



Role of Oracle Active Data Guard in High Availability Database Operations

Rab Nawaz
Department of Information
Technology, SZABIST,
Dubai, UAE

Tariq Rahim Soomro
College of Engineering & IT
Al Ain University of Science &
Technology, Al Ain, UAE

ABSTRACT

Data Guard is one of the high availability, disaster recovery and data protection architecture from Oracle. The Active Data Guard in Oracle Database 11g has brought significant changes in the existing architecture. This paper will study these changes and will see how active Data Guard is providing high availability with offloading workload to one or more standby databases and providing the opportunity to query real time data in standby databases. Finally this study will see how it provides the way to perform different write operations on primary database through standby databases.

Keywords

Oracle Active Data Guard, Data Guard, Stand by database, Fail over, switch over, Primary or secondary database, High availability in databases, data protection, Return on investment.

1. INTRODUCTION

Active Data Guard has been introduced by Oracle since its database version 11G R1 [1], before defining what exactly is

Active Data guard first we need to understand redo, principle of Data Guard, primary and standby databases, but it is no doubt established that what active data guard can do for Oracle Database? no other data protection technology can do for it. It provides reliability, performance, unique levels of simplicity and security of having an independent, exact replica of the production database. The principle of data guard has to ship (transfer) redo, and then apply redo, redo consists all the changes made in the database in the form of the log files that are required by oracle database to recover a database transaction [1]. Figure-1 below shows the principle of Data Guard.

Here in figure-1, 'a' generates the redo on primary database it's shipped (transfer) to secondary database with the redo transport services and 'b' here applies services update the standby after validating the redo. Here 'c' (primary database) writes process and writes the changes to data files, which are independent to Data Guard and here 'd' means if there is any outage due to network failure or some other reasons, the Data Guard will automatically resynchronize (resynch) the standby using archived redo generated at primary database.

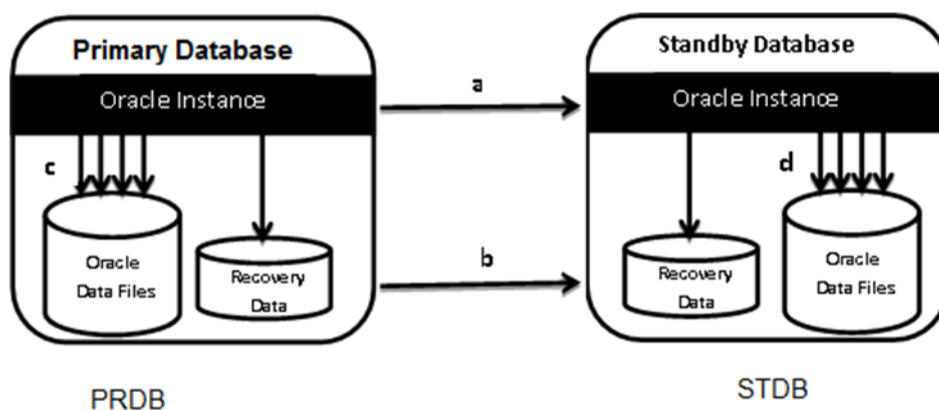


Fig 1: Principles of Data Guard



The primary database [2] is a production database referred to as primary, which is transmitting the redo to one or more replicas known as stand by databases, or we can say it's a database instance whose availability is required. The standby database is receiving and applying the redoes from primary database, it's continuously in recovery state, validating the redo and synchronizing the standby database with the primary database. Active Data Guard in Oracle Database 11g R2 supports up to 30 standby databases. There are three possible standby databases types [2].

1.1 Physical standby database

Physical standby database is an identical copy of primary database with same structure on disk on block by block bases. Physical standby database takes the redo data from primary database and applies to keep it "synched" with Primary.

1.2 Logical standby database

Logical standby database it contains production database as logical information, and the structure of data here can be different than the physical organization. Here the standby database logically kept synchronized using SQL Apply, which used to transforms the data from the primary database into SQL statements and latter on executes the SQL statements at the end of standby database. Here some restrictions might apply.

1.3 Snapshot standby database

Snapshot standby is consists of standby database, which is based on fully read write capabilities. It is created in such a way that physical standby database converted into a Read write snapshot standby database and capable of archiving and receiving redo data from primary database. The redo data is not allowed to applied until unless the conversion of snapshot standby into physical standby database. Here it is important that this should first discard local update that were made to the snapshot standby database. Also note that a snapshot standby database does not apply the redo data that it receives.

This paper will first discuss in section 2 what is active data guard? Section 3 will explore Offload and Workload of Active Data Guard on Primary Database. Section 4 will discuss the performance of Active data guard; and finally conclusion and future work will be discussed.

