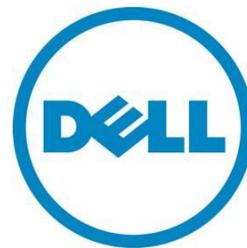

Oracle Database Disaster Recovery Using Dell Storage Replication Solutions

This paper describes how to leverage Dell storage replication technologies to build a comprehensive disaster recovery solution for an Oracle Database involving both synchronous and asynchronous options which helps customers to choose the best replication option depending on customer's infrastructure and feasibility.

Krishna Kamal Kapa

Ramamohan Reddy K

Manjunatha Reddy MC



Dell Oracle Database Solutions Engineering

This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

© 2013 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

August 2013 | Rev A00

Contents

Executive summary	5
1 Introduction	6
2 Technology Overview	7
2.1 Oracle Data Guard	7
2.2 Storage Replication	7
2.2.1 Synchronous Replication	7
2.2.2 Asynchronous Replication.....	7
2.3 Dell Storage Replication Features.....	8
2.3.1 Compellent Remote Instant Replay	8
2.3.2 EqualLogic Snapshot Replication	8
2.3.3 PowerVault Remote Replication.....	8
3 Dell Storage Disaster Recovery Solutions	9
3.1 Compellent Disaster Recovery Solution	9
3.1.1 Asynchronous Replication.....	9
3.1.2 Synchronous High Availability	10
3.1.3 Synchronous High Consistency	10
3.2 PowerVault Disaster Recovery Solution	12
3.2.1 Standard Remote Replication	12
3.2.2 Legacy Remote Replication	12
3.2.2.1 Synchronous Write Mode	13
3.2.2.2 Asynchronous Write Mode.....	14
3.2.2.3 Asynchronous Write Order Consistency Mode	14
3.3 Equallogic Disaster Recovery Solution	15
4 Test Environment	16
5 Test Scenarios	18
5.1 Dell Compellent Storage	18
5.1.1 Asynchronous Replication Disaster Recovery Simulation	18
5.1.2 Synchronous High Availability Replication Disaster Recovery Simulation	19
5.1.3 Synchronous High Consistency Replication Disaster Recovery Simulation.....	20
5.2 PowerVault Storage	20
5.2.1 Asynchronous Standard Remote Replication Disaster Recovery Simulation	20
5.2.2 Synchronous Write Mode Disaster Recovery Simulation	21
5.2.3 Asynchronous Write Mode Disaster Recovery Simulation.....	22
5.2.4 Asynchronous Write Order Consistency Disaster Recovery Simulation.....	22

Conclusion 23
References 23
Appendix A..... 24
Appendix B..... 25

Tables

Table 1. Comparison of different Compellent Replications 11
Table 2. Comparison of different PowerVault Replications 15
Table 3. Consistency group 16
Table 4. Database volume configuration..... 16

Figures

Figure 1. Asynchronous Replication with Replays 9
Figure 2. Synchronous High Availability Between Storage Arrays 10
Figure 3. Synchronous High Consistency Between Storage Arrays..... 11
Figure 4. Synchronous Write Mode Between Storage Arrays 13
Figure 5. Asynchronous Write Mode Between Storage Arrays 14
Figure 6. Asynchronous Write Order Consistency Between Storage Arrays 15
Figure 7. Test environment 17
Figure 8. Asynchronous Replication between storage arrays..... 18
Figure 9. Link Failure in High Availability Between Storage Arrays 19
Figure 10. Link Failure in Synchronous High Consistency Between Storage Arrays 20

Executive summary

Disaster Recovery (DR) is the business continuity solution to recover and protect a business infrastructure in case of natural disasters or human errors. Business continuity involves much more than recovering an IT infrastructure from a disaster. It also involves key aspects like Recovery Point Objective (RPO¹) and Recovery Time Objective (RTO²). Today's business continuity solutions must keep organizations up and running despite unplanned interruptions.

A comprehensive business continuity solution must address redundancy, data protection and recoverability of data according to the defined RPOs.

With Dell storage replication solutions, organizations can now deploy a reliable disaster recovery and business continuity solution and at the same time reduce overall operating costs and complexity.

This whitepaper describes how to leverage Dell storage replication solutions to build a comprehensive disaster recovery plan for an Oracle Database involving both synchronous and asynchronous options. This enables you to choose the best replication option depending on the infrastructure and feasibility.

The whitepaper also discusses the setting up of different replication methods between storage arrays for all the Dell storages and performing different disaster recovery test cases while validating the integrity of the Oracle Database.

The storage replication solutions discussed in this paper covers Dell Compellent, EqualLogic, and PowerVault storage systems.

¹ Maximum tolerable period in which data might be lost from an IT service due to a major incident.

² The duration of time and service level within which a business process must be restored after a disaster for business continuity.

1 Introduction

Dell storage replication solution portfolio of highly efficient array-based replication technologies enables advanced business continuity and disaster recovery solutions.

Dell storage replication solution creates an exact mirror copy of data on a remote system that can be mounted to rapidly recover from a failure. It provides minimal (asynchronous replication) to zero (synchronous replication) data loss environment for applications. It also reduces the RPO and RTO associated with traditional DR methods. Dell's storage replication solution enables non-disruptive scheduled maintenance, DR testing and site workload reallocations.

In the Dell storage replication solution, Dell storage is integrated with centralized DR management products to automate the process of setting up and configuring recovery plans. It also provides functionality to efficiently test these plans with no impact to production environments.

The key benefits of Dell storage replication solutions are

- Optimized management and minimal storage capacity using thin provisioning for snapshots.
- Ability to meet stringent business requirements for operational recovery.
- Advanced centralized recovery management integration.
- Rapid recovery from planned and unplanned outages.

2 Technology Overview

A disaster recovery solution is to apply the block changes from the primary site to the standby site using either an application level DR or a storage replication.

In the following sections we discuss about Oracle Data Guard and Dell storage replication solutions for achieving the disaster recovery.

2.1 Oracle Data Guard

Oracle Data Guard provides the management, monitoring, and automation software to create and maintain one or more standby databases to protect Oracle data from failures and data corruptions while also providing high availability for mission critical applications. 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, minimizing the downtime associated with the outage.

