ORACLE DATA GUARD DEPLOYMENT USING DELL/EMC CX3 STORAGE REPLICATION

BY WENDY CHEN
UJJWAL RAJBHANDARI
KAI YU

When deploying Oracle® Data Guard, administrators have traditionally created the standby database through Oracle Recovery Manager backup and restore—a method that can lead to significant downtime for large production databases. Using Dell/EMC CX3 storage replication software for this task can help significantly reduce primary database downtime during the deployment process.

Related Categories:

Business continuity

Database

Dell/EMC storage

Oracle

Storage

*Visit **DELL.COM/PowerSolutions** for the complete category index.*

Reliable, highly available systems have become essential for many enterprises, and any downtime for critical applications can lead to financial losses. Following best practices when using Dell™ PowerEdge™ servers, Dell/EMC CX3 Fibre Channel storage, and Oracle high-availability features can help administrators create highly available Oracle database systems. As illustrated in Figure 1, administrators can combine redundancy features to help maximize this availability:

- **Dell PowerEdge servers and Dell/EMC CX3 Fibre Channel storage:** Provides redundancy at the hardware level with features such as dual power supplies, dual host bus adapters (HBAs), and dual network interface cards (NICs) for Dell PowerEdge servers, and dual storage controllers for Dell/EMC CX3 Fibre Channel storage
- **Oracle Flashback:** Helps simplify data recovery by allowing administrators to shift data back and forth in time
- **Oracle Automatic Storage Management (ASM):** Helps protect against data loss by striping and mirroring data across multiple disks
- **Oracle Real Application Clusters (RAC):** Helps protect against single-instance component failures by enabling multiple instances to share access to one Oracle database
- **Oracle Data Guard:** Helps protect Oracle databases against primary site failures, disasters, errors, and data corruption

The Oracle Data Guard deployment process has traditionally utilized Oracle Recovery Manager (RMAN) when creating standby databases in Oracle RAC systems, which can require significant downtime for primary production databases. To help reduce this downtime, however, administrators can instead take advantage of Dell/EMC CX3 Fibre Channel storage replication features.

Oracle Data Guard disaster recovery features

Oracle Data Guard is designed to integrate with other Oracle software to provide disaster recovery for Oracle databases. When the primary database becomes unavailable, Data Guard can quickly switch a standby database to the production role, helping significantly reduce downtime caused during primary site outages. Implementing Data Guard for an Oracle RAC database can provide end-to-end data protection and high-availability capabilities.

A Data Guard–based configuration consists of one primary database and from one to nine standby databases (see Figure 2). All systems in this configuration must run an Oracle image built for the same platform. The primary and standby databases can be single- or multi-instance Oracle RAC databases. Databases within a Data Guard–based configuration connect and communicate with one another through Oracle Net over a network. Data Guard synchronizes the primary and standby databases by automatically transmitting redo data

> *"When the primary database becomes unavailable, Data Guard can quickly switch a standby database to the production role, helping reduce downtime caused during primary site outages."*

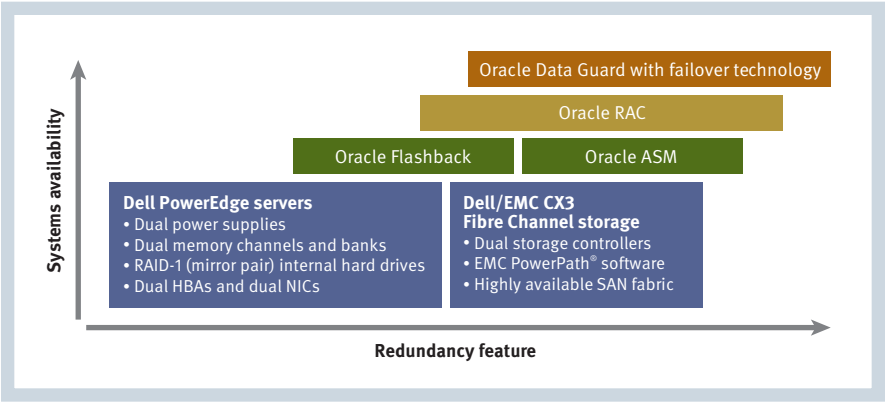

Figure 1. *Increasing systems availability by combining redundancy features*

from the primary database to the standby databases and applying redo data on the standby databases.

