extra cost, Sharding-JDBC directly connects database with Java application to get the best performance.

Sharding-Proxy is deployed as a stateless server and supports MySQL protocol at present. In the paper [What's Really New with NewSQL?](#), three types of NewSQL are introduced, among which Sharding-Proxy is a Transparent Sharding Middleware.

Sharding-Sidecar can be understood as a data panel in Service Mesh. The interaction between the application and the database provides a mesh layer. The concept of Database Mesh is similar to Service Mesh, and it focuses on how to connect data access applications to the database. Database Mesh will set up a huge grid system between applications and databases. Applications and databases need be placed in the grid system. They are all objects managed by the meshing layer.

Current Status

Meritocracy

Sharding-Sphere was incubated at Dangdang in 2015 and open sourced on [GitHub](#) in 2016. In 2017, Jingdong recognized its value and determined to sponsor this project. Sharding-Sphere has contributors and users from many companies; we have set up the PMC Team and Committer Team. New contributors are guided and reviewed by existed PMC members. When they are ready, PMC will start a vote to promote him/her to become a member of PMC or Committer Team. See the details [here](#). Contributions are always welcomed and highly valued.

Community

Now we have set development teams for Sharding-Sphere respectively in Jingdong, Dangdang and Bestpay. Companies like Sohu and Dataman have shown great interest in Sharding-Sphere. We hope to grow the base of contributors by inviting all those who offer contributions through The Apache Way. Right now, we make use of github as code hosting as well as gitter for community communication.

Core Developers

The core developers, including experienced open source developers and team leaders, have formed a group full of diversity.

Known Risks

Orphaned products

Three development teams from Jingdong, Dangdang and CHINA TELECOM Bestpay will spare no pains to work on Sharding-Sphere in the future with contributors from the growing community. Also, Sharding-Sphere is widely adopted in China by many [companies and organizations](#). Thus, it is very unlikely that Sharding-Sphere becomes orphaned.

Inexperience with Open Source

The current core developers all work for companies that have developed or contributed to many open source projects.

- Liang Zhang, PMC. He is the lead of two open source projects, Sharding-Sphere and Elastic-Job. Also Committer of Apache Dubbo (Incubator).
- Sheng Wu, PMC. He is the PPMC and committer of Apache [SkyWalking](#)(Incubating), Apache Zipkin(Incubating) contributor, CNCF [OpenTracing](#) member. Also contributed a lot other open source projects.
- Hongtao Gao, PMC. He is the PPMC and committer of Apache [SkyWalking](#)(Incubating) too. Also contribute a lot of features of Elastic-Job.

Therefore, we believe we have enough experience to deal with open source.

Homogenous Developers

The current core developers work across a variety of organizations including Jingdong, Dangdang and CHINA TELECOM Bestpay; some individual developers are accepted as core developers of Sharding-Sphere as well. Considering that Sohu and Dataman have shown great interest in Sharding-Sphere, we plan to encourage them to contribute and invite them as contributors to work together.

Reliance on Salaried Developers

At present, three of the core developers are paid by their employer to contribute to Sharding-Sphere project. It is estimated that the development of Sharding-Sphere will be continued with mainly salaried developers, and we will make efforts to attract more volunteers and grow the community.

Relationships with Other Apache Products

An automatic prober of Sharding-Sphere is introduced into [SkyWalking](#) to send APM data, and [SkyWalking](#) also use Sharding-Sphere to persist tracing data. Saga provided by [ServiceComb](#) is adopted by Sharding-Sphere as one of the distributed transaction processing engines. Sharding-Sphere integrates Apache Zookeeper as one of the service registration/discovery mechanisms.

A Excessive Fascination with the Apache Brand

We acknowledge the value and reputation that the Apache brand would bring to Sharding-Sphere. However, our primary interest is in the excellent community provided by Apache Software Foundation, in which all the projects could gain stability for long-term development.