There are two types of standby databases:

- **Physical standby database:** A physical standby database is kept synchronized with the primary database through **Redo Apply**, which recovers the redo data, received from the primary database and applies the redo data to the physical standby database.
- **Logical standby database:** A logical standby database is kept synchronized with the primary database through **SQL Apply**, which transforms the data in the redo, received from the primary database into SQL statements and then executes the SQL statements on the standby database.

2.2 Storage Replication

Storage Replication is the process of mirroring mission critical data from one site to another. The replication process can be synchronous or asynchronous as only replication can provide instant access to data and zero data loss (minimal Recovery Point Objective or RPO and minimal Recovery Time Objective or RTO).

2.2.1 Synchronous Replication

Synchronous replication guarantees data consistency (zero data loss) between replication primary and standby. This is achieved by ensuring write incident commitments at the primary and standby before a successful write acknowledgement is sent back to the host and its application.

In the event the write incidents cannot be committed at the source or destination, the write incidents are not committed at either location to ensure consistency. Synchronous replication provides continuous data protection.

2.2.2 Asynchronous Replication

In asynchronous replication, the manner and frequency in which the data is replicated diverges from synchronous replication. Instead of committing a write incident at both the replication source and destination simultaneously, the write incident is committed only at the source and an acknowledgement is sent to the storage host and application. At some point and depending on the frequency of replay, the accumulation of committed write incidents at the source volume is replicated to the destination volume.

2.3 Dell Storage Replication Features

Dell has different storage capabilities that can be used for disaster recovery solutions thereby reducing the time, cost, and data recovery efforts associated with traditional disaster recovery methods.

2.3.1 Compellent Remote Instant Replay

Compellent Remote Instant Replay (Remote Replication) replicates continuous snapshots, called “Replays” between local and remote sites, ensuring business continuity at a fraction of the cost of other replication solutions. Remote Instant Replay supports traditional Fibre Channel replication or cost-effective IP-based (iSCSI) replication which brings disaster recovery within the reach of different budgets.

2.3.2 EqualLogic Snapshot Replication

Equallogic snapshot replication is a point-in-time snapshot-based data replication method that provides data protection by copying volume data from the primary storage group to a remote secondary storage group. If the primary site should fail, data access can continue at the secondary site by failing over to the secondary site. Data is also recovered from data volumes created by replication in the secondary storage group. Replicas are similar to snapshots as they represent the contents of the volume at a specific point-in-time. Each replicated volume has a replica set that is the set of replicas created over time that delivers multiple recovery points.

2.3.3 PowerVault Remote Replication

PowerVault remote replication is a RAID controller module-level data replication between a local site and a remote site on a per-virtual disk basis. A remote replicated pair consists of two standard virtual disks, a primary virtual disk and a secondary virtual disk that contains identical copies of the same data that is part of the remote replication group. The remote replicated pair synchronizes at the same time as any other replicated pairs within the remote replication group.

3 Dell Storage Disaster Recovery Solutions

Dell storage disaster recovery solutions protect your business and business viability, and balance cost and risks enabling effective DR solutions.

3.1 Compellent Disaster Recovery Solution

[Compellent Remote Instant Replay technology](#) leverages space-efficient snapshots to provide a cost-effective solution for multi-site data protection. After the initial site synchronization, only incremental changes in data need to be replicated. This approach, also known as thin replication, not only reduces hardware costs, but minimizes bandwidth requirements and speeds recovery in the event of a disaster. You can choose between synchronous and asynchronous replication.

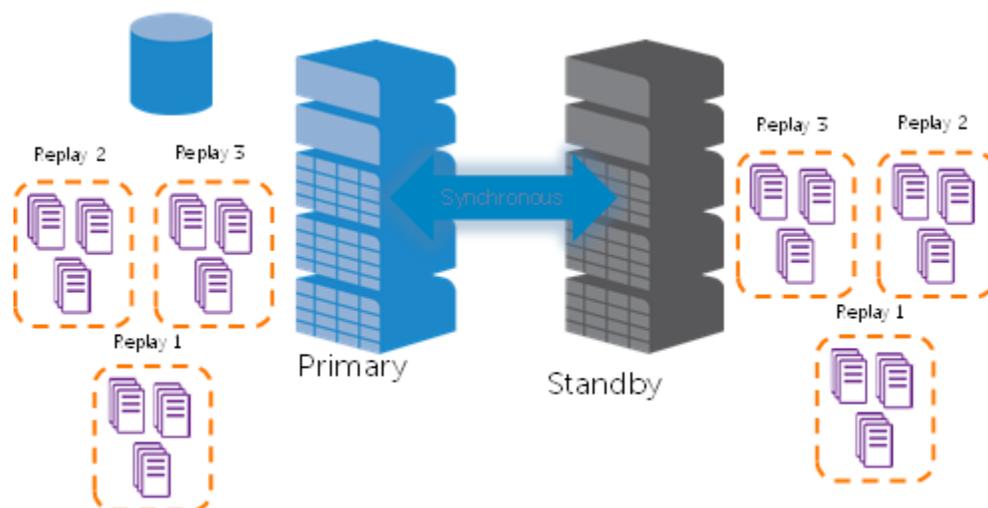
There are 3 types of replications available in Compellent.

3.1.1 Asynchronous Replication

In asynchronous replication, instead of committing a write at both the replication source and the destination simultaneously, the write is committed at the source only. An acknowledgement is then sent to the storage host and application. At some point in time, the accumulation of committed write at the source volume are replicated to the destination volume in batches and committed to the destination volume.

Dell Compellent Storage Center asynchronous replication is tied to the source volume's Replay schedule. When a Replay is created on the source volume, and that volume is configured for asynchronous replication, the volume's newly created Replay is replicated to the destination volume after it is frozen. Replays on a volume are either manually created or they are created as a scheduled task on a per volume basis. This type of replication may also be referred to as Point-in-time replication. This is a type of asynchronous replication that specifically leverages volume snapshots. Since asynchronously replicated transactions are not required to wait for write committals at the replication destination volume, the replication link and/or destination storage do not contribute to the application or the transaction latency at the source volume.

