IJCSMS International Journal of Computer Science and Management Studies, Vol. 12, Issue 02, April 2012
ISSN (Online): 2231-5268
www.ijcsms.com

Disaster Recovery of Data by Using Data Guard

Anil[1], Sanjay Kumar[2], Satish Kumar[3], Parveen Yadav[4], Balraj Sharma[5]

[1]Student M.Tech C.S.E, BITS Bhiwani, Haryana, India
anilverma.bits@gmail.com

[2]Asst. Professer, Vaish College of Engineering Rohtak, India

[3]Asst. Professer, Vaish College of Engineering Rohtak, India

[4]Assistant professor, BITS Bhiwani, Haryana, India
par.biran@gmail.com

[5]Lecturer, BITS Bhiwani, Haryana, India
par.biran@gmail.com

Abstract

Oracle Data Guard is the management, monitoring, and automation software infrastructure that creates, maintains, and monitors one or more standby databases to protect enterprise data from failures, disasters, errors, and data corruptions.

Data Guard maintains standby databases as consistent copies of the production database as far as transactions are concerned. These standby databases can be located at remote disaster recovery sites thousands of miles away from the production data center, or they may be located in the same city, same campus, or even in the same building. If the production database becomes unavailable because of a planned or an unplanned outage, Data Guard can switch any standby database to the production role, thus minimizing the downtime associated with the outage, and preventing any data loss.

The document explains the structure of a physical standby database with Oracle Data Guard in an SAP environment. It indicates all the steps needed to successfully install and configure an Oracle Data Guard system with a physical standby database and the logical order in which they must be carried out.
To enable you to operate the standby database (Oracle Data Guard), a description of how to configure the Data Guard Broker is also provided. In just a few steps this service allows you to swap the database roles. This means that in the event of a disaster, what is known as a switchover or failover is undertaken almost automatically. The database administrator can initiate the process with just one command.

Keywords: *Replication, Disaster recovery, backup, recovery, hot backups, online backups, dynamic changes, reconciliation, Disaster handling.*

1. Introduction

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruptions. Data Guard maintains these standby databases as transitionally consistent copies of the production database. Then, if the production database becomes unavailable because of a planned or an unplanned out-age, Data Guard can switch any standby database to the production role, thus minimizing the downtime associated with the outage. Data Guard can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability..

1.1 Data Guard Configurations

A Data Guard configuration consists of one production database and up to nine standby databases. The databases in a Data Guard configuration are connected by Oracle Net and may be dispersed geographically. There are no restrictions on where the databases are located, provided that they can communicate with each other. For example, you can have a standby database on the same system as the production database, along with two standby databases on another system.

You can manage primary and standby databases using the command-line interface or the Data Guard broker, who includes a graphical user interface called Oracle Data Guard Manager.

1.1.1 Primary Database

A Data Guard configuration contains one production database, also referred to as the primary database that functions in the primary role. This is the database that is accessed by most of your applications.

The primary database can be either a single-instance Oracle database or an Oracle Real Application Clusters database.

1.1.2 Standby Databases

A standby database is a transitionally consistent copy of the primary database. A standby database is initially created from a backup copy of the primary database. Once created, Data Guard automatically maintains the standby database by transmitting primary database redo data to the standby system and then applying the redo logs to the standby database.

Similar to a primary database, a standby database can be either a single-instance Oracle database or an Oracle Real Application Clusters database.

A standby database can be either a physical standby database or a logical standby database:

- *Physical standby database*

Provides a physically identical copy of the primary database, with on-disk database structures that are identical to the primary database on a block-for-block basis. The database schema, including indexes, is the same. A physical standby database is kept synchronized with the primary database by recovering the redo data received from the primary database.

- *Logical standby database*

Contains the same logical information as the production database, although the physical organization and structure of the data can be different. It is kept synchronized with the primary database by transforming the data in the redo logs received from the primary database into SQL statements and then executing the SQL statements on the standby database. A logical standby database can be used for other business purposes in addition to disaster recovery requirements. This allows users to access a logical standby database for queries and reporting purposes at any time. Thus, a logical standby database can be used concurrently for data protection and reporting.