2. WHAT IS ACTIVE DATA GUARD?

According to [3] an Oracle Active Data Guard is a solution, which enables read-only access of a physical standby database while continuously applying changes received from the production database for the purposes of querying, reporting, sorting, and Web-based access etc. When we use a for read-only purpose the physical standby database, while redo data is still being applied to the standby database, this situation is known as real-time query or "the Oracle Active Data Guard" option. It provides disaster recovery and data protection and allows users to access an up-to-date physical standby database [4]. Supplementary writes are also possible but with the condition that these writes must be redirected to a read write database with the help of database links. Configuration of active data guard is also important here. Configuration of Active Data Guard is not the scope of this paper, we have assumed that we already have two separate machines one is running production database that will be referred as PRDB and the other one is standby which will be pronounced as STDB, as we have shown in figure 1 above. Now question arises that how to achieve high availability with "Active Data Guard"? This is possible with two operations known as

"Switchover" and "Failover". A switch over operation is responsible to take care of transition role, either transition take over the primary and/or one of its standby databases. A switchover is done typically for planned maintenance of the primary system and guarantees no data loss. During a switchover the role of the primary database take place as a standby role, and the role of standby database take place as the primary role. With only Active Data guard we can achieve this, and this could be required like in case of patching or DB upgrades as an example. A failover operation occur, when the primary database become unreachable or fails for whatever the reasons and one of the standby databases take over the primary role and when primary database cannot be recovered in timely manner for whatever the reason then failover should be performed. In the time of the failover, it may or may not result in data loss depending on the protection mode [5].

3. ACTIVE DATA GUARD OFFLOADS / WORKLOAD ON PRIMARY DATABASE

Following are the few situations, through which we can reduce the burden from Primary databases:

3.1 Standby database as read only with real time data

Consider the situation when there was data guard, before introduction of active data guard. Data guard used to do one activity at a time, that means physical standby servers could be used either for read only or for applying redo logs. In that case for the purpose of querying the database standby database was being used only for applying redo logs from production and it was before totally waste of standby resources. Now after introduction of active data guard, one can at the same time open the database for read only purpose and redo logs may applied simultaneously, so it is possible to query real time data with Active data Guard and it is providing efficient use of physical server. This means the investment on standby server can get the return and providing no doubt offloading of burden on production server [1].

3.2 Backup from Standby Database

Here we can backup from physical standby Server rather than taking backup from production since physical standby Server is 'sync' with the production as the real time redo are being applied. Active data guard is helping hand here, where it helped to offload backup work from production database [6].

3.3 Operations on Primary database through Standby Database

There are several operations on primary database through standby database. For example Active Data Guard Standby Database writing data to primary database. For 'write' operation there will be use of production or primary database link in standby database. The operations will be redirected to Primary Database only through database links [7]. Following are few examples:

3.3.1 Example of Redirecting DML Query:

Once the Data Definition Language (DDL) creates the 'dblink' then is applied on the standby database. So one can use it to redirect writes to a remote database. For example [7][9]:



```
SQL> insert into emp@PRDB values (999,'WARUN','SUPER
GEEK',888,sysdate,1,0);
```

row created.

```
SQL> commit;
```

3.3.2 Examples of Redirecting DDL Queries:

Data Definition Language (DDL) statements cannot be executed directly over a 'dblink'. If it is necessary to redirect a DDL statement from the standby to a remote database, then one can call a remote procedure that uses dynamic SQL to perform remote DDL. However, one cannot invoke a remote procedure (even a read-only remote procedure) from a read-only database but can put the remote procedure call in a stored procedure. For example:

On the primary database:

```
CREATE OR REPLACE PROCEDURE do_ddl(STRING IN
varchar2) AS
BEGIN
execute immediate string;
END;
CREATE OR REPLACE PROCEDURE
call_do_ddl(STRING IN varchar2) AS
begin
do_ddl@PRDB(string);
end;
```

On the standby database:

```
SQL> exec call_do_ddl('create table foo2 (col1 number)');
```

PL/SQL procedure successfully completed.

```
SQL> select * from foo2@PRDB;
```

no rows selected

This confirms that the table has been created – otherwise the same query would have returned an ORA-942 error "table or view does not exist".

3.3.3 Example of Sequences:

If the application connected to the Active Data Guard standby and needs to use a sequence, then the application must also be redirected to the remote database. For example:

```
SQL> select customers_seq.nextval from dual@PRDB;
```

NEXTVAL

 1003

3.3.4 Operation of Data pump:

One cannot export directly from a read-only database because Oracle Data Pump needs to write to its master table. So it is obvious to say that Oracle data pump is directly exporting from active data guard standby database.

3.3.5 Operation of Network Export/Import:

To extract data from the standby one must run Data Pump from a read-write database using the NETWORK_LINK qualifier in the export or import command. It is also possible to export data from an Active Data Guard standby database and import it directly into read-write target database. This eliminates the need for the export dump file but the directory is still required for the log file of the export/import operation.

3.4 Reader Farm (situation where multiple standby database exits and all are

being queried concurrently to improve the reading performance of the application)

Reader farm is used in a situation where multiple standby database exits and all are being queried concurrently to improve the reading performance of the application. Here a Data Guard configuration can support multiple standby databases, and the basic configuration is shown in Figure-1 above and can be extended to boost the reader performance of the most demanding applications and beyond that what primary database alone could support. Figure-2 below shows an Active Data Guard configuration with a single primary database that supports read/write transactions; and multiple standby databases that provide read-only access to users [8].