Figure 1. Asynchronous Replication with Replays



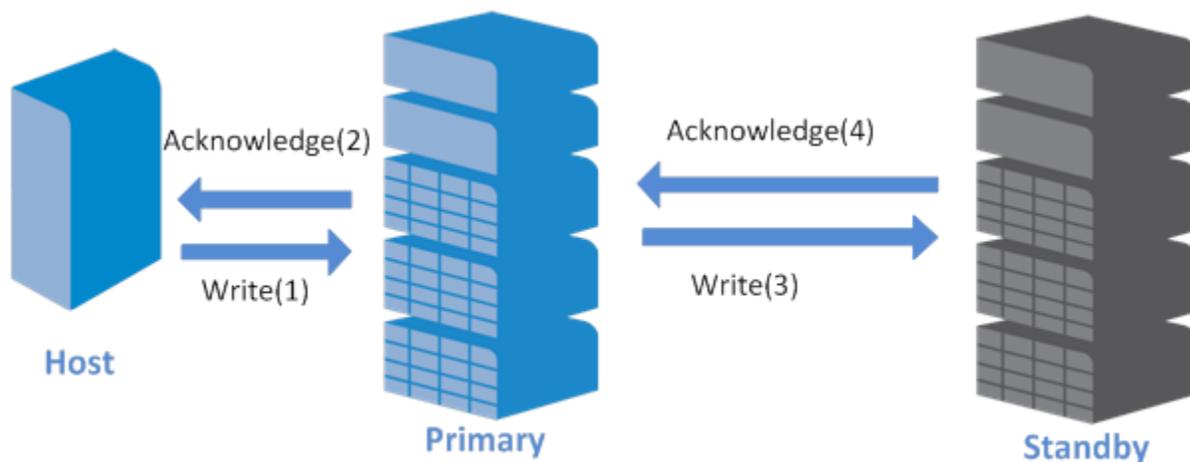
3.1.2 Synchronous High Availability

In High Availability synchronous replication mode, write incidents from the application are written to the primary first and acknowledged back to the application. The changes are then synchronized to the standby site.

In this replication mode, Check Pointing³ is initiated to track all changes to the primary volume when the standby volumes are not available or latency between the sites is too high.

Once the standby storage is available, the active replay from the primary is copied to the standby and the replication continues.

Figure 2. Synchronous High Availability Between Storage Arrays

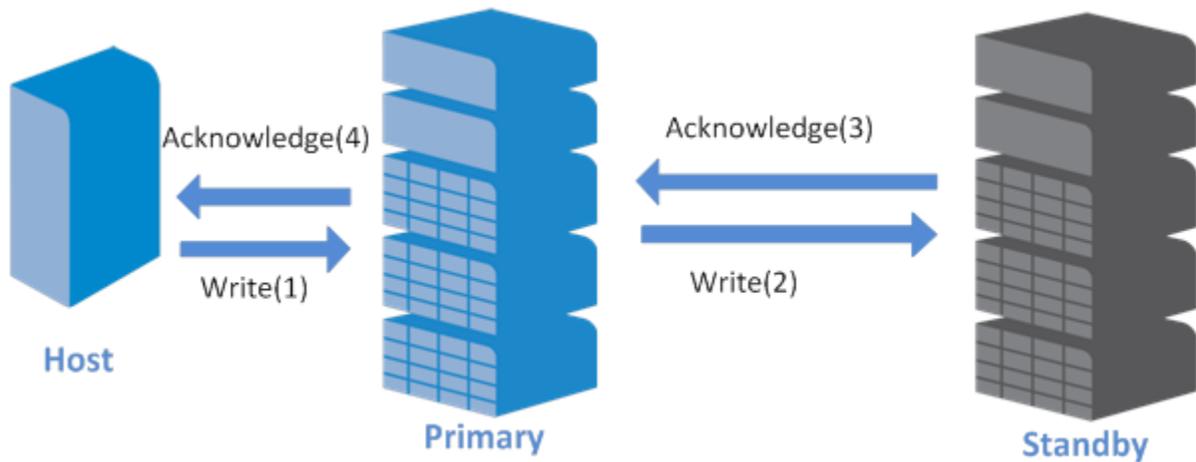


3.1.3 Synchronous High Consistency

In high consistency synchronous replication mode, write incidents from the application are written to both the primary and the standby, and then acknowledged by the application. Write incidents from the application are failed if either the primary volume or standby volume is not available. All write incidents are failed to ensure that the primary and the standby are always identical. This policy ensures there is no data loss at the standby site in the event there is a failure at the primary site.

³ Check Pointing is the functionality that tracks changes upon communication failure in order to minimize the amount of data that has to be copied to the standby when communication is reestablished.

Figure 3. Synchronous High Consistency Between Storage Arrays



The following table shows the differences of Compellent replication modes and the level of data protection.

Table 1. Comparison of Different Compellent Replications

Mode	Risk of Data Loss	Protection	Transport	No Response from Standby
Synchronous High Consistency	Zero Data Loss	Double Failure Protection	SYNC	Stall primary until acknowledgement is received
Synchronous High Availability	Zero Data Loss	Single Failure Protection	SYNC	Primary never waits for standby acknowledgement
Asynchronous	Depends on replay schedule time	Minimal Data Loss	ASYNC	Changes applied during replay cycle

3.2 PowerVault Disaster Recovery Solution

The PowerVault Storage Management Software (MDSM) provides two types of remote replication methods, standard remote replication and legacy remote replication.

Note: Remote replication operations are supported in conjunction with the asynchronous standard remote replication feature, but a given virtual disk cannot participate in both legacy remote replication and asynchronous standard remote replication pairs.

3.2.1 Standard Remote Replication

Standard remote replication is an asynchronous remote point-in-time replication method and is an enhanced replicating option to the legacy remote replication premium feature. It captures the state of the source virtual disk at a particular point-in-time using an embedded snapshot process and includes just the data that has changed since the last image capture. In this method, the remote storage array is not fully synchronized with the local storage array. If any disaster of the primary storage site occurs, the transition to the remote storage array might lose some transactions.