The primary and secondary server instances send their own history and status to the monitor server instance.

1.1.3 Configuration Example

Instance Figure 1-1 shows a Data Guard configuration that contains a primary database instance that transmits redo data to physical and logical standby databases that are both in remote locations from the primary database instance. In this configuration, a physical standby database is configured for disaster recovery and backup operations, and a logical standby database is configured primarily for reporting, but it can also be used for disaster recovery. You could configure either standby database at the same location as the primary database. However, for disaster recovery, Oracle Corporation recommends that you configure standby databases at remote locations.

Figure 1-1 shows a typical Data Guard configuration in which archived redo logs are being applied to both physical and logical standby databases.

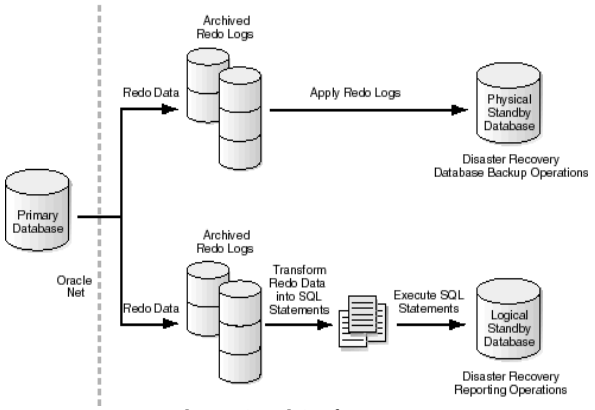

Figure 1-1 Typical Data Guard Configuration

2. Data Guard process architecture

As shown in the following figure, Data Guard uses several processes of the Oracle database instance to achieve the automation necessary for disaster recovery and high availability.

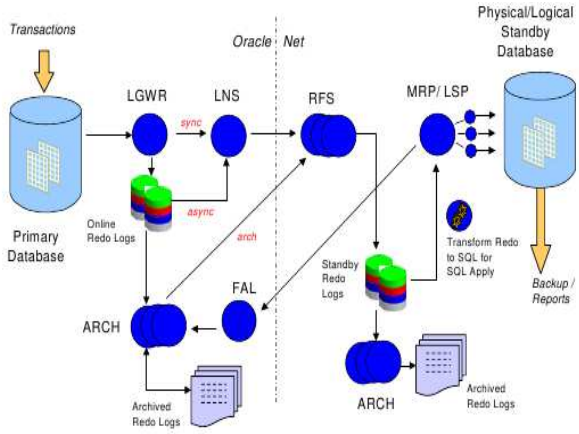

Figure 1-2 shows a typical Data Guard configuration

On the primary database, Data Guard uses the Log Writer (LGWR) process or multiple Archive (ARCH) processes to collect transaction redo data. In order to ensure isolation from network disruptions, the Log Writer process uses specialized background processes, called Log Writer Network Server (LNS) process, to synchronously or asynchronously transmit the redo data to the standby

database. The Archive processes transmit the redo data to the standby database directly. The primary database also has the Fetch Archive Log (FAL) process to provide a client-server mechanism for transmitting archived logs to the standby database following a loss of communication between the primary and standby database(s), for automatic gap resolution and resynchronization.

On the standby database, Data Guard uses one or more Remote File Server (RFS) processes to receive redo data from the primary database, the Managed Recovery Process (MRP) to apply redo data to the physical standby database, and the Logical Standby Process (LSP) to apply SQL-translated redo data to the logical standby database.
If the Data Guard Broker is enabled, Data Guard also uses the Data Guard Broker Monitor (DMON) process
to manage and monitor the primary and standby databases as a unified configuration.