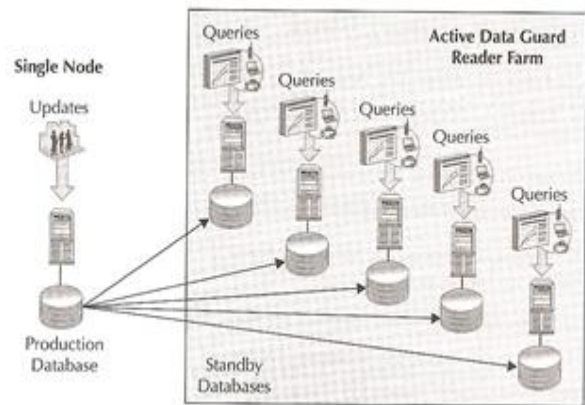


Fig 2: Reader Farm

4. PERFORMANCE OF ACTIVE DATA GUARD STANDBY DATABASE

Performance of active data guard with the help of standby database can be attained in following ways:

4.1 Gap between primary and active Data Guard Standby database (Apply Lags).

When one use from a primary database to a physical standby database a real-time query to offload queries, then to ensure that it is within acceptable limits one has to monitor apply lag. The current apply lag may be different, in the form of elapsed time, between "the same change was first visible on the primary" and/or "the last applied change became visible on the standby". This metric is almost computed to the nearest second. To obtain the apply lag, query is performed like this: V\$DATAGUARD_STATS view [9].

4.2 Use of "Statspack" for standby using primary database.

Normally the Automatic Workload Repository (AWR) would be used to gather information about the performance on a database but as AWR writes its information back to the database it cannot be used on an Active Data Guard standby, which is because it is read-only. In order to gather performance information on an Active Data Guard standby one need to configure "Standby Statspack". With Active Data Guard, one can use "statspack" from the primary database to collect data from a standby database that is read-only and performing recovery and queries. Like AWR, "Statspack" cannot be executed directly on a physical standby due to its read-only nature. Therefore tuning the performance



of the apply process and SQL queries involves manually collecting statistics [9].

5. CONCLUSION

The active data guard is no doubt an important answer to the issues discuss earlier. Active Data Guard not only provides a simple solution to improve the performance of high-availability database applications by providing backups to a synchronized replica of the production database, offloading the overhead of ad hoc queries and reporting. The simplicity and performance of Active Data Guard makes it possible to be deployed for a wide variety of business applications. Active Data Guard provides a high return on investment while simultaneously protecting against data loss and downtime.

6. REFERENCES

- [1] Larry Carpenter, Joe Meeks, Charles Kim, Bill Burke, 2009, Oracle Data Guard 11g Handbook, 978-0-07-162148-9
- [2] Scott Jesse, Bill Burton, Bryan Vongray, 2011, Oracle Database 11g Release 2 High Availability, 978-0-07-175207-7
- [3] Steven Chan, 2008, Comparing Oracle Data Guard Vs. Active Data Guard for EBS Environment, https://blogs.oracle.com/stevenChan/entry/comparing_oracle_data_guard_vs_active_data_guard_f
- [4] Virginia Beecher, Viv Schupmann, Janet Stern, 2011, Oracle Database High Availability Overview, 11g Release 2 (11.2), E17157-08
- [5] Oracle Data Guard Broker , From Oracle Documentation, http://docs.oracle.com/cd/B28359_01/server.111/b28295/sofo.htm
- [6] Oracle-Base (www.Oracle-base.com) <http://www.oracle-base.com/articles/11g/data-guard-setup-11gr2.php>, Retrieved January 2013
Forman, G. 2003. An extensive empirical study of feature selection metrics for text classification. *J. Mach. Learn. Res.* 3 (Mar. 2003), 1289-1305.
- [7] Oracle Maximum Availability Architecture (White Paper September 2011), Michael T. Smith http://www.google.ae/url?sa=t&trct=j&q=write%20operations%20with%20oracle%20active%20data%20guard&source=web&cd=1&ved=0CCsQFjAA&url=http%3A%2F%2Fwww.oracle.com%2Ftechnetwork%2Fdatabase%2Ffeatures%2Favailability%2Fmaa-wp-11gr1-activedataguard-1-128199.pdf&ei=W6XuUPCdCOvs0gWMtIDwDg&usq=AFQjCNHwumyrbQzD5xJf5_kD7HUqGvjhQ&bvm=bv.1357700187,d.d2k
- [8] Data Disk (www.datadisk.co.uk), http://www.datadisk.co.uk/html_docs/oracle_dg/active_dg.htm, Retrieved January 2013
- [9] Larry M. Carpenter , October 2011, Oracle OpenWorld Active Data Guard Hands on Lab