In general, the database applications data is spread across multiple virtual disks and hence, it is very important that all virtual disks related to that application are replicated to the secondary site in a way that makes group-level recoverability possible.

Dell PowerVault replication consistency group for consistent data synchronization of the primary virtual disk and secondary virtual disk is referred to as remote replication group and a remote replication pair.

A remote replication group is a container that can house several replicated pairs, which is comprised of two virtual disks, a primary virtual disk and a secondary virtual disk. It is associated with a local storage array and remote storage array virtual disks having similar synchronization settings.

A remote replicated pair is comprised of two virtual disks, a primary virtual disk and a secondary virtual disk, that contain identical copies of the same data. The replicated pair is a part of a remote replication group, which allows the replicated pair to synchronize at the same time as any replicated pairs within the replication group.

3.2.2 Legacy Remote Replication

Legacy remote replication is a continuous real-time data replication between the primary and secondary storage arrays. In this mode, any changes made on the source virtual disk are reflected instantly on the target virtual disk. Hence the primary virtual disk and the secondary virtual disk are identical at every moment.

In this mode, one or more replicated pairs that consist of a primary virtual disk at the primary site and a secondary virtual disk at a secondary remote site are created, and then a replication relationship between the two virtual disks is established.

Note: Legacy remote replication is supported only on FC storage arrays and hence the replication distance is restricted to meet the latency and application performance requirements.

When the primary virtual disk receives a write request from a host, the RAID controller module first logs information about the write to a replication repository virtual disk. It writes the data to the primary virtual disk and then the RAID controller module initiates a remote write operation to copy the affected data blocks to the secondary virtual disk at the remote site. Then the RAID controller module

sends an I/O completion indication back to the host to confirm that the data is copied successfully to the secondary storage array. The time of sending the I/O completion acknowledgement is dependent on the write mode selected during the remote replication creation.

PowerVault MD storage management software provides two write modes - synchronous write mode and asynchronous write mode. In asynchronous write mode, write consistency mode group is available to maintain the same order of write requests on the secondary virtual disk to failover during the primary site disaster.

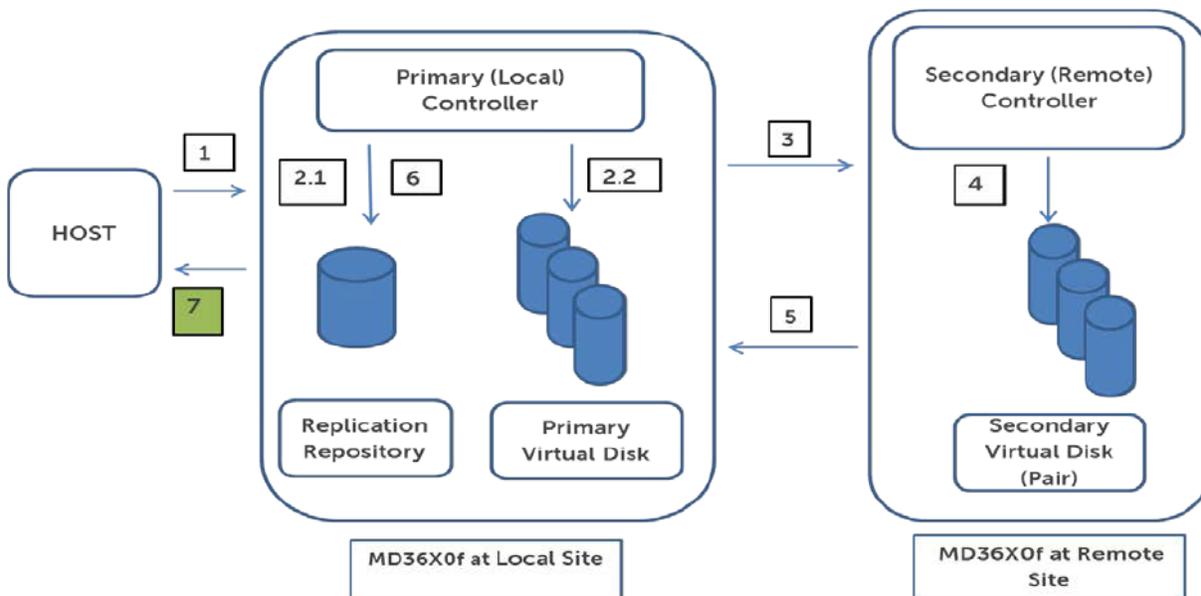
Write consistency group is used to retain the same order of write requests on the secondary virtual disk as they are on the primary virtual disk. If one virtual disk that exists in a write consistency group becomes unsynchronized, all of the virtual disks in the write consistency group become unsynchronized. In addition, any write activity to the remote secondary storage array is prevented to protect the consistency of the remote data set.

There are three modes of legacy remote replication - synchronous, asynchronous and asynchronous with write order consistency replication modes.

3.2.2.1 Synchronous Write Mode

Synchronous write mode offers the best chance of full data recovery from the secondary storage array in the event of a disaster at the primary storage site. This mode of replication occurs at the expense of host I/O performance. In this mode, the primary storage array waits until the data has been successfully copied to the standby storage before sending an I/O completion back to the host system.

