

Providing high-availability for object-relational databases in distributed data centers

Samokhin Nikita Y. (ITMO University, Saint Petersburg)

Vepritskaya Victoria M. (ITMO University, Saint Petersburg)

High availability (HA) is a feature of a system, which aims to ensure the required level of operational performance, usually uptime, for a period higher than normal. HA for databases includes three main properties as follows:

- Replication - sharing information to ensure consistency between redundant resources;
- Failover - switching to a redundant or standby unit upon a failure or an abnormal termination of a unit in action;
- Load balancing - distribution of workloads across multiple resources to optimize resource use, maximize throughput, minimize response time, and avoid overload of any single resource.

The main goal of this work is to suggest an appropriate approach to provide HA for (geographically) distributed databases. The amount of data centers (DC) is growing daily, especially in Russia. A geographically distributed DC is a DC that has more than one physical components separated and located far from each other. A database used in such DCs is considered as a distributed database. However, it could be distributed within a single DC (or parts of it).

Database HA is provided by means of replication. Replication can be used on many database management systems (DBMS) with a master-slave or master-master relationship between the original and the copies. It is the multi-master replication scheme that the current research is focused on. In multi-master replication, updates can be submitted to any database node, and then ripple through to other servers. The most common challenge in a multi-master replication is transactional conflict prevention or resolution. Ideally, when data is replicated between database servers, so that the information remains consistent throughout the database system and users cannot tell or even know which server in the DBMS they are using, the system is said to exhibit replication transparency and the database is considered to have HA.

The object-relational database management system (ORDBMS) PostgreSQL 11 is used as an example. Pacemaker is used to provide failover and load-balancing. Bucardo 5.5 software is used for performing the asynchronous replication in multi-master scheme. The whole system is operating within the Naulinix operating environment.

The study was carried out within the ITMO University project of developing a geographically distributed cloud system. The study has been financially supported by the Ministry of Education and Science of the Russian Federation.

Автор: _____

Самохин Н.Ю.

Научный руководитель: _____

Вефрицкая В.М.

Профессор ЦИИЯ ИТМО: _____

Рябухина Ю.В.