White Paper

Deploying Oracle® Data Guard on Oracle Database 10*g* R2 Enterprise Edition Real Application Clusters with Red Hat® Enterprise Linux® 4 on Dell™ PowerEdge™ Servers and Dell|EMC™ Storage

Abstract

Oracle Data Guard is an effective and widely adopted data protection and disaster recovery solution for Oracle databases. One of the biggest challenges in Data Guard deployment is the instantiation of the standby database which is traditionally done by Oracle Recovery Manager (RMAN) backup and restore. However, this method can lead to significant downtime for large production databases. This article outlines a Data Guard deployment solution utilizing DELL|EMC data replication technology, which can significantly reduce primary database downtime during Data Guard deployment.

July, 2007

July 2007 Rev. A01

Table of Contents

Introduction

Dell™ PowerEdge™ servers and Dell|EMC™ storage systems are ideal choices to deploy highly reliable and sustainable Oracle 10*g* databases. This white paper is intended to help IT professionals design and configure Oracle 10*g* Data Guard solutions using Dell servers and storage that apply "best practices" derived from laboratory and real-world experiences. This white paper documents Dell's recommended approach for implementing a tested and validated solution for Oracle Data Guard on 10*g* R2 Real Application Cluster (RAC) database, Dell PowerEdge 9th generation servers, Dell|EMC storage systems and Red Hat Enterprise Linux 4.

Dell Solutions for Oracle Database 10*g*

Dell Solutions for Oracle Database 10*g* are designed to simplify operations, improve utilization and cost-effectively scale as your needs grow over time. In addition to providing price/performance leading server and storage hardware, Dell Solutions for Oracle Database 10*g* include:

- **Dell Configurations for Oracle** – in-depth testing of Oracle 10*g* configurations for the most in-demand solutions; documentation and tools that help simplify deployment
- **Integrated Solution Management** – standards-based management of Dell Solutions for Oracle 10*g* that lower operational costs through integrated hardware and software deployment, monitoring and update
- **Oracle Licensing** multiple licensing options that can simplify customer purchase
- **Dell Enterprise Support and Professional Services for Oracle** – offerings for the planning, deployment and maintenance of Dell Solutions for Oracle 10*g*

For more information concerning Dell Solutions for Oracle Database 10*g*, please visit www.dell.com/oracle.

Overview of this White Paper

The balance of this white paper will provide the reader with a detailed view of Oracle Data Guard on 10*g* R2 Real Application Cluster database, best practices for configuring the hardware and software components and pointers for obtaining more information.

Introduction

IT systems today have an increasing need to be highly reliable and available, in order to support global businesses. IT systems can not afford to be down which may cause huge financial loss. It is critical to implement high availability (HA) solutions to protect mission critical applications such as Oracle databases. To help maximize systems availability for Oracle databases, various best practices can be implemented using Dell PowerEdge server technologies, Dell|EMC CX3 storage array technologies, and Oracle 10g HA technologies.

As illustrated in Figure 1, system availability improves with additional redundancy features used in the Oracle database solution stack.

- **Dell PowerEdge servers and Dell|EMC CX3 storage** provide redundancy at the hardware level with technologies such as dual power supplies, dual host bus adapters, and dual communications interconnect.
- **Oracle Flashback technology** enables self-service repair operations to recover from logical corruptions while the database is online.
- **Oracle Automatic Storage Management (ASM)** protects against data loss by striping and mirroring data across multiple disks.
- **Oracle Real Application Clusters (RAC)** protect against single instance component failures by enabling multiple instances to share access to one Oracle database.
- **Oracle Data Guard** is a disaster recovery solution for Oracle database. It protects Oracle database against primary site failures, disasters, errors and data corruption.