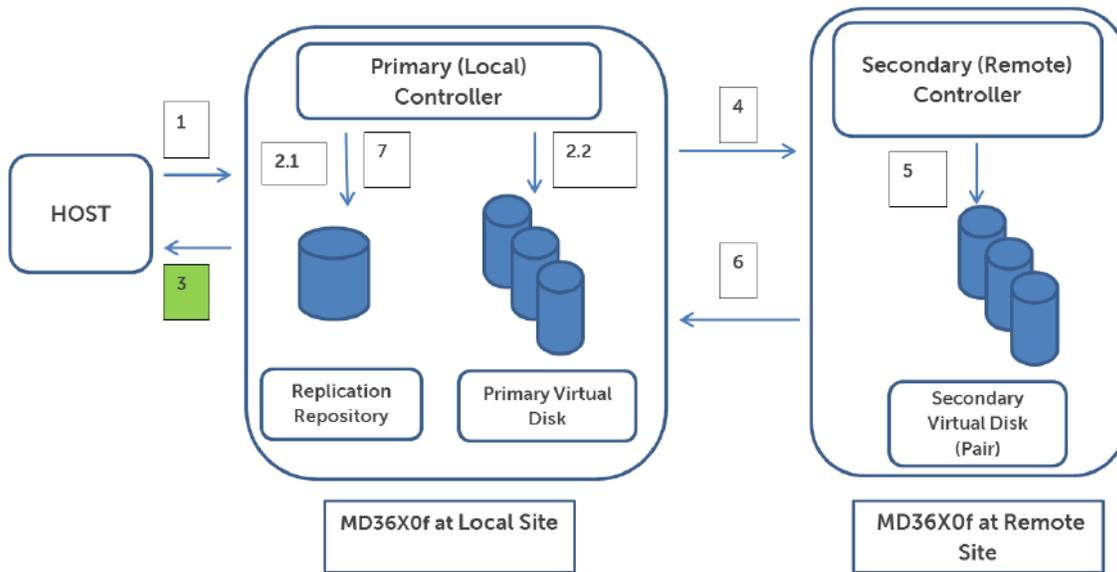
Figure 4. Synchronous Write Mode Between Storage Arrays



3.2.2.2 Asynchronous Write Mode

Asynchronous write mode offers faster host I/O performance, but the mode does not guarantee that the data is successfully written to the secondary virtual disk before indicating a successful write to the host system. In this mode, the primary storage array sends an I/O completion to the host system before the data is successfully copied to the standby storage array.

Figure 5. Asynchronous Write Mode Between Storage Arrays



3.2.2.3 Asynchronous Write Order Consistency Mode

For database applications that are out of order, write incidents cause the database recovery to fail and lead to data inconsistency and data loss. Databases operate under a strict sequencing protocol of write incidents to tablespaces and logs. The order of the writes to each of the database tablespace objects that reside on the virtual disks must be the same as the remote write incidents applied at the remote storage array. In order to ensure that write incidents to the secondary storage array are completed in the same order as the primary storage array, the remote replication provides asynchronous write order consistency replication mode by enabling add to write consistency mode group.

Figure 6. Asynchronous Write Order Consistency Between Storage Arrays

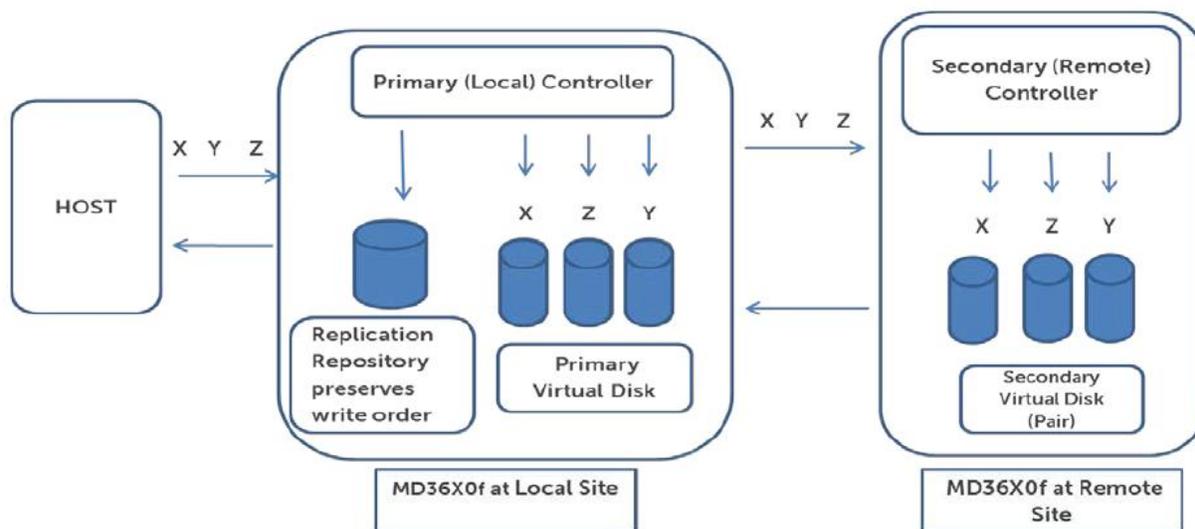


Table 2. Comparison of Different PowerVault Replications

Mode	Replication Method	Repository	Synchronization	Communication	Distance
Remote Replication(Standard)	Point-in-Time	Multiple	Writes Only Changed data	iSCSI and Fiber Channel	Unlimited
Remote Replication(Legacy)	Continuous	Single	Writes all data	Fiber Channel	Restricted

3.3 Equallogic Disaster Recovery Solution

EqualLogic operates a point-in-time snapshot SAN replication model. Pointer based snapshots are replicated to the secondary PS series group where large amounts of data can be replicated over lower bandwidth links efficiently.

EqualLogic SyncRep is the simultaneous writing of volume data across two different storage pools in the same PS Series Group, resulting in two hardware-independent copies of volume data. SyncRep only acknowledges write completion when both the source and target volume is synchronized. If one pool is not available due to a power failure or other resource outage, you can still obtain volume data from the other pool.

For synchronous replication, see [Dell EqualLogic PS Series Array: Understanding Synchronous Replication](#).

4 Test Environment

In our study we performed an Oracle database storage replication between a primary and standby storage using different storage features. We simulated a few disaster scenarios at the primary site and performed the recovery of Oracle database at the standby site and verified the integrity of the Oracle Database. The configurations set up for each Dell storage are described below:

1. Two storages arrays, one as the primary and the other as the secondary for performing the replication of the Oracle database for each of the Dell storage configurations is set up.
2. The latest Storage Center 6.3.1 for Compellent or the latest MDSM software for PowerVault storage is installed on both the primary and standby storages.
3. A Consistency Group is created using the volumes in Table 3 in each of the storage.

Table 3. Consistency group

Consistency Group
DATA
REDO1
REDO2
FRA

4. Replication Feature license on each system is installed for all the storages.
5. A two-node Oracle 11gR2 grid Infrastructure and a RAC database referred to as the production database environment is deployed for each of the Dell storage solutions.
6. Using Dell Quest Benchmark factory, a 300GB database is created to perform the replication test cases. The database is configured with 4 volumes as described in the following table.