Documentation

A complete set of Sharding-Sphere documentations is provided on shardingsphere.io in both English and Simplified Chinese.

- [English](#)
- [Chinese](#)

Initial Source

The project consists of three distinct codebases: core, example and document. The address of three existed git repositories are as follows:

- <https://github.com/sharding-sphere/sharding-sphere>
- <https://github.com/sharding-sphere/sharding-sphere-example>

- <https://github.com/sharding-sphere/sharding-sphere-doc>

Source and Intellectual Property Submission Plan

The codes are currently under Apache License Version 2.0 and have been verified to have no intellectual property or license issues before being released to open source by Dangdang in 2016. Dangdang will provide SGA and all committers will sign ICLA after Sharding-Sphere is accepted into the Incubator.

External Dependencies

As all dependencies are managed using Apache Maven, none of the external libraries need to be packaged in a source distribution. All dependencies have Apache compatible licenses except MySQL (GPL-2.0).

We will remove MySQL dependencies in future. MySQL JDBC driver is adopted by Sharding-Proxy to connect MySQL now; We will use SPI to load JDBC driver, so MySQL JDBC driver is no longer provided on Sharding-Sphere.

- Guava Apache-2.0
- Guava Retrying Apache-2.0
- commons-codec Apache-2.0
- commons-pool Apache-2.0
- commons-dbcp Apache-2.0
- netty Apache-2.0
- curator Apache-2.0
- grpc Apache-2.0
- protobuf BSD 3-clause
- lombok MIT
- groovy Apache-2.0
- snakeyaml Apache-2.0
- spring-context-support Apache-2.0
- spring-context-test Apache-2.0
- spring-boot-starter Apache-2.0
- spring-boot-configuration-processor Apache-2.0
- spring-boot-starter-test Apache-2.0
- slf4j MIT
- logback EPL-1.0
- junit EPL-1.0
- hamcrest BSD 3-clause
- mockito MIT
- h2 MPL-2.0/EPL-1.0
- mysql GPL-2.0 Will remove before first apache release, use SPI instead of it
- postgresql BSD
- mssql-jdbc MIT
- HikariCP Apache-2.0
- ANTLR BSD
- [OpenTracing](#) BSD

Required Resources

Git Repositories

- <https://github.com/sharding-sphere/sharding-sphere.git>
- <https://github.com/sharding-sphere/sharding-sphere-example.git>
- <https://github.com/sharding-sphere/sharding-sphere-doc.git>

Issue Tracking

The community would like to continue using [GitHub](#) Issues.

Continuous Integration tool

Travis

Mailing Lists

- Sharding-Sphere-dev: for development discussions
- Sharding-Sphere-private: for PPMC discussions
- Sharding-Sphere-notifications: for users notifications, and notifications from [GitHub](#)

Initial Committers

- , Liang Zhang, zhangliang@apache.org
- , Hao Cao,
- , Sheng Wu, wusheng@apache.org
- , Hongtao Gao, hanahmily@apache.org
- , Yonglun Zhang
- , Juan Pan
- , Jun Zhao
- , Ling Yue
- , Xiaoguang Ma
- , [QingYang](#) Chen
- , Yongsheng Peng, pengys@apache.org

Affiliations

- JD: Liang Zhang, Yonglun Zhang, Juan Pan, Jun Zhao
- Dangdang: Hao Cao, Ling Yue
- CHINA TELECOM Bestpay: [QingYang](#) Chen
- Individuals: Sheng Wu, Hongtao Gao, Xiaoguang Ma

Sponsors

Champion

- Roman Shaposhnik (rvs at apache dot org)

Mentors

- Craig L Russell (clr at apache dot org)
- Benjamin Hindman (benh at apache dot org)
- Willem Ning Jiang (ningjiang at apache dot org)
- Von Gosling (vongosling at apache dot org)

Sponsoring Entity

We are expecting the Apache Incubator could sponsor this project.