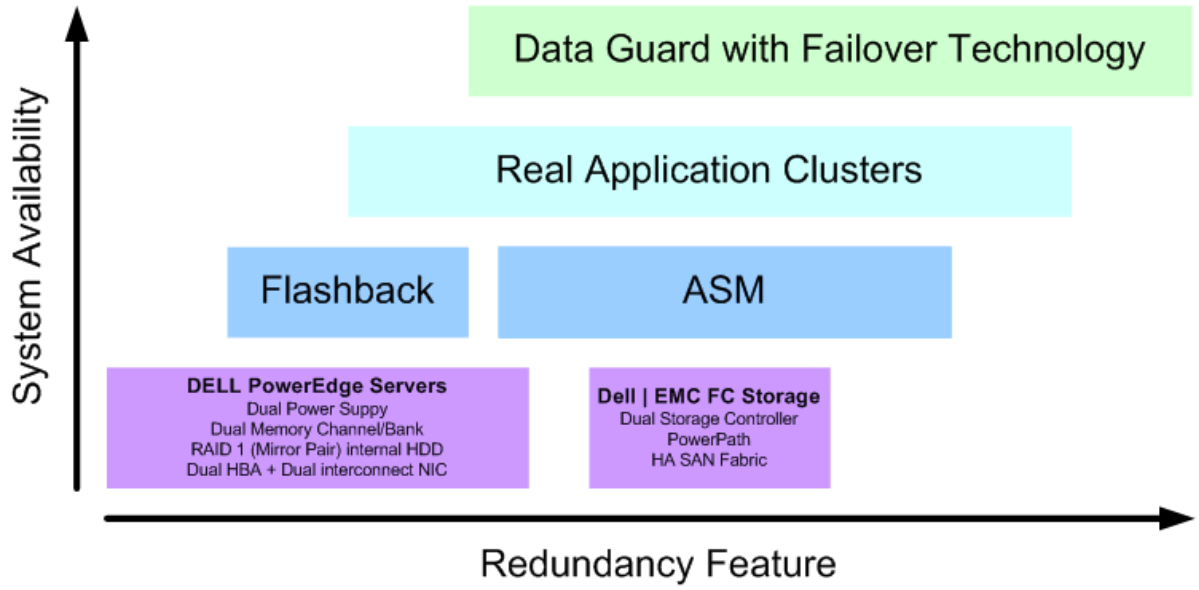

Figure 1 - System Availability vs. Redundancy Feature

This white paper introduces a method for deploying Data Guard using a standby RAC database with minimum downtime to the primary RAC database.

Oracle Data Guard Overview

When a primary database becomes unavailable, Data Guard can quickly switch the standby database to the production role, thus greatly reducing downtime caused by primary site outage. Data Guard is highly integrated with Oracle Database 10*g* and it is a primary choice for data protection and disaster recovery of Oracle databases. Implementing Data Guard for Oracle Database 10*g* with RAC provides end to end data protection and high availability as illustrated in Figure 2.

