

Author: Bill Ennis

Oracle Standby is a popular failover option used. It features the use of the transaction logs by transferring them from the Production server to the standby and then applying those transaction logs so that the Standby server is kept in sync.

Things to keep in mind before getting started:

In order to use Standby there are a couple of pre-requisites to keep in mind. The first is that the database must be in archive log mode. If there are no archive logs there will be nothing to transfer over to the Standby Server. Another very important application restriction is that you cannot use non-logged operations. If operations are not logged they will not be present in the transaction logs. If they are not in the transaction logs they will never be sent to and applied in the Standby server. FYI - Non-logged operations are sometimes used when using Oracle's SQL Loader utility to load data into the database.

Tranferring the Archive Logs over to the Standby Server

Oracle 8I makes the transfer of the archive logs to the Standby machine seamless and transparent. You can actually configure Oracle to copy the archive log files over the network to the standby machine. This is called running in Duplex archive log mode. Oracle will also retry the transfer for archive logs in the event of a periodic network outage. You can configure Oracle to write in both locations (Production and Standby) as mandatory meaning that both writes are done successfully or none - but in most cases this is not recommended.

Deciding on a Window of Recoverability

A very important decision you must make when using Oracle Standby is how often you intend to perform log switches. By default Oracle will switch the active transaction log (called a redo log) when it is full. Redo log sizes are configurable. In addition the amount of activity going on within the system usually varies. So if you don't explicitly perform log switches you never really know when they are going to happen. For this reason in all cases that I have used Oracle Standby I have run a regularly scheduled job on the Production server that switches the redo log for me (usually at 5 minutes). This interval is key, for it determines the worst case scenario of the amount of data you may be willing to lose. The worst case is the event that you totally lose the capability to get to the archive logs in the event of a disaster and must activate the Standby server without the last archive log.

Operational Issues

The person that is responsible for maintaining the operational aspects of your database needs to be made aware that you are using Oracle Standby. This is because there are certain situations where changes made to the Production instance are not completely carried out on the Standby Server. This mainly involves when you make structural changes to the database that require operating system actions such as the creation of new datafiles or the dropped use of datafiles.

Monitoring

When I have a Standby Server in operation I want something to keep an eye on things to make sure they are going as planned. For this reason I have created an application that monitors the status of the Standby server and keeps track of the last archive log that was applied. This application then queries the production instance for the last archive log that it has created. If things match up all is well; if not this application will fire off a notification (usually to an email enabled pager or cell phone) to indicate that there is a problem.