Each Data Guard standby database can be either a physical standby database or a logical standby database. A physical standby database is a physically identical copy of the primary database, which Data Guard synchronizes through media recovery. A logical standby database contains the same logical information as the primary database, but the physical structure of the database can be different. In this case, Data Guard converts the redo log data into SQL statements and executes these statements on the standby database.

Traditional Oracle Data Guard deployment using Oracle Recovery Manager

During deployment of a physical or logical standby database for Data Guard–based systems, administrators typically create a backup copy of the primary database using RMAN. To create an Oracle RAC 10*g* physical standby database with Oracle ASM for an Oracle RAC 10*g* primary database with ASM, they would perform the following steps:

1. Deploy a placeholder standby database with the same database name as the primary database.
2. Configure Oracle Net on the primary and standby systems by adding both the primary and standby database Transparent Network Substrate (TNS) aliases in the tnsnames.ora files.
3. Prepare the primary database by enabling forced logging and the ARCHIVELOG mode, which are prerequisites for using Data Guard.
4. Prepare the standby database by configuring the database initialization parameters that control the automatic fetching of redo log files when there is a redo log gap and that prepare the standby database for database role changes.
5. Create a primary database backup copy using the RMAN command `backup database` with the clause `include current controlfile for standby`, then create the standby database using the RMAN command `duplicate target database for standby`.
6. Configure the primary database by setting the initialization parameters to prepare the primary database for database role changes.
7. Start the redo apply process on the physical standby database.

If administrators want to create a logical standby database, they must first create a physical standby database as described in the preceding list, then transition it to a logical standby database.

The initial creation of the standby database, as outlined in step 5, can be a significant task. After creating a backup copy of the primary database, administrators create the standby database using the RMAN command `duplicate target database for standby`. This command automates restoring the standby control file and the primary database backup, and leaves the database mounted so administrators can start the redo apply process to recover the standby database.

