Mindtree

Cloud-enabled multi-master HA through Galera & MySQL Reference Architecture

Abstract: This paper is for enterprise solution architects, Database & Big Data (Infrastructure) architects, and Hadoop Administrators who manage infrastructure operations on Big Data platforms. We are going to discuss a multi-master, cloud-ready HA solution using Galera and MySQL for which we created a reference architecture and implemented it for multiple customers.

Introduction

Why Clustering?

It's required if the database is running on a single server/node that becomes the single point of failure. However, we can recover the data that has been lost during the failure in a number of ways, but we cannot recover the full data. Moreover, for the end user, it results in latency.

Clustering in databases refers to the ability of several database instances to work in such a way that for the end user, they will appear as only one instance. In case, for any reason, one instance goes down, another can take up all the processes and the end-user will observe neither latency nor data loss. Clustering could help improve I/O throughput for your database service and share CPU & memory loads.

Fault tolerance is another major pros of clustering. As there are more than one instances running, clustering offers an alternative instance in the case of individual instance failure.

Finally, Load Balancing is another feature of clustering that distributes the load among instances and directs users automatically to the instance with the least load.

Why Galera Cluster?

Galera Cluster is based on InnoDB SE. The main advantage of using Galera Cluster over regular MySQL replication is that Galera is an open source synchronous replication. Galera addresses a number of issues pertaining to MySQL replication such as write conflicts and replication lags.

MySQL Galera is a shared-nothing architecture as it uses certification-based replication (Note 1) to achieve synchronous replication. Galera also supports auto refresh, auto discovery, and auto sync features.

Problem Definition

Below are the few of the major pitfalls MySQL/DBRD replication process:

- Asynchronous replication will result in data loss.
- MySQL 5.5 changes back to asynchronous state when in trouble.
- Single-threaded replication.
- Master-slave architecture: read-write replication lag will result in inconsistent data.
- Cold standby during disaster situation.

High-Level Solution

In order to address the limitations of the cluster products mentioned above, we tried out Galera (Note 2) cluster in our lab. After setting up the reference architecture, the same solution was proposed to a number of customers which were implemented accordingly.

Solution Details

A number of customers, who were using database products like Oracle, MS SQL etc., started showing inclination to MySQL due to multiple reasons. Oracle, MS

1.1 Reference Architecture

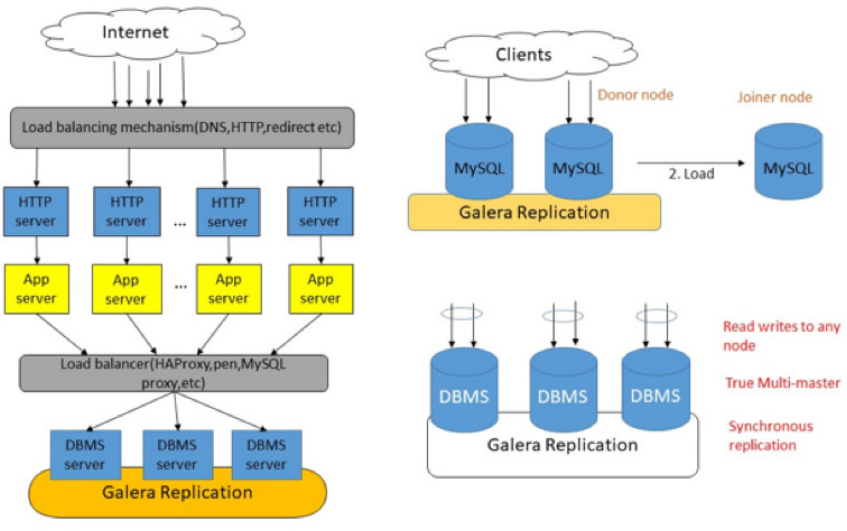

SQL etc. support the concepts of database clustering. So, in order to provide MySQL-based solution as per customer's interest, we had to offer the cluster service in MySQL as well and that is when we began evaluating Galera.

After going through the successful demonstration of this HA product and the architecture that we set up in lab, the customers are keen to move from other DB cluster solution to this. In addition, we saw that the customers are interested in adding more nodes to the Galera cluster.

MySQL-Galera Architecture:

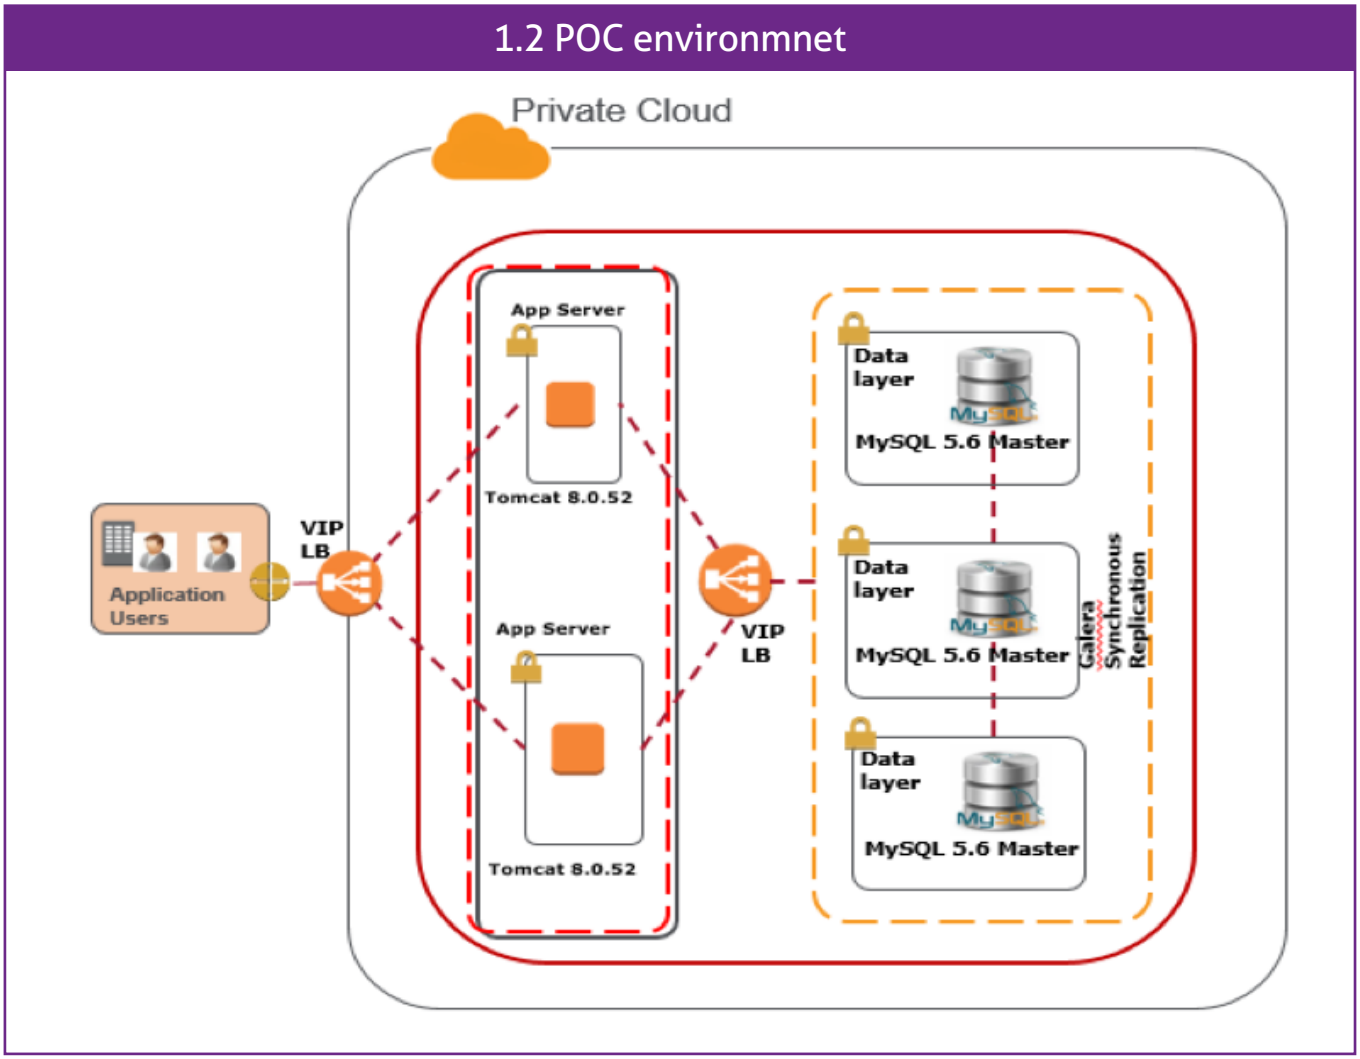

The above POC environment was designed using below components,

- Private Cloud
- OS - Centos 7.0
- Middleware - Apache Tomcat 8.0.52
- Database - MySQL 5.6 Master
- Load Balancer - HAProxy – Distributing Load

Using the tool "MySQL Workbench", we demonstrated

- High Availability
- No slave lag
- No lost transactions
- No more data inconsistency
- Smaller client latencies
- Read scalability and write throughput improvement.

Galera Success Stories

One of our customers, a big player in telecom vertical, has multiple applications running based on Drupal. For managing the databases for that application, Mindtree has implemented quadruple master Galera Cluster for MySQL.

Mindtree Capability in Galera-space (on Prem + on Cloud)

- Install & configure
- Tuning
- Troubleshooting
- Day-to-day management
- Automation
- Dockerization

What is galera cluster ?

MySQL Galera cluster is a multi-master that is on top of synchronous replication. In Galera, the certification-based replication is used to achieve synchronous replication along with transactions ordering techniques and group communication.

Summary

Galera Cluster is one of the most robust and all-encompassing solutions in cluster era. Though it's slightly complicated to set up for the first time, once deployed, it provides significant improvement in high-availability, transparency and scalability for the MySQL ecosystem.

References

- Note 1 - http://galeracluster.com/documentation-webpages/gettingstarted.html
- Note 2 - http://galeracluster.com/products/
- https://www.digitalocean.com/community/tutorials/how-to-configure-a-galera-cluster-with-mysql-5-6-on-ubuntu-16-04
- https://severalnines.com/blog/9-tips-going-production-galera-cluster-mysql
- https://www.sebastien-han.fr/blog/2012/04/01/mysql-multi-master-replication-with-galera/
- http://www.oracle.com/us/corporate/pricing/technology-price-list-070617.pdf