3. Requirements

A second database host, configured in exactly the same way as the first, is needed to operate Oracle Data
Guard. This means:
•Same operating system, e.g. AIX 5.3 ML6 on both hosts
•Same parameter settings on database and operating system (e.g. nfiles)
•Identical Oracle version, currently the 11.2.0.1 database patch set with individual patches as recommended in SAP note.
•Identical file system structure, especially for SAP data and Oracle home.
The databases must be operated in ARCHIVELOG mode
•Use of server parameter files (SPFILE)

4. Benefits of Data Guard

- *Disaster recovery and high availability* - Data Guard provides an efficient and comprehensive disaster recovery and high availability solution. Automatic failover and easy-to-manage switchover capabilities allow quick role transitions between primary and standby databases, minimizing the downtime of the primary database for planned and unplanned outages.
- *Complete data protection* - A standby database also provides an effective safeguard against data corruptions and user errors. Storage level physical corruptions on the primary database do not spread to the standby database. Similarly, logical corruptions or user errors that cause the primary database to be permanently damaged can be resolved. Finally, the redo data is validated at the time it is received at the standby database and also when applied to the standby database.
- *Efficient utilization of system resources* - A physical standby database can be used for backups and read-only reporting, thereby reducing the primary database workload and saving valuable CPU and I/O cycles. In Oracle Database 10g Release 2, a physical standby database can also be easily converted back and forth between being a physical standby database and an open read/write database. A logical standby database allows its tables to be simultaneously available for read-only access while they are updated from the primary database. A logical standby database also allows users to perform data manipulation operations on tables that are not updated from the primary database. Finally, additional indexes and materialized views can be created in the logical standby database for better reporting performance.
- *Flexibility in data protection to balance availability against performance requirements* - Data Guard offers the Maximum Protection, Maximum Availability and Maximum Performance modes to help enterprises balance data protection against system performance requirements.
- *Protection from communication failures* - If network connectivity is lost between the primary and one or more standby databases, redo data cannot be sent from the primary database to those standby databases affected. Once connectivity is reestablished, the missing redo data is automatically detected by Data Guard and the necessary archive logs are automatically transmitted to the standby databases. The standby databases are resynchronized with the primary database, with no manual intervention by the administrator.
- *Centralized and simple management* - Data Guard Broker automates management and monitoring tasks across the multiple databases in a Data Guard configuration. Administrators may use either Oracle Enterprise Manager or the Broker's own specialized command-line interface (DGMGRL) to take advantage of this integrated management framework.
- *Integrated with Oracle database* - Data Guard is available as an integrated feature of the Oracle Database (Enterprise Edition) at no extra cost.

5. Disadvantages of Data Guard

- *Broken network connection between the observer and the primary database.*

Completely unusable, If the connection is lost between the observer and the primary database, or there are network failures that cause the primary database to be isolated, the observer attempts a fast-start failover.

- *Instance failures.*
  If a single-instance primary database (either RAC or nonRAC), or if all instances of a RAC primary database fail, the observer attempts a fast-start failover.
- *Shutdown abort.*
  If a single-instance primary database (either RAC or nonRAC), or if all instances of a RAC primary database are shut down with the ABORT option, the observer attempts a fast-start failover. Fast-start failover will not be attempted for the other types of database shutdown.
- *Offline data files.*
  If the observer determines that one or more data files in the primary database have been taken offline by the database because of I/O errors, the observer attempts a fast-start failover.
- *Corrupted Dictionary.*
  Dictionary corruption of a critical database object. Currently, this state can be detected only when the database is open
- *Corrupted Control file.*
  Each database that needs to be log-shipped must be set up through a separate maintenance plan.

6. Conclusions

The objective of this work is to identify the possibilities in the technology and to write an algorithm to make the log shipping from production database at one location to the standby database at remote geographic location fastest.

The approach behind this research will be two pronged.

Compressing the redo logs to the optimum in advance. Doing the shipping of compressed logs parallel to the remote standby database location and uncompressing them to the original forms so that they can be applied on the standby database.

References

[1] Data Guard Broker
[2] Data Guard Concepts and Administration
[3] Oracle data base high Availability best practices 10g release 2 documents.
[4] Setup Flashback Database on Data Guard Physical Standby Database for SAP Customers.