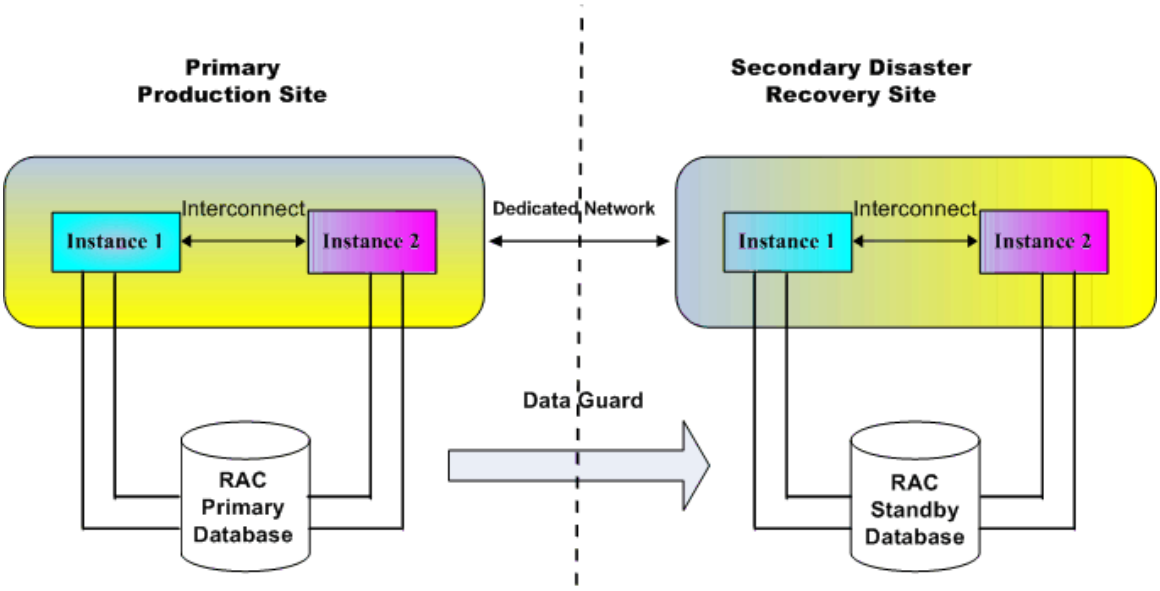

Figure 2 - Oracle Database 10*g* RAC with Data Guard

A Data Guard configuration consists of one primary database and one or more (up to nine) standby databases. All members of a Data Guard configuration must run an Oracle image that is built for the same platform. Primary database and standby databases can be single instance databases or multi-instance RAC databases. Databases within a Data Guard configuration are connected and communicated via Oracle Net over the network. Data Guard synchronizes the standby database with the primary database by automatically transmitting redo data from the primary database to the standby database, and applying redo data on the standby database.

There are two types of Data Guard standby databases – physical standby database and logical standby database.

- A physical standby database is a physically identical copy of the primary database. Data Guard synchronizes the physical standby database via media recovery.
- A logical standby database contains the same logical information as the primary database, although the physical structure of the database can be different. Data Guard converts the redo log data into SQL statements, and executes these SQL statements on standby database.

Traditional Data Guard Deployment Method using RMAN

Steps for Creating an Oracle 10*g* RAC Standby Database with ASM using RMAN

To create an Oracle Database 10*g* RAC standby database with ASM, for an Oracle primary RAC database with ASM, the steps can be outlined as:

1. **Deploy standby dummy RAC database**: Deploy a standby RAC database with the same database name as the primary database.
2. **Configure Oracle Net:** On both primary and standby systems, add both primary database and standby database TNS aliases in the tnsnames.ora files.
3. **Prepare primary database:** Enable forced logging and ARCHIVELOG mode on primary database which are the prerequisites for Data Guard.
4. **Prepare standby database:** A set of standby database initialization parameters control the automatic fetching of redo log files when there is a redo log gap, as well as preparing standby database for database role changes.
5. **Initial instantiation of standby database:**
    - **Backup primary database using RMAN:** Create a primary database backup copy using RMAN backup command with the INCLUDE CURRENT CONTROLFILE FOR STANDBY clause.
    - **Create standby database using RMAN:** Use RMAN command DUPLICATE TARGET DATABASE FOR STANDBY to create standby database.
6. **Configure primary database:** Pre-configure initialization parameters to prepare the primary database for database role changes.
7. **Start up physical standby database:** Start up redo apply on physical standby database.
8. **Additional steps to create logical standby database:** A logical standby database is created by first creating a physical standby database as outlined above before it can be transitioned to a logical standby database.

Disadvantages of Using RMAN to Instantiate Standby Database

A major step during deployment of a standby database, as outlined in the step 5 of the previous section, is the initial instantiation of the entire standby database. It consists of two tasks. First, RMAN creates a backup copy of the primary database. Then RMAN creates the standby database using the backup copy, by issuing RMAN command DUPLICATE TARGET DATABASE FOR STANDBY. This command automates the following steps:

- Restores the standby control file.
- Restores the primary database backup.
- RMAN leaves the database mounted so user can start redo apply to recover standby database.

To fully synchronize the standby database with the primary database, all archived redo logs, which are generated after the primary database backup start time, are needed for Data Guard recovery. For large primary databases, typically multiple terabytes in size, as well as active primary databases that generate large amount of redo data in a short amount of time, deploying Data Guard using RMAN is a challenging task. This is because performing a full database backup, restore, and recovery can require significant time to complete, even with disk-based backup. Administrators often have to request a prolonged primary production database downtime of many hours or even days to stop new archived redo generation during Data Guard deployment.

Dell|EMC Storage Replication Technology Overview

Dell|EMC storage replication technology provides local and remote storage array-based data replication capabilities. Inside a Dell|EMC CX3 storage array, physical storage is partitioned into RAID groups, then logical units (LUNs). Different LUNs are assigned to different server hosts. The contents of LUNs can be replicated inside the same CX3 array to another LUN of the same size. Alternatively, the contents of LUNs may be replicated to another LUN of the same size on a different CX3 array.

EMC MirrorView is remote data replication software that supports ongoing mirroring of data, including changes, on one LUN in one array (called the primary LUN) to secondary LUN(s) on other array(s). MirrorView runs on the CX3 storage system so no host resources are required to replicate the data. MirrorView offers both synchronous and asynchronous remote mirroring capabilities.