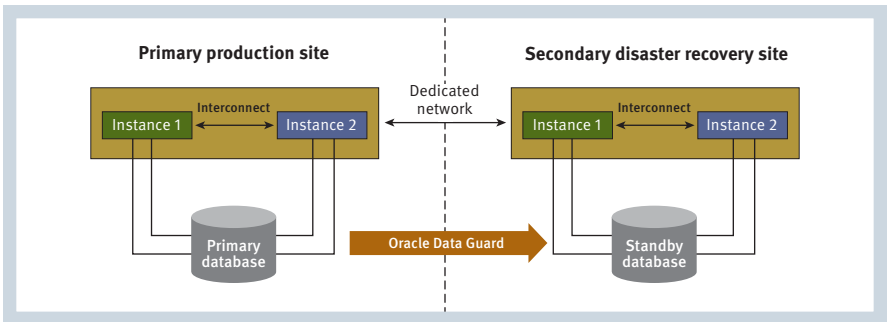

Figure 2. *Using Oracle Data Guard with an Oracle RAC database*

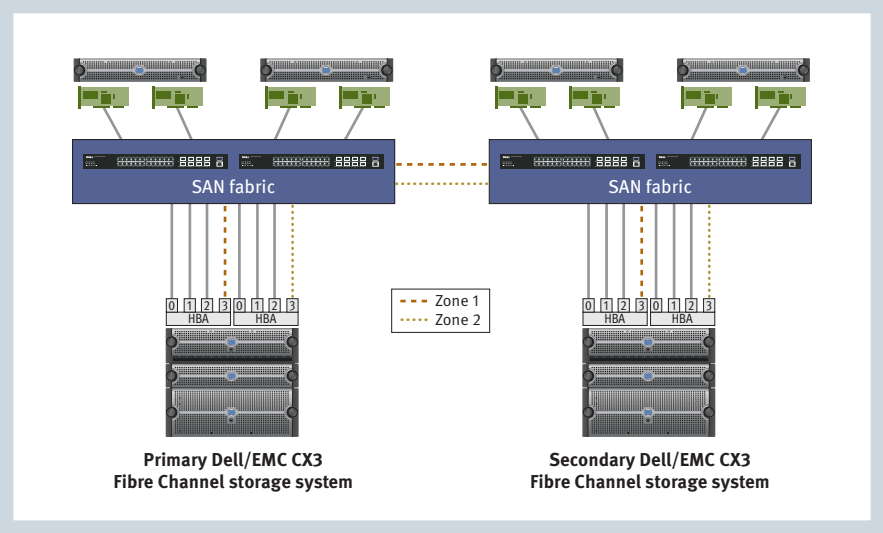

Figure 3. *Creating an EMC MirrorView remote mirror connection through a Fibre Channel–based SAN fabric*

Fully synchronizing the standby database with the primary database requires all archived redo logs, which are generated after the primary database backup, to enable Data Guard recovery. For multi-terabyte primary databases, as well as active primary databases that generate a large amount of redo data in a short amount of time, deploying Data Guard using RMAN can be challenging: a full database backup, restore, and recovery may require significant time to complete, even with disk-based backup. Administrators often must request primary database downtime of hours or even days to stop new archived redo generation during this deployment.

Dell/EMC CX3 storage replication features

Dell/EMC CX3 Fibre Channel storage replication features provide local and remote array-based data replication capabilities. Administrators can partition each physical storage system into RAID groups, then logical units (LUNs). They can then assign different LUNs to different servers, and can replicate LUN contents to another LUN of the same size in the same array or in a different array.

EMC® MirrorView® software supports ongoing mirroring of data, including changes, on one LUN in an array (called the primary LUN) to secondary LUNs on other arrays. Because MirrorView runs on the storage system, no host resources are required to replicate the data. It offers two types of remote mirroring capabilities:

- **MirrorView/Synchronous (MirrorView/S):** MirrorView/S 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 the primary storage system are acknowledged only after the secondary storage systems have committed the data.
- **MirrorView/Asynchronous (MirrorView/A):** Designed for low-bandwidth requirements, MirrorView/A utilizes delta sets to collect writes to the primary storage system that occurred within a specific period of time, which helps reduce the bandwidth needed between the primary and secondary sites by only transmitting the latest versions of data.

Dell/EMC CX3 Fibre Channel storage–based replication software such as MirrorView can provide an effective means of moving large volumes of enterprise data to remote standby sites. Administrators can take advantage of MirrorView to support the implementation of Data Guard standby databases by replacing the traditional RMAN method when creating the initial standby database copy.

Oracle Data Guard deployment using EMC MirrorView

To use MirrorView when creating a Data Guard standby database, administrators can follow the same steps described in the "Traditional Oracle Data Guard deployment using Oracle Recovery Manager" section in this article, but use MirrorView in step 5 instead of RMAN. The same procedure applies to both MirrorView/S and MirrorView/A.

Setting up EMC MirrorView

MirrorView uses a front-end port, called the mirror port, on each storage processor (SP) as a communication channel between the storage systems in a remote mirror configuration. For MirrorView to work correctly, a Fibre Channel connection must exist between the primary and secondary storage systems, either directly or through a switch fabric. The 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, the 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 an example remote mirror connection through a Fibre Channel–based storage area network (SAN) fabric.

Administrators can use the EMC Navisphere® Manager Web browser–based management tool to set up and manage MirrorView configurations and create two-way connections between storage systems.

