High Availability of Services

Business critical systems require maximum accessibility of the services offered. We therefore deploy our systems in central clusters of the application and database servers. Clustered solutions increases performance of the system by sharing the traffic between individual servers and simultaneously ensures better accessibility of the system even in the case of breakdown of individual applications or database servers. Moreover, permanent monitoring of the server status needs to be ensured so as to remove any problem which arises, or which is expected to arise, as soon as possible.

For the most critical systems with a large number of connected clients we use independent and geographically separated data centres and employ the mechanism of distributed replication of databases. In the case of breakdown or unavailability of the whole primary data centre, the system requests are redirected to a backup data centre, and thus, it might well happen that the client does not notice the breakdown of the data centre at all.

Database and High Availability of the System

Building highly available applications with high data volumes is a big challenge. The application must be resistant to failure of any component of the system. One of the most critical elements in this case is the database. This is because it holds all the data necessary for the application and if the data is not available the functionality of the entire system is broken.

The failure of the database or its components can have fatal consequences on the availability of the system. A critical failure is especially corruption of data. Of course, each database will have its backup / restore plan so the data is not lost forever. However, the restoration of large volumes of the data will definitely take a long time; longer than is acceptable for a high availability system.

Even the short blackouts of the system caused by the failure of the database server are unpleasant. Failures of the database server can occur quite frequently and the system should be able to handle such failures without impacting on end users.

Our company develops and supplies high availability systems to its customers. In order to conform to high availability, the databases of its systems are based on the high availability features of the Oracle Database.

The possible failure of the database system can be caused, for example, by failure of any component of the database hardware (server, storage, network etc), corruption of the stored data files, human error, etc. The Oracle software has the answer to many of those challenging situations. Here we discuss some of them.

Failure of the Database Server

The oracle Real Application Cluster feature allows running of the database on more than one server node. If one node fails, the remaining nodes take over the processing of the requests. No manual intervention is needed and no downtime is required (although there is some very short time when the system becomes unresponsive because of the automatic reconfiguration).

Corruption of the Stored Data

The databases of our systems contain large volumes of the data (usually terabytes). In case of corruption of data or complete failure of storage, the recovery of such an amount of data would take hours. We split our data into partitions according to relevance. We use the Partitioning feature of the database. Each partition of the data is stored in a dedicated file. In the case that corruption occurs, we restore the files with the most relevant data first. Such restoration takes a few minutes. After that restoration, our systems are able to work and offer all functions which are required to be highly available. The recovery of the less relevant data is performed in parallel with the running system. The reports over the less relevant data (for example older data) become available later, when the reported data is recovered. What is important is that the system offers important functionality within a short time after the incident.

Production System Crash

Crashes of the complete systems do not happen very often. However, high availability systems should also consider such a scenario. In the case that the crash takes place at the site where the system resists, there is no way to recover the system quickly at that site. The only solution in this case is to have a secondary site containing the same system and to switch the production operations into that site. The challenge in such a scenario is synchronization of the data stored in the database between the primary (production) site and the secondary (standby) site. During normal operations, the data at the primary site is changed continuously. There is a need to transport those changes to the secondary site. The solution is the DataGuard database feature. The DataGuard feature ships change-vectors from the primary site and then applies them on the secondary site. Based on the configuration this feature guarantees minimal or even zero data loss in case of a crash of the primary site. This feature also offers an alternative to the recovery of the system in case of the corrupted data files, as the corruptions are not transferred.

Human Error

The DataGuard feature helps us also in the case of human error. Application of the change vectors is delayed at the secondary site. In the case of human error, the application is stopped prior to applying the error at the secondary site. The production in such a case can be redirected into the secondary site with minimum data loss. Our systems use the above discussed technologies to achieve high availability of database systems (and many others). We protect our system by the DataGuard feature, which as more than one secondary site, some of which are located at the primary site location and provide an alternative to recovery of corrupted data files, some of them serve as the fix for the possible human error. There are also geographically distributed secondary sites which are prepared to take over the work in the case of a primary site disaster. We use the Real Application Cluster feature to minimize downtime in the case of server failures. We use partitioning to improve manageability and reduce the downtime necessary for maintenance (one partition is processed at a time; the data for all other partitions is accessible).

Conclusion

The high availability of the systems is a challenge which must be addressed especially at the database level. There are various technologies helping us to address that challenge. We have decided to base the high availability of our databases on Oracle technologies. The features such as Oracle Real Application Clusters, Oracle Data Guard and Oracle Partitioning allow us to implement high availability for our database.