- MirrorView/Synchronous uses synchronous writes to maintain an exact block-for-block copy of the primary storage data in real time at a secondary location. Server writes to primary storage system are acknowledged only after all secondary storage systems commit the data.
- MirrorView/Asynchronous is designed for low-bandwidth requirements. MirrorView/Asynchronous utilizes Delta Sets to collect writes to primary storage system occurred within a specific period of time. Delta Sets reduce the bandwidth needed between primary site and secondary site by only transmitting the latest version of data.

When it comes to replicating Oracle data, Dell|EMC CX3 storage-based replication, such as MirrorView, is an effective method for moving large volumes of enterprise data to remote standby sites. EMC MirrorView can be leveraged for supporting the implementation of Data Guard standby databases, by replacing the traditional method of using RMAN to instantiate the initial standby database copy.

Standby Database Instantiation using EMC MirrorView

This section outlines the general steps involved to deploy Data Guard standby database using EMC MirrorView. The same procedure applies to both MirrorView/A and MirrorView/S. The revised Data Guard deployment procedure follows the same steps as described in the section "Steps for Creating 10g RAC Standby Database with ASM using RMAN", except that MirrorView is used in step 5 to instantiate the standby database, rather than RMAN.

Setting Up MirrorView

MirrorView uses a front-end port on each CX3 storage processor (SP) as a communication channel between the storage systems in a remote mirror configuration. This port is called the mirror port.

For MirrorView to work correctly, a Fibre Channel connection, either direct or through a switch fabric, must exist between the primary and secondary storage systems. SP A mirror port on the primary storage system and the SP A mirror port on the secondary storage system must be zoned together. Similarly, SP B mirror port on the primary storage system and the SP B mirror port on the secondary storage system must be zoned together. Figure 3 shows a sample fabric remote mirror connections.