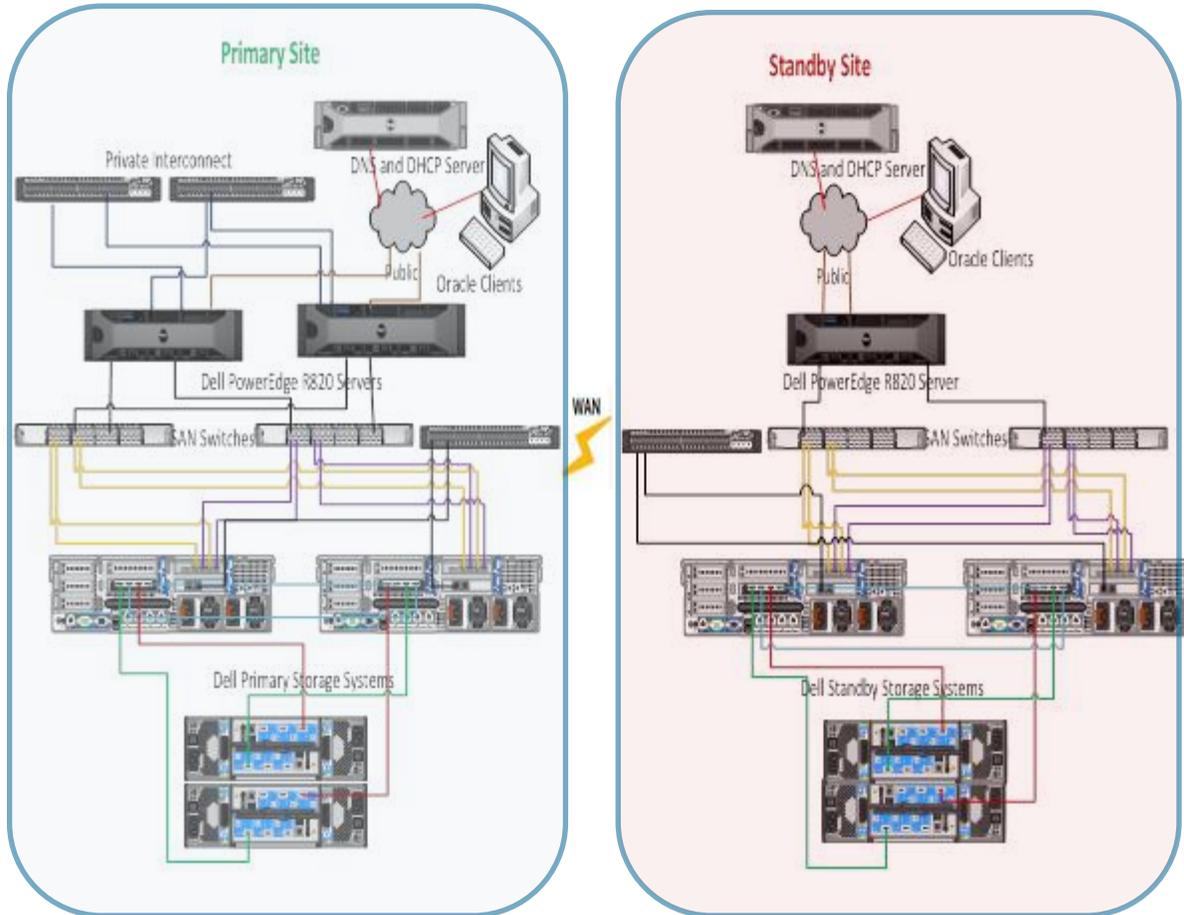
Table 4. Database volume configuration

Volumes	File Type
DATA	Oracle database data files
REDO1	Redo log files and control files
REDO2	Multiplexed redo logs and control files
FRA	Flashback logs and Archived Logs

7. An Oracle 11gR2 Grid Infrastructure with only the RAC binaries are deployed with a secondary storage as the standby site to perform the replication of data.
8. After creating the Consistency Group, a snapshot schedule is created from each of the Dell storage for auto generation of the storage snapshots. These storage snapshots are used during our test cases.

The figure 7 provides the detailed architecture used for this study.

Figure 7. Test environment



5 Test Scenarios

Several test scenarios have been executed on all the Dell storage types to perform disaster recovery across sites and validate the database on the standby sites.

5.1 Dell Compellent Storage

In Dell Compellent storage, we tested use cases by simulating either a production disaster or a link breakdown between the primary and the standby sites. We verified the failover scenarios by mounting the Oracle database on the standby site and then checked for consistency.

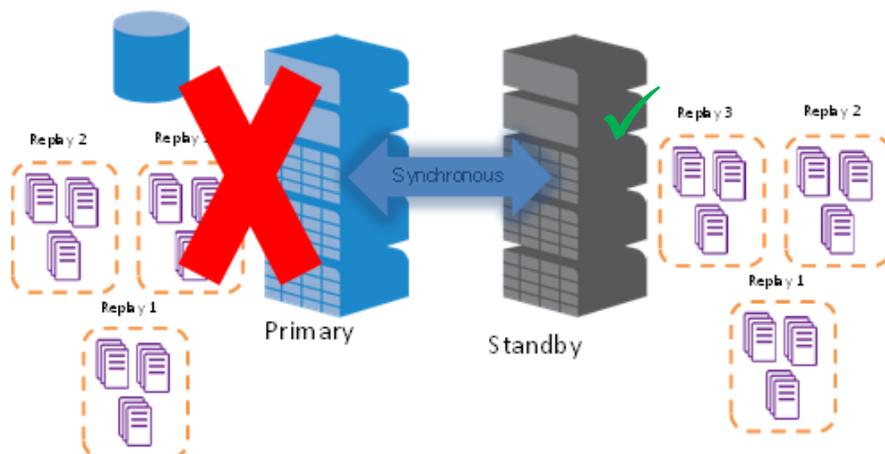
5.1.1 Asynchronous Replication Disaster Recovery Simulation

Asynchronous replication copies replays from the source volume to the destination volume after they are frozen. Copying replays is scheduled for a specific time.