"Dell/EMC CX3 Fibre Channel storage replication features provide local and remote array-based data replication capabilities."

"Dell/EMC CX3 Fibre Channel storage–based replication software such as EMC MirrorView can help significantly reduce downtime for critical database systems."

Activating EMC MirrorView

To use MirrorView, administrators must create a remote mirror of the primary database LUNs. When a standby database is used, only the LUNs with the contents of the primary database—such as the data files, online redo log files, and archived redo log files—need to be replicated to the standby site. For Oracle RAC primary databases, the contents of the Oracle Cluster Registry or the Oracle Cluster Synchronization Services Voting Disk do not need to be replicated. After creating the remote mirror of the primary LUNs, administrators can add a secondary image to the remote mirror for data synchronization. The secondary image LUN must not be part of a storage group, and it must be the same size as the primary LUN.

Oracle databases typically span multiple LUNs, and because each LUN is replicated individually, the contents of one LUN might not be content-consistent with other LUNs. The MirrorView consistency group feature can help address this issue. A consistency group is a set of mirrored-pair LUNs that represent content-consistent data at a particular point in time. Administrators cannot individually fracture, synchronize, or promote members of a consistency group; they must perform these actions on all members of a consistency group at once.

When secondary images are synchronized with the primary source, a fracture operation is required to stop updates from the primary image to the secondary mirror image. When using MirrorView/S, administrators should first put the primary database in hot backup mode or shut it down before fracturing the consistency group. Doing so helps ensure Oracle database consistency by matching the same System Change Number in the data file headers and control files. When using MirrorView/A, administrators should run an additional synchronization update after putting the primary database in hot backup mode or shutting it down and before fracturing the consistency group. This additional synchronization is required because MirrorView/A completes the mirroring process in batch mode, rather than using the real-time synchronization of MirrorView/S.

After fracturing, administrators should remove the secondary images from the mirror groups before rejoining them to the storage group on the standby storage to present them to the standby hosts. They should then restore the standby control file to the initial standby database copy to convert it to a standby database.

Minimizing primary database downtime

During standby database deployment, administrators can choose not to take the primary database offline by putting it in hot backup mode. If administrators do choose to shut down the primary database, they can quickly and easily fracture a consistency group. When using MirrorView/A, to help reduce the time for the additional synchronization after shutdown, administrators should run a synchronization immediately before the shutdown.

Efficient Oracle Data Guard deployment with Dell/EMC CX3 storage

Creating Oracle Data Guard standby databases using RMAN has typically been a challenging and time-consuming task that can require significant downtime, especially for large production databases. By providing a viable alternative to RMAN for standby database creation, Dell/EMC CX3 Fibre Channel storage–based replication software such as EMC MirrorView can help significantly reduce this downtime for critical database systems.

Wendy Chen *is a systems engineer on the Dell Database, Applications, and HA team. Her interests include Oracle database backup and recovery, disaster recovery, and Oracle RAC, and she has eight years of experience in Oracle database administration, application development, and Oracle solutions engineering. Wendy has an M.S. in Computer Science from Brigham Young University.*

Ujjwal Rajbhandari *is a systems engineer in the Database and Application Engineering Department of the Dell Product Group. He is responsible for validating Oracle RAC solutions on Dell PowerEdge servers and Dell/EMC Fibre Channel storage. Ujjwal has a B.E. in Electrical Engineering from the Indian Institute of Technology, Roorkee, and an M.S. in Electrical Engineering from Texas A&M University.*

Kai Yu *is a senior systems engineering consultant on the Dell Database and Applications Solutions engineering team. He specializes in Oracle RAC, database migration, and the Oracle E-Business Suite, and has previously been an Oracle database and applications database administrator. Kai has a B.E. and M.E. in Computer Science from Huazhong University of Science and Technology and an M.S. in Computer Science from the University of Wyoming.*

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell/EMC storage:
DELL.COM/EMC

Dell and Oracle solutions:
DELL.COM/Oracle