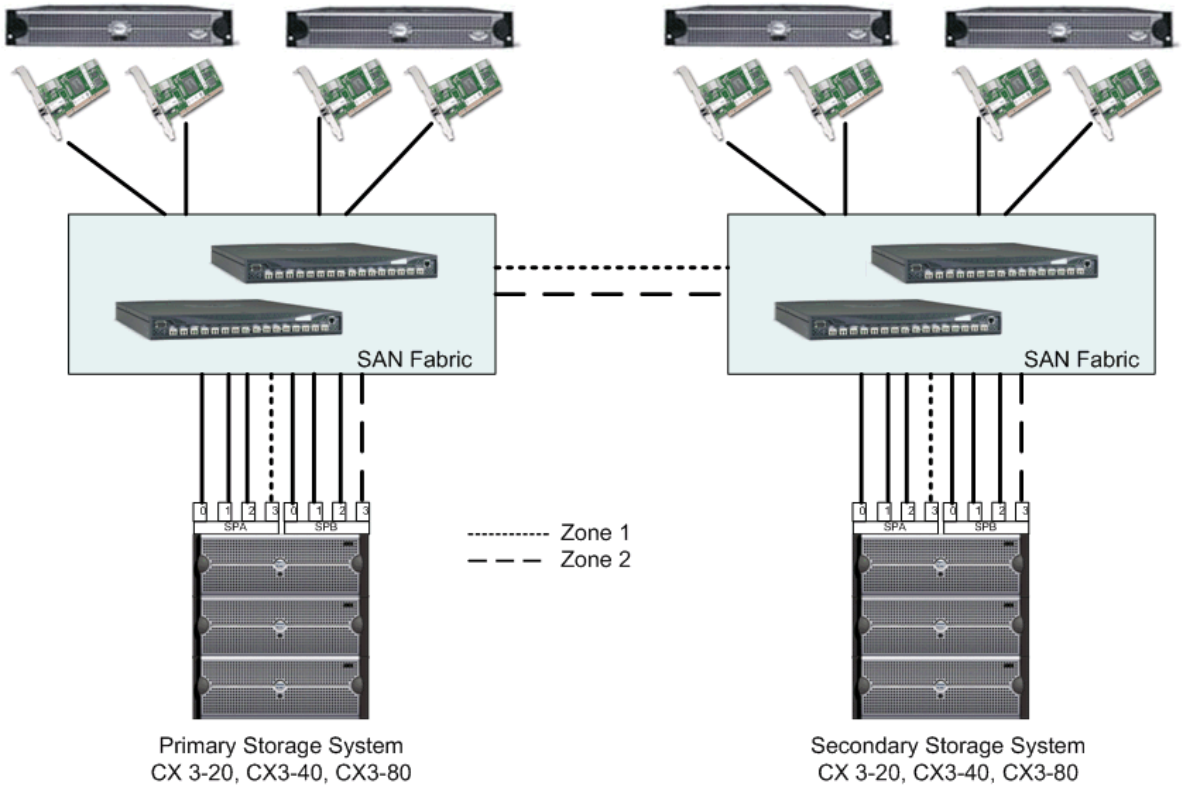

Figure 3 - Sample Remote Mirror Connection Through Fabric

MirrorView configurations can be set up and managed through EMC Navisphere Manager, a Web-based management tool. A usable, two-way connection between the MirrorView storage systems can be established by EMC Navisphere Manager.

Activating MirrorView

To use MirrorView, a remote mirror of the primary database LUNs must be created. For instantiating the standby database, only the LUNs with the contents of the primary database, such as the datafiles, the online redo log files and the archived redo log files, need to be replicated to the standby site. For RAC primary databases, the contents of the Oracle Cluster Registry (OCR) or Voting Disk (CSS) do not need to be replicated. After creating the remote mirror of the primary LUNs, a secondary image can be added to the remote mirror for data synchronization. The secondary image LUN must not be part of a storage group and it must be precisely the same size as the primary LUN.

Oracle databases typically span multiple LUNs where the contents of the LUN might not be content-consistent with another LUN as each LUN is replicated individually. This issue can be addressed by the MirrorView consistency group feature. A consistency group is a set of mirrored pair LUNs that represent content consistent data at the same point-in-time. All members in a consistency group cannot be individually fractured, or synchronized, or promoted. These actions have to be performed on all members of a consistency group as a whole.

When secondary images are synchronized with the primary source, a facture operation is required to stop updates from the primary image to the secondary mirror image. For MirrorView/S, the primary database should be first put in hot backup mode or offline prior to fracturing the consistency group. This will help ensure Oracle database consistency by matching the same SCN (System Control Number) in all datafiles headers and the control files. For MirrorView/A, an additional synchronization update should be issued after putting the primary database in hot backup mode or shutting down the primary database, and before the consistency group is fractured. The additional synchronization is required for MirrorView/A because the mirroring process is completed in batch mode as compared to the real-time synchronization in MirrorView/S.

After fracturing, the secondary images should be removed from the mirror groups prior to re-joining to the storage group on the standby storage for presentation to the standby hosts. The standby controlfile should be restored to the initial standby database copy to convert it to a standby database.

Minimal Downtime

During standby database deployment, administrators can choose not to take the primary database offline by putting it in hot backup mode.

If administrators choose to shut down the primary database, the time taken to fracture a consistency group is a matter of seconds. In the case of MirrorView/A, in order to reduce the time taken for the additional synchronization after shutdown to a matter of a few seconds, it is recommended to issue a synchronization immediately before the shutdown.

Conclusion

The instantiation of the Data Guard standby database has been a challenging task especially for large production databases due to the downtime required when using RMAN. Dell|EMC storage replication technology such as MirrorView is an alternative method for deploying Data Guard. Real world experiences and lab experiments have shown that utilizing MirrorView can significantly reduce the primary database downtime during Data Guard deployment. We have observed this reduction on primary database downtime to be from hours or days, to just a matter of seconds or minutes.

Figures Index

Additional Information

Additional Information on Data Guard and EMC storage replication technologies can be found in the following resources.

1. "CLARiiON Remote Data Replication and Implementing Remote Disaster Protection for an Oracle Deployment", EMC Power Link, January 6, 2004. https://powerlink.emc.com/nsepn/webapps/btg548664833igtcuup4826/km/live1/en_US/Offering_Technical/White_Paper/H1053_CLARiiONDB_ldv.pdf

2. EMC MirrorView/Synchronous for Navisphere Administrator's Guide.

3. EMC MirrorView/Asynchronous for Navisphere Administrator's Guide.

4. Oracle Data Guard Concepts and Administration 10*g* Release 2 (10.2).