1. Create a consistency group for each of the disk groups associated with the database.
2. Create a replay schedule for a specific duration depending on the bandwidth available.
3. From the Compellent Enterprise Manager using [Wiki](#), set up the Asynchronous replication from the primary storage center to the standby storage center for the consistency group database volumes.
4. Place the database in hot backup using [Appendix B](#) before the replay and after the replay is completed, remove the database from the hot backup mode. The replay is copied to the standby storage center.
5. Simulate the disaster by shutting down the primary site.
6. Create volumes from the replays on the standby storage center and map the volumes to the standby host.

The database was consistent till the point of the replay.

Figure 8. Asynchronous Replication between storage arrays



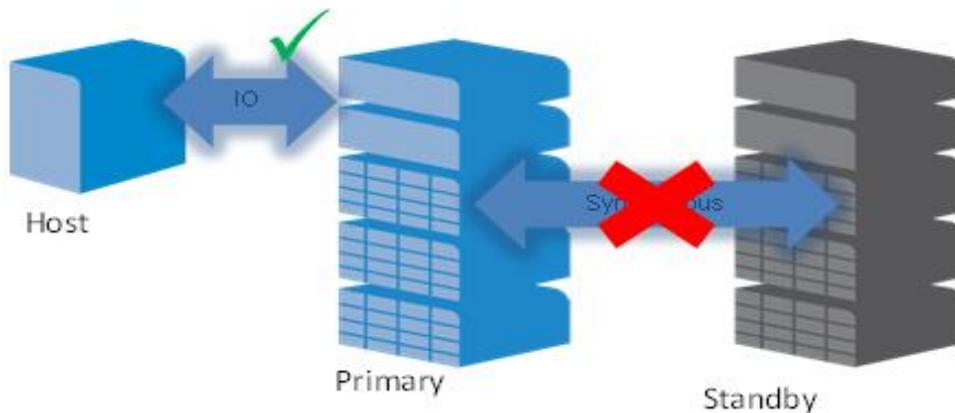
5.1.2 Synchronous High Availability Replication Disaster Recovery Simulation

This solution provides zero data loss for single failure scenarios. Writes from the application are written to the primary volume first and acknowledged back to the application, and then the changes are synchronized to the standby site.

Synchronous replication accepts writes to the source volume even if the destination volume is unavailable (or when latency is too high) to avoid interrupting service. However, if writes are accepted to the source volume, the destination volume data becomes stale.

1. Create consistency groups for each of the disk groups associated with the database.
2. From the Compellent Enterprise Manager using [Wiki](#), set up the synchronous high availability replication between the primary storage center and the standby storage center.
3. Simulate a link failure between the primary and the secondary sites and perform some changes on the primary site. The primary storage center journals all the changes and creates an active replay.
4. Establish the link and the active replay is copied to the standby site.
5. Verify the changes on the standby site by enabling the standby site [Appendix A](#)

Figure 9. Link Failure in High Availability Between Storage Arrays



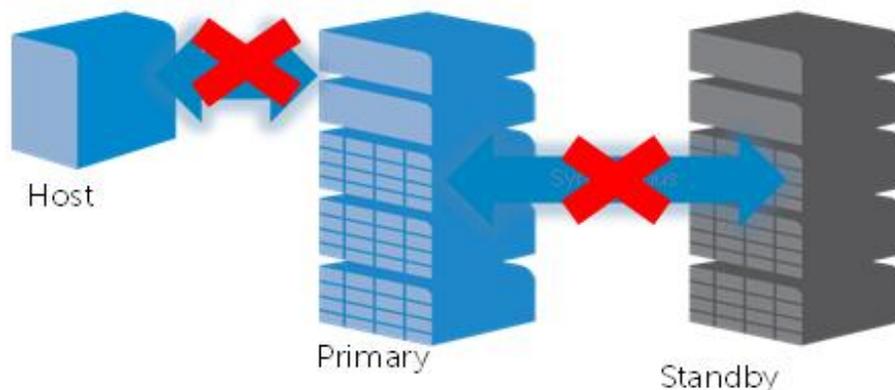
5.1.3 Synchronous High Consistency Replication Disaster Recovery Simulation

This solution provides zero data loss for double failure scenarios. In High Consistency synchronous replication mode, writes from the application are written to both the primary and the secondary volume at the same time before the acknowledgment is sent back to the application.

Writes from the application are failed if either the primary volume or replica volume is not available. All writes are failed to ensure that the primary volume and the replica volume are always identical.

1. Create Consistency Group for each of the disk groups associated with the database.
2. From the Compellent Enterprise Manager using [wiki](#), Set up the synchronous High Consistency replication between the primary storage center and the standby storage center Simulate a disaster at the primary site and mount the standby site to verify the data replication is consistent.

Figure 10. Link Failure in Synchronous High Consistency Between Storage Arrays



5.2 PowerVault Storage

In PowerVault storage, we tested use cases by simulating both a primary site disaster and a communication link breakdown between the primary and secondary storage sites. We verified the failover scenarios by mounting the Oracle database on the standby site and then checked the consistency.

5.2.1 Asynchronous Standard Remote Replication Disaster Recovery Simulation

Asynchronous standard remote replication is intended to operate in configurations with relatively high-latency and/or low-bandwidth communication links. This method decouples primary-side virtual disk performance from link performance, and therefore enables the use of lower cost links.

1. Create Remote Replication Group which is associated with both the local and remote storage.

2. Create Remote Replicated Pair for each database virtual disk on the primary storage within the existing remote replication group. Complete the similar associated tasks on the remote storage array using the complete Replicated Pair.
3. Place the database in backup mode before creating the remote replicated pair and once it is created, take the database out of backup mode.
4. The initial full synchronization occurs between the local storage array and the remote storage array when the replicated pair is created and the consecutive point-in-time images at regular intervals.
5. Set up the Asynchronous Standard Remote Replication between the primary storage array and the secondary storage array from the PowerVault MD Storage Manager (Enterprise Management) using [Wiki](#).
6. Simulate the communication link failure between the primary and secondary storage array. Some transactions are done at the primary site and primary remote replication group starts the re-synchronization process immediately after communication has been restored.
7. Simulate the failover of the primary site to the secondary site by changing the replication group role to secondary. The users are connected to the secondary site.

5.2.2 Synchronous Write Mode Disaster Recovery Simulation

This mode guarantees that the remote virtual disks are identical to the primary side virtual disks with each and every write and provides the highest level of data protection.

In synchronous write mode, if the primary array suffers a catastrophic failure or disaster, application data on the secondary array can be recovered, and only the write requests that were outstanding during the failure causes content discrepancies. The use of a synchronous remote replication allows applications to be made disaster-tolerant through remote replication.

1. Create a virtual disk in the secondary storage array that is greater than or equal to the capacity of the primary virtual disk to be in the replicated pair.
2. Create a replicated virtual disk pair between the primary virtual disk and secondary virtual disk and then create the replication relationship between the two virtual disks.
3. Set up the Synchronous Legacy Remote Replication between the primary storage array and the secondary storage array from the PowerVault MD Storage Manager (Enterprise Management) using [Wiki](#).
4. Simulate a link interruption between local primary and remote secondary RAID controller module when the I/O is occurring on the primary virtual disk.
5. The replicated pair changes to unsynchronized status and sends an I/O completion to the primary host.

6. After the connectivity is restored, the data blocks changed during the link interruption are copied to the secondary virtual disk and the replicated pair changes from unsynchronized state to synchronization in progress status.

5.2.3 Asynchronous Write Mode Disaster Recovery Simulation

This mode accepts writes to the primary virtual disk if the secondary virtual disk is unavailable and the performance of the primary-side application is not gated by communication link performance.

Asynchronous write mode removes the restriction that the remote write has to complete along with the local primary write but this mode does not guarantee that data was successfully written to the secondary virtual disk before indicating a successful write to the host system.

1. Create a virtual disk in the secondary storage array that is greater than or equal to the capacity of the primary virtual disk to be in the replicated pair.
2. Create a replicated virtual disk pair for legacy remote replication between primary virtual disk and secondary virtual disk by selecting the asynchronous write mode for the remote replication.
3. Set up the asynchronous legacy remote replication between the primary storage array and the secondary storage array from the PowerVault MD Storage Manager using [Wiki](#).
4. Simulate a disaster at the primary storage site and test the failover to secondary storage site and verify the data on the secondary site.

5.2.4 Asynchronous Write Order Consistency Disaster Recovery Simulation

This mode guarantees the remote secondary storage site is recoverable in case of any disaster occurs on the primary storage site.

In asynchronous write order consistency mode ensures writes to the secondary storage array are completed in the same order as on the primary storage array by enabling add to write consistency mode group. If one remote replication in a write consistency group becomes unsynchronized, all of the remote replications in the group become unsynchronized.

1. Create a virtual disk in the secondary storage array that is greater than or equal to the capacity of the primary virtual disk to be in the replicated pair.
2. Create a write consistency group and add virtual disks to the consistency group to guarantee the order of write requests on the secondary virtual disk.
3. Create a replicated virtual disk pair for legacy remote replication between primary virtual disk and secondary virtual disk.
4. Setup the Asynchronous Write Order Consistency Legacy Remote Replication between the primary storage array and the secondary storage array from the PowerVault MD Storage Manager (Enterprise Management) using [Wiki](#).
5. Simulate a disaster at the local storage site and test the failover to remote storage site and verify the data consistency on the remote site.

Conclusion

Dell offers multiple Disaster Recovery solutions based on different storage technologies to reduce the RTO and RPO of business critical applications and databases. This paper provides guidance to setup a complete disaster recovery solution for Oracle Databases involving both Synchronous and Asynchronous replication modes.

These solutions can help to achieve the following benefits

- Complete/Point-in-time database recovery
- Enables Disaster Recovery for Standard Edition Oracle Database
- Reduces licensing cost
- Minimal host resource utilization to enable DR solution

References

[Disaster Recovery of Oracle Database Standard Edition \(SE\) and Standard Edition One \(SE1\) on Dell EqualLogic PS Series iSCSI Storage Systems](#)

[Asynchronous Remote Replication Feature on Dell PowerVault MD3 Storage Arrays](#)

[Dell PowerVault MD3600f/MD3620f Design Guide for Disaster Recovery Using Remote Replication](#)

[Dell PowerVault MD3600f/MD3620f Remote Replication Functional Guide](#)

[Dell Compellent Storage Center 6.3](#)

For Command line Utility refer the below

[Dell PowerVault Modular Disk Storage Arrays CLI Guide](#)

[Dell Compellent CompCU](#)

Appendix A

Enable Standby Database

- Create Volumes from the replay and map to the standby servers <Wiki4>
- Rescan the disks and mount the ASM disk groups using ASMCA or sqlplus.
- Make sure that the Grid Infrastructure is installed and configured, and the oracle home binaries are installed in the target node.
- Create Pfile, Password file, adump directories on the standby host. These files can be copied from the production instances also.
- Modify the scsi id's in the multipath.conf file with the scsi id's of the discovered luns and restart the multipathd service
- Mount all the diskgroups using asmca
- Manually change the oracle binary ownership and permission. As root user:

```
cd $ORACLE_HOME/bin
chgrp asmadmin oracle
chmod 6751 oracle
ls -l oracle
```
- Ensure the ownership and permission is now like:

```
-rwsr-s--x 1 oratest asmadmin
```
- Start the database using

```
$ srvctl start database -d dupdb
```

Appendix B

Creating a storage snapshot backup

1. Log into your Oracle database server using sqlplus and issue the following commands.

```
SQL > ALTER SYSTEM ARCHIVE LOG CURRENT;
```

```
SQL > ALTER DATABASE BEGIN BACKUP;
```

2. Create replays from the Console or command line.

3. When the replay has been taken take the database out of backup mode using following commands.

```
SQL > ALTER DATABASE END BACKUP;
```

```
SQL > ALTER SYSTEM ARCHIVE LOG CURRENT;
```